

DECODING CULTURAL CONSUMPTION FROM DIGITAL TRACES

A Dissertation

Presented to the Faculty of the Graduate School
of Cornell University

in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by

Minsu Park

August 2019

© 2019 Minsu Park
ALL RIGHTS RESERVED

DECODING CULTURAL CONSUMPTION FROM DIGITAL TRACES

Minsu Park, Ph.D.

Cornell University 2019

With an unprecedented scale of users interacting with online platforms and mobile devices, *digital traces* (i.e., detailed behavioral data generated from those sources) provide us with an unparalleled opportunity to explore new scientific approaches that enable novel insights about the patterns of cultural consumption, which has an impact on social and social-psychological outcomes. Through discussion of three projects, I show that by leveraging the large-scale digital traces (i) individuals manage mood through self-exposure to external stimuli such as music, e.g., people listen to more relaxing music late at night and more energetic music during normal business hours, including mid-afternoon when affective expression is lowest, (ii) cross-cultural convergence is more advanced in cosmopolitan countries with cultural values that favor individualism and power inequality, and (iii) the diversity of musical tastes are associated with whether one is following high-profile news media and how much one is 'into' music rather than socioeconomic covariates such as income and education.

BIOGRAPHICAL SKETCH

Minsu Park completed his Bachelor of Science in Mechanical and System Design Engineering at Hongik University. His interest diverged to Computational Social Science and Social Computing with his passion in understanding the social and social-psychological mechanisms that underlie how people construct meaning and consequently pursue action. He finished his Master of Science in Culture Technology from KAIST by studying behavioral and perceptual differences between the depressed and non-depressed individuals on social media. He continued his journey with a particular focus on cultural consumption which has an impact on social and social-psychological outcomes at Cornell University. He currently develops and applies quantitative and computational techniques to collect, repurpose, and analyze data from diverse sources, including mobile devices (*small data*; e.g., granular-level of mobility and physiological logs), online platforms (*big data*; e.g., large-scale text and behavioral data such as messages and interactions from online social media), and online experiments to study the dynamics of cultural consumption.

To Dami

ACKNOWLEDGEMENTS

I could not have completed my PhD without the unwavering support of people, to whom I am immensely grateful. Academically, my biggest debt is owed to my advisor and mentor, Michael Macy. No other person has had such a massive influence on my development as an independent researcher. His constructive but demanding critique of my work has pushed me to think clearly about my research. Whenever I felt I was hitting a wall, he patiently guided me to find my own solutions. He always supported me to pursue my own research ideas and shape my own research program. Michael also provided thoughtful advice and incredible care when I was at critical and difficult junctures in my life. I especially appreciate that he afforded me endless supports to let me be with my family in NYC throughout the years.

I would also like to thank Mor Naaman who was my initial advisor and then became a committee member and has since been an exceptionally supportive mentor to me. He never stopped being generous and kind with professional and social supports as well as his feedback on my work throughout my entire years at Cornell. Because of him, in particular, I could have chances to work with great people outside of academia through multiple internships and could meet other researchers with common research interests outside of Cornell.

Other mentors have each provided their own indispensable forms of support. Drew Margolin pushed me to prepare my work with values beyond the new methodological approach and data. While working on his qualifying exam questions, I could better structure my research and knowledge based on the theoretical implications of my work and its connections to the literature. Ingmar Weber invited me to Qatar Computing Research Institute as a research associate right before I started my PhD. Since then he has been one of my role models in

the academic community. I deeply admire his contagious passion for research and incredible productivity that energized and encouraged me to complete this dissertation and more. I spent terrific summers with internships, during which I was privileged to work with colleagues and mentors who profoundly influenced my thinking and research. In particular, I would like to thank Michael Gilbert, Julia Haines, and Elizabeth Churchill at Google and Jennifer Thom and Henriette Cramer at Spotify.

I had the opportunity of working with and alongside many bright and inspiring colleagues in Information Science, Michael's Social Dynamics Lab, and Mor's Social Technologies Lab. Beyond our research, they have uniquely enriched my experience at Cornell in many ways, including endless couch-surfing. I also owe a special thank-you to our outstanding department administrators.

Finally, there is Dami, to whom I can never say thank you enough. I decided to pursue a PhD to move to NYC for her study in fashion design—yet, the decision would never have materialized without the support from my family, so that I owe a huge debt to them as well. Regardless of the initial motivation, I could quickly realize that I like to do research and there are particular topics I genuinely love to explore for a lifetime. I thank her for bringing me such a life-change opportunity. I would also never be able to complete my PhD without her encouragement and dedication, especially after we had our demanding explorer, Uri. Because of her, I could move forward. This dissertation is dedicated to her.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
List of Tables	ix
List of Figures	x
1 Introduction	1
1.1 Motivation	1
1.2 Overview	3
2 Global music streaming data reveal diurnal and seasonal patterns of affective preference	4
2.1 Introduction	5
2.2 Results	8
2.3 Conclusion	14
2.4 Methods	16
2.4.1 Dataset Description	16
2.4.2 Chronotypes	17
2.4.3 Completed Plays	18
2.4.4 Measures of Within- and Between-individual Affective Preferences	19
2.4.5 Seasonal Variation	21
2.4.6 Group Baseline Comparisons	22
2.4.7 Other Psychological Features in Music Attributes	23
2.4.8 Effects of Daylight Saving Time	24
3 Cultural values and cross-cultural video consumption on YouTube	26
3.1 Introduction	26
3.2 Cultural Values	28
3.2.1 Individualism-collectivism (IDV)	29
3.2.2 Uncertainty Avoidance Index (UAI)	30
3.2.3 Power Distance Index (PDI)	30
3.2.4 Masculinity (MAS)	32
3.3 Materials and Methods	32
3.3.1 YouTube Data Collection	32
3.3.2 Bipartite Co-consumption Network	33
3.3.3 Measure of Cultural Openness	35
3.3.4 Economic, Linguistic, and Technological Measures	37
3.4 Results	38
3.5 Discussion	43
3.6 Conclusion	45

4	Understanding Diversity of Musical Tastes via Online Social Media	46
4.1	Introduction	46
4.2	Related Literature	49
4.2.1	Changing Status of the Omnivore Thesis	50
4.2.2	Technology and Music Listening Practice	52
4.2.3	Research Questions	53
4.3	Method	54
4.3.1	Initial Data Collection	55
4.3.2	Socioeconomic Status	55
4.3.3	Genre and Subgenre Information Collection	58
4.3.4	Measuring Diversity	59
4.3.5	Into-ness and Openness	60
4.3.6	Data Validation and Preparation	62
4.4	Results	65
4.4.1	Diversity Measure	65
4.4.2	Correlates of Musical Diversity	67
4.5	Discussion	70
4.6	Final Remarks	72
5	Conclusions	74
5.1	Mood Management through Cultural Consumption	74
5.2	Diurnal and Seasonal Patterns of Variety-seeking	75
5.3	Diffusion of Cultural Interests	76
5.4	Effects of Collective Traces and Recommendations	77
5.5	Data Science Tools for Large-scale and High-dimensional Observational Data	77
5.6	Inference of Sociodemographic Attributes of Online Users	78
5.7	Final Thoughts	79
A	Supplementary Information of Chapter 2	
	(Global music streaming data reveal diurnal and seasonal patterns of affective preference)	81
B	Supplementary Information of Chapter 3	
	(Cultural values and cross-cultural video consumption on YouTube)	88
	Bibliography	89

LIST OF TABLES

3.1	Descriptive statistics of network structure by video category. . .	39
3.2	OLS regression model of cultural betweenness among 58 countries.	40
3.3	OLS regression model of cultural closeness among 58 countries. .	41
4.1	Fifteen variables used to explain the measured musical diversity scores and genre- and subgenre-level of diversity scores. The distributions accompanying each variable begin at zero and end at the adjacent maximum. Many variables are not normally distributed.	63
4.2	Multiple regression coefficients of individual factors on the musical diversity of genre and subgenre.	68
A.1	Detailed statistics related to baseline comparisons in musical intensity. We report additional information including exact probability values, degrees of freedom, confidence intervals, and effect sizes for baseline comparisons between groups to support our results related to the diurnal and seasonal patterns of affective preference in musical intensity based on temporal music consumption of one million Spotify users over a year. We performed all the tests using Welchs two-sample t-test (two-sided) to correct for unequal size and variance between paired samples.	81
A.2	Proportion of sample in each demographic group from each of the 51 countries based on the World Factbook. Countries are ordered by the proportion of sampled users, which reflects the country’s relative demographic distributions compared to the world population distribution, not Spotify’s user distribution over the globe. Although the population distribution in the World Factbook breaks populations under age of 25 in age 0–14 and age 15–24, we merged the two age groups (0–14 and 15–24) into one (13–24) as users need to be 13 or older to sign up for Spotify.	82
A.3	Country names included and excluded for analyses. Countries collected from YouTube in this study.	83
B.1	Country names included and excluded for analyses.	88

LIST OF FIGURES

2.1	Millions of global music plays reveal diurnal affective patterns.	9
2.2	Diurnal affective patterns are robust across diverse geographic regions, demographic groups and chronotypes.	10
2.3	Affective preference is associated with seasonal variation in day length.	13
3.1	Construction of the bipartite network of video co-consumption on YouTube.	35
4.1	Multidimensional scaling for distance between genres.	61
A.1	Seasonal variations in affective preferences revealed by music streaming exhibit robust patterns across different geographic regions and user groups.	84
A.2	Affective preferences revealed through music selection varies from month to month, but daily differences are seasonally robust.	85
A.3	The hourly distribution of plays from global music streaming data.	85
A.4	Principal component analysis (PCA) of 11 musical attributes identified a first principal component corresponding to musical intensity.	86
A.5	Regression discontinuity analysis reveals no impact of daylight saving time (DST) transitions on musical intensity.	87

CHAPTER 1

INTRODUCTION

1.1 Motivation

With advances in mobile technology and Internet access, more people are consuming and engaging in culture than ever before. In Google searches for the year of 2016, search queries for cultural activities including music, movie, and sports eclipsed even those of news and politics [49]. Central to this shift is the relocation of cultural consumption channels from physical spaces to online spheres, whether from movie theatres or DVD-ROMs to on-demand streaming services; from the video game arcade to disc-less console downloads; from the book to the e-book; and from the operator, ticket booth, or community bulletin board to on-line platforms (from booking a restaurant and concert to even forming physical activity groups).

With this change (where digital technologies have become embedded in our day-to-day lives), examining cultural consumption through digital traces is relevant from several perspectives. For example, the way people access cultural products may affect their relationships to the product and its content: whereas in prior decades certain cultures might be polarized into hierarchies of “upper” and “lower-class,” the growth of the Internet has made culture less class-stratified [45, 96]. Alongside consumption activities themselves are new forms of information sharing and collaboration between consumers: new styles and trends are recounted in social media, and consumers often become producers of new cultural content. Do these changes affect our understandings of culture, how it should be presented and consumed, and its relationships with other

types of culture? All of which are new but the change offers useful sources to gather data that can provide valuable research opportunities.

This change has transformed not only the ways we participate in cultural activities but also the ways in which scientists understand and analyze cultural consumption behaviors. Digital traces of our everyday consumption activities are time-stamped recordings of *when* and *with whom* each individual consumed what specific cultural products. With the introduction of new tools and methodologies, large-scale digital traces enable researchers to explore microscopic behaviors at macroscopic scales [47]. With an unprecedented breadth and depth and scale [72], it also can offer a more complete picture of cultural consumption behaviors beyond what has been learned from previous studies, which have largely relied on small homogeneous samples observed in the artificial setting of the lab or indirect retrospective accounts obtained through survey responses [47, 51, 76]. These pictures of cultural consumption may provide critical insights into our understanding of broader moral, social, and cultural values that drive societies [131].

Yet, leveraging large-scale digital traces for understanding cultural consumption behaviors faces several challenges. (a) Many domain knowledge in the behavioral and social sciences is based on qualitative measures. The challenge is how to computationally operationalize this knowledge so that it is amenable to quantitative analysis. (b) Raw data (including activity, interaction, and product information) is massive but typically does not directly measure a specific consumption behavior (and its associated factors, e.g., socioeconomic status) and specific product characteristics (e.g., arousal level of a song). New advanced computational techniques are required to infer a measure that accu-

rately reflects a theoretical construct from raw data, often from heterogeneous data sources. (c) Digital traces are primarily observational. To produce reliable scientific results, extensive robustness checks or additional analysis (e.g., causal analysis beyond correlation) might be required. This dissertation describes our attempts in addressing these challenges while addressing some important theoretical questions.

1.2 Overview

In this dissertation, we present novel computational methods to derive new insights from digital traces that can help us better understand cultural consumption behaviors and produce theoretical and practical implications. We consider three key aspects of cultural consumption behaviors: cultural consumption (1) for achieving a certain psychological end and the impact of (2) social hierarchies and (3) cultural values on cultural preferences and tastes. In Chapter 2, we leverage worldwide Spotify data to study diurnal and seasonal patterns of affective preference across 51 countries, revealing a previously unknown dynamics of human emotion. In Chapter 3, we revisit the omnivore thesis to investigate the relationships between socioeconomic status and construction of cultural tastes using Last.fm and Twitter data. In Chapter 4, we demonstrate that cosmopolitan culture is integrated with a specific national culture (e.g., cultural values that favor power inequality and tolerance for uncertainty) as a response to globalized cycles of production and consumption. The main contributions are summarized in each chapter, and then high-level conclusion and future directions are summarized in the last chapter.

CHAPTER 2

GLOBAL MUSIC STREAMING DATA REVEAL DIURNAL AND SEASONAL PATTERNS OF AFFECTIVE PREFERENCE

Originally published in Nature Human Behaviour (2019) [95]

Abstract People manage emotions to cope with life's demands. Previous research has identified affective patterns using self-reports and text analysis, but these measures track the expression of affect, not affective preference for external stimuli such as music, which affects mood states and levels of emotional arousal. We analysed a dataset of 765 million online music plays streamed by 1 million individuals in 51 countries to measure diurnal and seasonal patterns of affective preference. Findings reveal similar diurnal patterns across cultures and demographic groups. Individuals listen to more relaxing music late at night and more energetic music during normal business hours, including mid-afternoon when affective expression is lowest. However, there were differences in baselines: younger people listen to more intense music; compared with other regions, music played in Latin America is more arousing, while music in Asia is more relaxing; and compared with other chronotypes, 'night owls' (people who are habitually active or wakeful at night) listen to less-intense music. Seasonal patterns vary with distance from the equator and between Northern and Southern hemispheres and are more strongly correlated with absolute day length than with changes in day length. Taken together with previous findings on affective expression in text, these results suggest that musical choice both shapes and reflects mood.

2.1 Introduction

Individuals manage mood to function productively and cope with the demands of daily routines [52, 125]. The way in which a person chooses to regulate their mood has consequences for mental health, interpersonal functioning and personal well-being [53]; social networking, exercise and meditation generally have positive consequences, while cigarettes, drugs and alcohol can be detrimental [142]. People may also choose to regulate their mood through media consumption, including movies, TV, books and music. Among these media, music is unique in predating recorded history as a universal component of human life [28, 117], one that both reflects and alters levels of emotional arousal [125, 57, 67], energy, wakefulness [124] and tension [125, 67]. Music is also uniquely omnipresent, serving as a background soundtrack to both leisure and work activities [88], with reported listening time averaging up to 44% of waking hours [51]. While consumption of other media may also be useful for understanding emotion management, the omnipresence of music affords a singular opportunity to identify diurnal and seasonal patterns in listener's musical choices, at a very high level of temporal granularity and across diverse cultures and demographic groups.

Previous research on music consumption has relied largely on self-reports, surveys and laboratory experiments, with severely restricted numbers of participants, observation periods and geographic ranges, and without representative or naturalistic musical stimuli [51]. These limitations can now be overcome due to the rapidly growing use of mobile devices and music-streaming services worldwide. Almost half (45%) of Internet users aged 16-64 actively access licensed music throughout the day using streaming services¹⁵ on a variety of

devices, such as mobile phones, computers and smart speakers [102, 90, 25]. Of equal importance, detailed sonic and affective attributes are now available for millions of individual songs [51].

The growth of text-based social media has enabled a growing number of large-scale studies of global affect using text analysis. Recent studies used Twitter and Facebook data to take ‘the pulse of the nation’ [83], for cross-cultural comparisons of diurnal and seasonal patterns of positive and negative affective expressions [46], to measure affective responses to events [126] and track the consequences of shared emotionally salient news feed content [68].

Music listening differs from what people write in that it offers insight not only into what people may be feeling but also what they may want to feel. Put another way, people can choose which music to consume to achieve a desired mood (along, of course, with purposes unrelated to mood management, such as learning to sing or play the song). While previous studies of social media postings make it possible to track daily and seasonal patterns of affective expression, music consumption offers an unprecedented opportunity to identify global patterns of affective preference. Affective expression exposes others (the readers) to the writer’s emotional content; conversely, the choice of music is a ‘revealed preference’ [112] for exposure to emotional content created by others. In short, tracking the temporal patterns of affective preference can offer a more complete picture of the emotional rhythms in human behaviour, beyond what has been learned from previous studies of affective expression.

To that end, we report hourly, daily and seasonal patterns of affective preference based on musical choices on a global scale. This descriptive account does not attempt to answer important questions about the motivations that shape

listening behaviour, the emotional effects of music exposure or the latent cognitive strategies in mood management. Instead, we contribute an empirical foundation for future investigations by tracking the affective content of the music people choose to listen to, broken down by hour, day and month, and by user demographics and global locations.

Accordingly, we analysed hourly, daily and seasonal changes in affective preferences as revealed by the choice of online music streamed via Spotify around the clock across 51 countries. For each listener with at least 25 completed plays, we collected up to 1,000 completely played tracks (mean $M = 771.9$; $s.d. = 336.8$). The set of listeners comprised a stratified random sample of one million worldwide Spotify users, matching each country's age and gender distribution on Spotify with current data from the Central Intelligence Agency's The World Factbook [20]. This sample included a total of 765 million tracks played between 1 January and 31 December 2016. Completed plays measure active self-exposure to music, excluding any songs the user may have sampled and discarded (see 'Completed plays' in the Methods for more details).

Spotify offers a way to analyse each track using 11 highly correlated audio attributes: acousticness, danceability, duration, energy, instrumentalness, liveness, loudness, mode, speechiness, tempo and valence. Principal component analysis (PCA) identified a latent construct that accounts for 29.4% of the variance in the correlation matrix (see 'Musical intensity measured by audio features of a track' in the Methods for more details). This principal component corresponds to musical intensity, ranging from highly relaxing (acoustic, instrumental, ambient, and flat or low tempo) to highly energetic (strong beat, danceable, loud and bouncy).

Aggregate temporal patterns in music consumption confound within-individual diurnal rhythms with between-individual differences in the hours when individuals with different baseline preferences for musical intensity tend to listen to music. Accordingly, we removed between-individual differences by mean-centring each individual's intensity scores such that every individual has the same baseline affective preference. We then restored between-group differences (for example, when comparing men and women or days of the week) by adding the group mean as a constant to the scores of each individual group member (see 'Measures of within- and between-individual affective preferences' in the Methods for more details). Thus, the reported temporal dynamics reflect changes over time for a prototypical group member, while differences in the intercept reflect between-group differences in baseline intensity scores.

2.2 Results

Figure 2.1 reveals qualitatively identical patterns of affective preference for musical intensity on a global scale across days of the week. Musical intensity levels were highest between 08:00 and 20:00, and lowest around 03:00, with a 5-h rise (between 03:00 and 08:00) and a 7-h decline (between 20:00 and 03:00). Maximum intensity was sustained for 12h (from 08:00 to 20:00), while minimum intensity reversed quickly and lasted only about 1h (from 03:00 to 04:00 on weekdays and 04:00 to 05:00 on weekends). Although the timings of minimum and peak intensity were nearly identical for all 7d, the baseline intensity level was higher on Friday and Saturday than on other days, especially in the evening when weekend social activities are likely ($M = 0.879$ and 0.883 for Fri-

day and Saturday, compared with $0.820 < M < 0.852$ for other days; $P < 0.001$ for all pairwise comparisons; all tests for equal means throughout the paper use Welch’s t-test to correct for unequal size and variance between paired samples; see Supplementary Table A.1 for additional statistical details). The morning dip on Saturday and Sunday was delayed by 1h (from 03:28 to 04:28), suggesting that listeners may have been sleeping in.

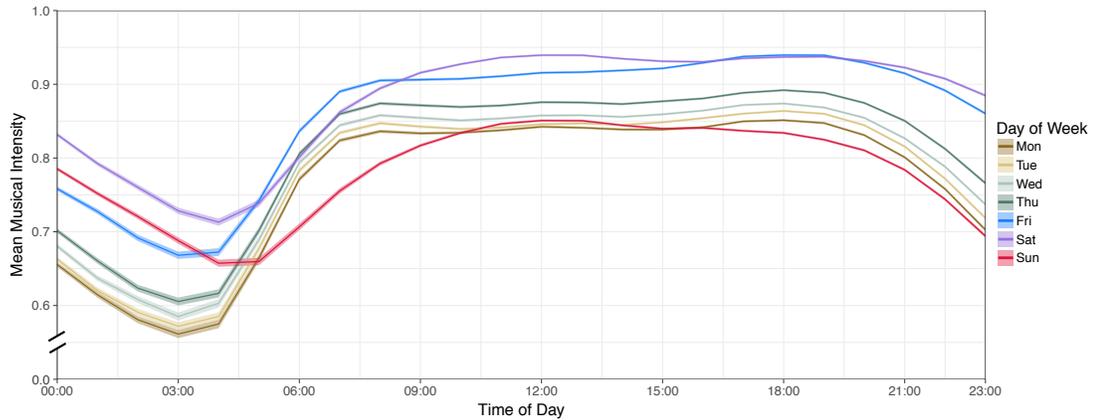


Figure 2.1: Millions of global music plays reveal diurnal affective patterns.

Overall, the diurnal pattern is remarkably similar to the temporal changes in positive affect reported in previous research using sentiment analysis of time-stamped Twitter messages⁴ to measure user’s affective expression. Nevertheless, we discovered one striking exception: people the world over continue to choose highly intense music throughout the day, despite the mid-afternoon slump that is registered by what they write on Twitter. The dynamic congruence with positive affect and non-congruence with negative affect suggest an intriguing hypothesis for future research: listeners select arousing music that matches their positive mood and offsets their negative mood.

Figure 2.2 shows that the diurnal pattern is highly consistent across five geographic regions—Latin America, North America, Europe, Oceania and Asia

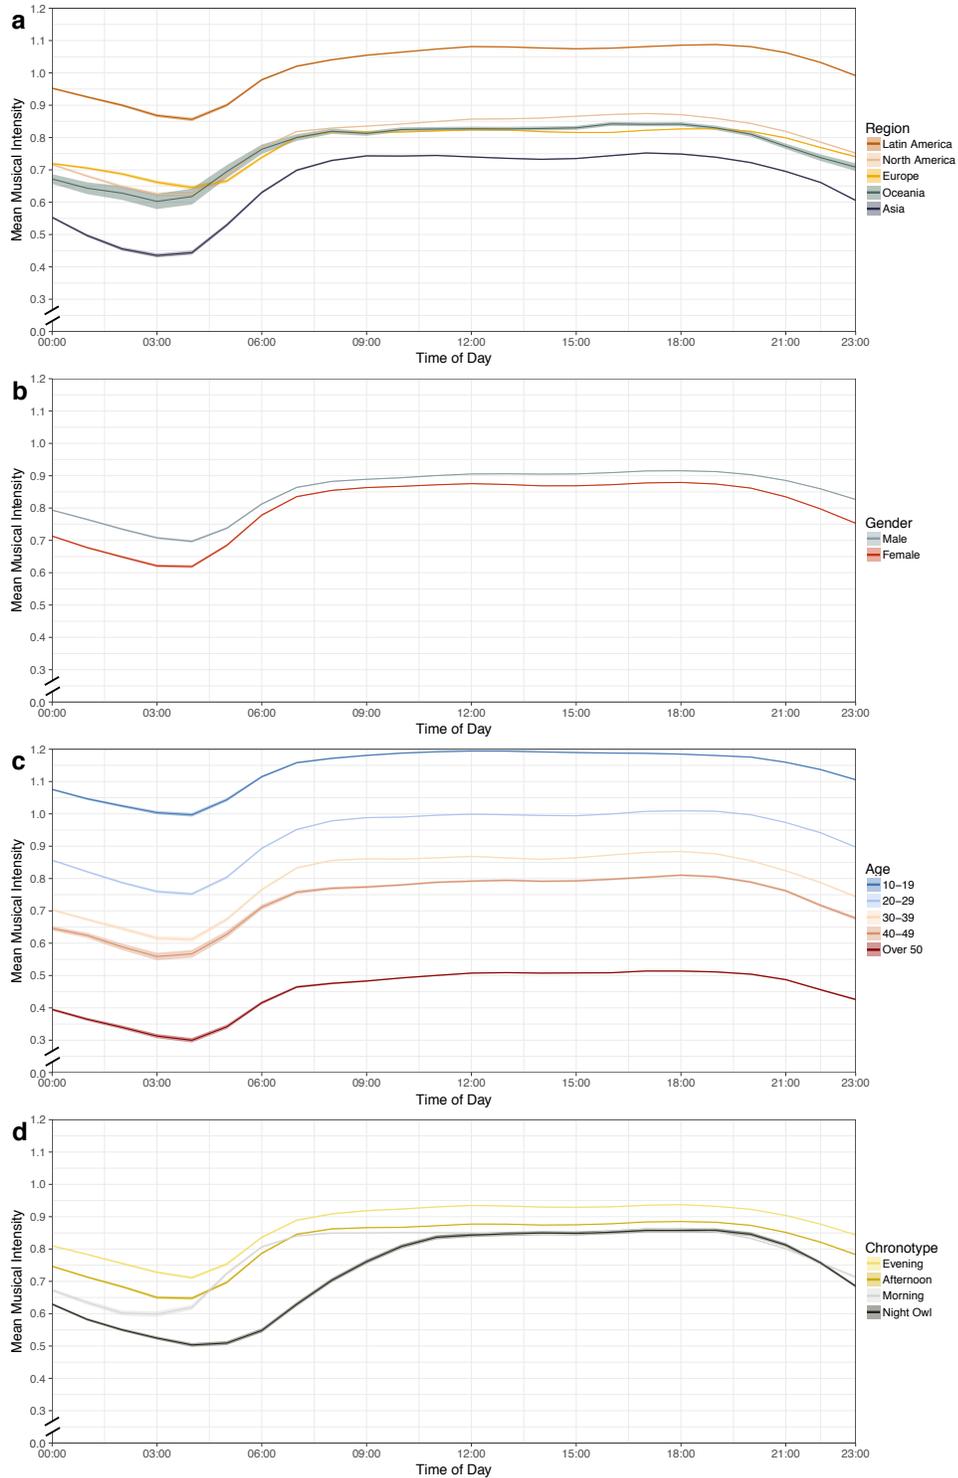


Figure 2.2: Diurnal affective patterns are robust across diverse geographic regions, demographic groups and chronotypes.

(Fig. 2.2a)—and across demographic groups based on gender (Fig. 2.2b), age (Fig. 2.2c) and chronotypes (Fig. 2.2d). Although the overall temporal pattern is highly robust, there are interesting between-group baseline differences. Music played in Latin America ($M = 1.053$) is relatively more intense, and music in Asia is more relaxed ($M = 0.698$) compared with Oceania ($M = 0.807$), Europe ($M = 0.804$) and North America ($M = 0.830$; $P < 0.001$ for eight pairwise comparisons of Latin America with the four other regions and Asia with the four other regions; see Supplementary Table A.1 for additional statistical details). This result corroborates and extends survey- and experiment-based findings that show cultural differences in affective preferences [128]. These studies suggest that there may be cultural differences in preferences for high-arousal positive affective states, such as excitement or enthusiasm, and low-arousal positive affective states, such as calm and peacefulness, between, for example, Western and East Asian cultures.

Across the globe, intensity scores also differ by age and gender. As people get older, they listen to less-intense music ($M = 1.162, 0.970, 0.841, 0.769$ and 0.484 , respectively, for the five age groups from 10–19 to over 50; $P < 0.001$ for all pairwise comparisons; see Supplementary Table A.1 for additional statistical details). Intensity scores were lower for music streamed by women ($D = -0.037$; $t = -26.04$; $d.f. = 1,033,792$; $P < 0.001$), especially in the evening. Curiously, however, this global gender difference masks large gender differences on opposite sides of the equator, as reported in Supplementary Fig. A.1a. Women stream music with higher intensity than men in the Southern Hemisphere ($D = 0.017$; $t = 6.50$; $d.f. = 262,409$; $P < 0.001$), while the pattern is the opposite in the Northern Hemisphere ($D = -0.054$; $t = -32.31$; $d.f. = 771,029$; $P < 0.001$).

The temporal dynamics are also similar across three of four chronotypes. Chronotypes were defined by when users are most actively listening, in six-hour increments beginning at midnight. The outlier group is the night owls whose baseline music intensity scores ($M = 0.684$) are lower than the scores for the other three chronotypes, with group averages increasing with the time of day during which users are most likely to listen ($M = 0.834$ for morning people, $M = 0.861$ for afternoon people and $M = 0.903$ for evening people; $P < 0.001$ for all pairwise comparisons; see Supplementary Table A.1 for additional statistical details). These diurnal patterns among chronotypes closely resemble the previous findings [46] based on affect words in Twitter messages, suggesting that music consumption is closely aligned with the emotions people express. However, there is an interesting difference with affective expression in the behaviour of night owls who tend to prefer more relaxing music overall, yet display a larger increase in musical intensity during the daytime ($D = 0.412$; $t = 239.66$; $d.f. = 2,648,000$; $P < 0.001$ for the comparison between 04:00 and 18:00) compared with the daytime increase for the other 3 chronotypes ($D = 0.280$; $t = 344.11$; $d.f. = 4,300,469$; $P < 0.001$). A possible explanation is that night owls may need stronger musical stimuli to stay alert during the day.

Figure 2.3 reports weekly and monthly changes in music consumption that suggest that people have seasonal music preferences [66, 100]. Previous research using self-reports found that listeners prefer highly arousing music during warmer months and serene music in colder seasons [100, 69], but these studies were based on self-reports from small samples in specific countries. Figure 2.3 confirms these results on a global scale, except during winter weeks when music listening is dominated by ceremonial holiday music for Christmas and Carnival. Intensity scores peak around the summer solstice ($D = 0.078$;

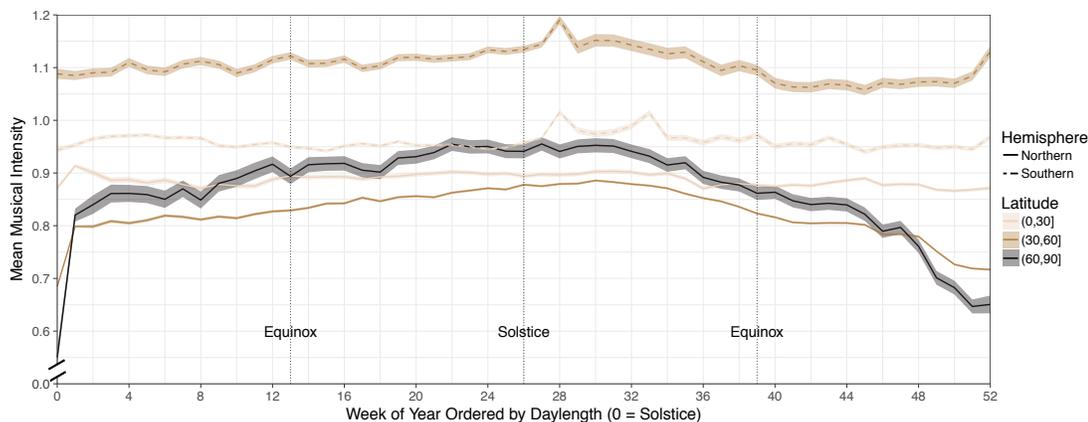


Figure 2.3: Affective preference is associated with seasonal variation in day length.

$t = 507.83$; $d.f. = 107,747,995$; $P < 0.001$ for the mean difference in intensity between summer weeks 24–28 and all other weeks). Intensity scores then decline with day length, but the seasonal variation decreases with proximity to the equator. Remarkably, music associated with late-December holidays is associated with a steep winter decline in intensity in the Northern Hemisphere and a sharp uptick in the Southern Hemisphere, suggesting that seasonal variation associated with holiday music can depend decisively on day length at the time of the holiday ($D = -0.049$; $t = -304.82$; $d.f. = 116,364,849$; $P < 0.001$ for winter weeks 48–0 compared with other seasons in the Northern Hemisphere; $D = 0.087$; $t = 109.51$; $d.f. = 2,347,689$; $P < 0.001$ for week 28 compared with other seasons in the Southern Hemisphere). The other summer uptick in the south at latitudes under 30S is Carnival on 7 February.

The results in Fig. 2.3 resemble the seasonal patterns reported in previous studies based on affective expression in global Twitter messages [58, 46]. However, while Golder and Macy [46] found that positive mood covaries with change in day length, not absolute day length, we found that absolute day

length (the interval between sunrise and sunset) is a better predictor of musical intensity ($r = 0.029$; $P < 0.001$) than change in day length ($r = -0.007$; $P < 0.001$; difference in the Pearson's correlations = 0.036; Steiger's $z = 743.585$; $P < 0.001$; $n = 764,992,760$). The same result holds when excluding holiday songs ($r = 0.014$; $P < 0.001$ for absolute day length; $r = -0.008$; $P < 0.001$ for change in day length; difference in the Pearson's correlations = 0.023; Steiger's $z = 464.790$; $P < 0.001$; $n = 752,692,716$). This indicates that seasonal variations in affective music choices are more strongly influenced by seasonal activities that depend on temperature, weather, and indoor and outdoor daylight than by seasonal changes in the timing of sleep relative to the dawn signal that synchronizes the circadian pacemaker (see 'Seasonal activities and choice of music' in the Supplementary Information for additional details). Longer days are also associated with warmer temperatures, with peak temperature often lagging behind the solstice (depending on the location relative to land, water and prevailing winds). Peak music intensity also lags behind the solstice, suggesting that the mechanism that drives musical preference may be the activities associated with temperature as well as daylight.

2.3 Conclusion

In conclusion, data from on-demand music streaming now make it possible to study music consumption across highly diverse cultures, including countries whose music consumption is rarely studied. The findings reveal diurnal and seasonal patterns of affective preference that are highly robust across different user groups as well as countries that differ both geographically and culturally. Additional robustness tests are reported in the Supplementary Information, in-

cluding seasonal patterns by different user groups (Supplementary Fig. A.1) and diurnal patterns broken down by day of the year (Supplementary Fig. A.2).

Although the robustness of the results is encouraging, there are important limitations. First, despite the reliance on a stratified random sample that reflects local census distributions of age and gender, the sample is potentially biased towards individuals who have access to streaming services and devices, particularly in lower-income countries. Second, the data are observational, and without randomized trial experiments, temporal patterns of musical intensity cannot directly test whether and when listeners use music to reflect rather than influence their mood. The relative importance of mood management and mood expression is likely to depend heavily on the cultural activities to which music provides an accompaniment, such as parties and holiday rituals.

In addition, we have data only on the intensity level of the music people choose to consume, not the affective states of the listeners. We were therefore limited to comparisons with affective expression among a different set of users on a different platform and during an earlier time period. Nevertheless, our diurnal and seasonal results show a remarkable similarity to results based on sentiment analysis of Twitter messages [46]. There are differences as well. Positive emotion in Twitter messages dips around 15:00 while the consumption of arousing music does not, suggesting that music can also be used as a mid-afternoon stimulant. While diurnal mood patterns on Twitter point to the sleep cycle as the synchronizing mechanism, listening behaviour suggests that temporal variations in preferences for affective stimuli through music may be more closely aligned with the temporal organization of daytime and night-time activities. For example, we found that listeners across the globe prefer quiet, low-intensity, re-

laxing music late at night and high-intensity, energetic music with a strong beat throughout the day, including late afternoon when affect expressed in writing is depressed. The comparisons suggest the possibility that music choices may reflect the emotional rhythms of daily and seasonal activities to which music contributes by shaping as well as expressing mood.

2.4 Methods

2.4.1 Dataset Description

This study uses redacted retrospective data collected between 1 January and 31 December 2016 from music-streaming instances at Spotify—a popular streaming service for music, podcasts and video. Spotify provides 11 sonic and mood attributes (for example, acousticness, loudness, valence and energy), available through their API (<https://beta.developer.spotify.com/documentation/web-api/reference/tracks/get-audio-features/>). We obtained data for 764,992,760 streams from a stratified random sample of 991,035 users across 51 countries. The sample excludes users with fewer than 25 plays and was stratified to match each countrys age and gender distributions and population size, based on current data from Central Intelligence Agencys The World Factbook [20]. The sample excludes countries where Spotify is unavailable, or with too few users after sampling to measure cross-cultural patterns. This stratified sampling adjusts the sampling frame to reflect the population distribution, since the distribution of Spotify users does not necessarily reflect the underlying population distribution. As a result, the stratified sample represents world population distribution,

not Spotify user distributions over the globe. The mean age of this sample (not the service) was 37.1 years (*median* = 29 years; *s.d.* = 23.9 years) and 49.2% were female. Demographic distributions for each country can be found in Supplementary Table A.2. A user's geo-location (for example, city, country, region and continent) was assigned based on the most commonly occurring geo-grid—one-tenth decimal degree by one-tenth decimal degree of pairwise latitude and longitude (approximately 100km²)—based on Internet Protocol address. Using the Python pytzwhere library, the geo-grids were matched with time zones to normalize all time stamps to local time and adjust for daylight saving time (DST). Age and gender were obtained from current Spotify user profiles.

2.4.2 Chronotypes

Following Golder and Macy [46], users were allocated to four six-hour chronotypes based on the time when the user was most active on Spotify, beginning at midnight. Some 15.1% were morning people (06:00 to 12:00); 44.8% were afternoon people (12:00 to 18:00); 35.1% were evening people (18:00 to 00:00); and 5.0% were night owls (00:00 to 06:00). These chronotypes are similarly distributed across gender and age. The baseline intensity of music played by night owls differs from the other three chronotypes, as reported in Fig. 2.2d (see also Supplementary Fig. A.3 for the distribution of plays across different times of day).

2.4.3 Completed Plays

In contrast with radio-like streaming services, Spotify is a user-driven on-demand service with a vast catalogue from which users search for and choose songs they want to listen to. Spotify reports that more than 80% of listening on Spotify in 2016 (when we collected the data) was initiated by user selection and not through algorithmic personalization [118]. Users can also exercise selection by choosing which songs to play to completion and which to skip. We limited the analysis to completed (or non-skipped) plays to focus on the music people actively choose to listen to, excluding what they choose to skip.

Musical intensity measured by audio features of a track. Music provides listeners with an affective experience through various musical features, ranging from song lyrics to the emotional attributes of audio features. Musicologists argue that audio features (particularly biopsychological cues, such as arousal) have better cross-cultural applicability without the language constraints of lyrics [6]. Spotify's track-specific audio attribute data are considered the gold standard in music information retrieval [3]. Spotify provides 11 common audio features: acousticness, danceability, duration, energy, instrumentalness, liveness, loudness, mode, speechiness, tempo and valence (see descriptions in Supplementary Table A.3). The attributes are highly correlated, and PCA identified a latent structure, with the first principal component unambiguously interpretable as a measure of intensity that explains 29.4% of the variance. We excluded the second principal component, which explained an additional 12.1% of the covariance but did not have a meaningful interpretation including shared characteristics related to known musical attribute dimensions that people usually perceive, such as arousal (similar to our intensity measure), valence and

depth [50], among others. Supplementary Fig. A.4 displays the locations of the 11 Spotify attributes on the PCA coordinate space for the first two principal components. Song-specific intensity scores range from -7.70 to 3.96 and are strongly associated with musical acousticness ($r = -0.765$), energy ($r = 0.867$) and valence (negative to positive emotion; $r = 0.643$; all of the Pearson correlations are significant at $P < 0.001$; $n = 13,578,157$). Factor loadings show that tracks with high intensity tend to be fast, loud, vocal (that is, not instrumental), happy, cheerful and euphoric (see Supplementary Table A.3 for the complete set of factor loadings).

2.4.4 Measures of Within- and Between-individual Affective Preferences

Temporal changes in affective preference were measured as the average intensity level of the music that a user listened to in each of the $247 = 168h$ of the week. Failure to disaggregate within- and between-individual affective preferences makes changes over time uninterpretable due to the confounding of individual diurnal rhythms and temporal changes in the composition of active users on Spotify. Between-individual variation in intensity scores (that is, the average level of intensity in the music that a user listened to) captures how individuals differ from one another in their baseline affective preferences, regardless of the time of day or day of the week. Between-individual baseline intensity (*BIntensity*) scores were averaged over the scores for tracks played during 168 time points for each user, across all hours (which therefore does not vary from hour to hour):

$$BIntensity_u = \overline{Intensity_u} = \frac{1}{\|H\|} \sum_{h \in H} Intensity_u(h)$$

The within-individual intensity score (*WIntensity*) for a person-hour measures the signed difference between an individual's mean intensity score for that hour and their baseline score (as defined above). Within-individual differences in intensity scores measure how a given individual's affective preference varies over time, after removing differences in baseline scores between individuals who are active at different times, leaving only the change over time that is within each individual:

$$WIntensity_{u,g}(h) = Intensity_u(h) - BIntensity_u + \frac{1}{\|UH(g)\|} \sum_{(u,h) \in UH(g)} Intensity_u(h)$$

where u and h pairs indicate user-hours, and $UH(g)$ is the set of all user-hour combinations in the group g (where g can be a day of the week, region, demographic group or chronotype) for which the within-individual pattern is measured.

The final term in $WIntensity_{u,g}(h)$ is the grand mean across all user-hours in g . Note that the final term is $\frac{1}{\|U(g)\|} \sum_{u \in U(g)} BIntensity_u$ for groups g (such as region, demographics or chronotype) as the grand mean across all user baseline intensities in group g . Adding back the group-specific grand mean restores between-group differences while preserving within-individual temporal changes, since adding this constant to the mean-centred data for each individual member of that group does not affect the within-individual temporal dynamics. However, care should be taken in trying to interpret between-group differences by visual inspection of the figures, since the number of observations varies greatly over

the course of the day (see Supplementary Fig. A.3). Thus, a group with much higher musical intensity scores late at night (when listening is less frequent), and only slightly lower scores during the day, might have a significantly lower baseline score than might be inferred simply by imagining a horizontal line fitted to the figure.

Plots in the main text show the mean within-individual intensity scores across different groups for each of 24h over 7d (that is, 168 hourly observations per user):

$$WIntensity_g(h) = \frac{1}{\|UG(h)\|} \sum_{(u,g) \in UG(h)} WIntensity_{u,g}(h)$$

where u and g pairs are the subset of users in group g who were active during hour h , and $UG(h)$ is the set of all users in group g who were active during hour h . These scores reveal diurnal patterns in affective preferences over the course of a day.

2.4.5 Seasonal Variation

The seasonal analysis parallels the diurnal analysis, except that intensity scores are averaged over person-weeks (or person-days for Supplementary Fig. A.2) instead of person-hours. The analysis tests the hypothesized emotional effects of changing day length. The length of the day at a given location varies sinusoidally over the year, with higher amplitude waves the farther one moves from the equator, resulting in long summer days and short winter days in extreme latitudes, and consistent day length near the equator. The day length at a given

location on a given day is governed by the day of the year and the latitude at that location.

Two models are widely used to estimate day length. Although the Center for Biosystems Modeling (CBM) [35] reports more accurate day length estimation than the Brock model [15] when compared with the Astronomical Almanac, this only applies to low and mid-latitudes, with CBM accuracy declining rapidly poleward of 60. Therefore, we use both models, the CBM for < 60 and the Brock model for ≥ 60 .

The Northern and Southern hemispheres have winter and summer six months apart, which makes interpretation of day length patterns awkward when the person-week (or person-day) affective preference is plotted against calendar dates. Instead, the x-axis in Fig. 2.3 is ordered by day length, starting and ending with the winter solstice, with the longest day at the mid-point. The x axis begins with 21 December 2016 for countries in the Northern Hemisphere and 20 June 2016 for those in the Southern Hemisphere, with the summer solstice (20 June in the north and 21 December in the south) at the mid-point, and the day preceding the winter solstice on the far right (see also Supplementary Figs A.1 and A.2).

2.4.6 Group Baseline Comparisons

In the main text, we report baseline differences in mean musical intensity scores across groups in different group categories (for example, day of the week, age, gender, chronotype and geographical region). We performed all statistical tests of group differences in baselines using the unadjusted data, not the mean-

centred data points with adjusted baselines. However, in the figures that report mean-centred within-individual results, we facilitated visual inspection (both of variations around the baseline and of baseline comparisons) by adding back the mean for each group. The group means were also computed from the unadjusted data and did not reflect the mean-centring used to identify within-individual temporal variation.

2.4.7 Other Psychological Features in Music Attributes

Based on a hierarchical PCA on 25 computer-generated attributes for 102 music excerpts across diverse genres and styles, previous research [37] has shown that computer-generated sonic and affective features can similarly capture latent dimensions of human-perceived attributes [50] on the same music excerpts: arousal (the first principle component, indicating music that is danceable and loud), valence (the second; positive and happy) and depth (the third; instrumental and low tempo). While the arousal dimension has very similar characteristics to our intensity measure (for example, positive correlations with danceability and loudness, and negative correlations with acousticness), none of our lower-ranked PCA dimensions was directly matched with the other two dimensions. This is not surprising, given that we applied PCA to 11 audio features generated from a large body of popular music (that is, hundreds of millions of complete songs by millions of artists) while previous work relied on 25 features in hundreds of excerpts from commercially unreleased songs that were previously curated for balance across genres and styles. A curated pool of music excerpts may be suitable for the fine assessment of music preferences and validation of automated feature extraction, but the latent feature structures should

not be expected to match those of actual listening behaviours.

2.4.8 Effects of Daylight Saving Time

The transition to daylight saving time (DST) provides an opportunity to tease apart the effects of day length from the potential confound of biorhythms associated with the lightdark and wakesleep cycles. DST radically shifts the lightdark cycle, but there is only a very small change in day length, which affords the opportunity to use regression discontinuity for causal inference [127]. In our dataset, 31 countries had DST in 2016. We labelled each day of the year relative to the start and end dates for DST for a given country. For instance, Sunday 13 March 2016 was the start date of DST in the United States. Accordingly, 12 March, 13 March and 14 March were labelled -1 , 0 and 1 , respectively. We took mean intensity scores across 31 countries for each labelled day. For each DST start-date and end-date-based daily intensity score, we conducted two tests: (1) non-parametric discontinuity estimation using the smoothing parameter (bandwidth) proposed by Imbens and Kalyanaraman [64, 41] (IK bandwidth) for discontinuity at the DST start or end dates; and (2) McCrary's test [81] for possible discontinuity around the DST start or end dates. Supplementary Fig. A.5a shows the result of the non-parametric discontinuity estimation based on the start date of DST. This indicates discontinuity around New Year's Day and Christmas, but no discontinuity at the start date of DST. This was statistically confirmed by McCrary's test ($z = 0.200$; $P = 0.842$) and by a regression using the local approach with default IK bandwidth ($z = -1.101$; $P = 0.271$; $R^2 = 0.144$). Supplementary Fig. A.5b also shows no discontinuity at the end date of DST, which was statistically confirmed using local linear regression ($z = -0.855$;

$P = 0.392$; $R^2 = 0.399$) and McCrary's test ($z = 0.195$; $P = 0.846$).

CHAPTER 3

CULTURAL VALUES AND CROSS-CULTURAL VIDEO CONSUMPTION ON YOUTUBE

Originally published in PLoS ONE (2017) [94]

Abstract Video-sharing social media like YouTube provide access to diverse cultural products from all over the world, making it possible to test theories that the Web facilitates global cultural convergence. Drawing on a daily listing of YouTube's most popular videos across 58 countries, we investigate the consumption of popular videos in countries that differ in cultural values, language, gross domestic product, and Internet penetration rate. Although online social media facilitate global access to cultural products, we find this technological capability does not result in universal cultural convergence. Instead, consumption of popular videos in culturally different countries appears to be constrained by cultural values. Cross-cultural convergence is more advanced in cosmopolitan countries with cultural values that favor individualism and power inequality.

3.1 Introduction

The recent upsurge of nationalist movements opposing open borders and free trade brings new urgency to questions about the effects of social media on cultural convergence. Video-sharing social media like YouTube provide access to diverse cultural products from all over the world [4]. Unlike traditional media such as television, CDs, or books [18], content on social media (e.g., video clips and music videos) is readily accessible across countries that differ in national GDP [103, 9], geographic location [103, 119], language [111], and religion [119].

Nevertheless, the ability to easily obtain social media content does not mean consumers take advantage of the opportunity. Although technologies increasingly facilitate cross-border flow of media content, previous studies support the “cultural proximity hypothesis” that consumption reflects cultural values that in turn shape cultural norms about socially acceptable content, such that consumers prefer products closer to their own culture [9, 122, 123, 110]. However, these studies focused on consumption of tangible cultural products like books and CDs [122], not content on social media that can be easily downloaded from the Web. Hyperlinks on web pages [9, 123] were also studied extensively, but hyperlinks are generated by producers of online content who vie for the attention of the public and hyperlinks themselves do not reveal consumption patterns of online content.

An important exception is a study [4] showing that Korean pop (or K-pop) music videos on YouTube are highly popular in countries whose cultures differ sharply from Korea as well as in countries that are culturally very similar to Korea. However, it remains to be seen whether this finding generalizes beyond one type of media content produced in only one country.

Using co-consumption of popular videos on YouTube, this study extends research on the cultural proximity hypothesis by examining the relationship between cultural values and cultural openness. Drawing on a daily listing of YouTube’s most popular videos across 58 countries, we investigate the consumption of popular videos in countries that differ in cultural values [62], language, gross domestic product (GDP), and Internet penetration rate.

We chose YouTube because it is the most popular platform for media consumption on the Web, with more than one billion viewers every day, watching

hundreds of millions of hours of content [139, 93]. Video over Internet Protocol is forecast to represent 82 percent of all download traffic by 2020 [24]. Our research addresses why some YouTube videos (e.g., Gangnam Style) are globally consumed while others are limited to a single country, despite the existence of a technological infrastructure for global cross-cultural communication. To find out, we recorded the 50 most popular videos listed by YouTube for the past day for each of 74 countries over six months. “Popularity” is based on YouTube’s undisclosed algorithm [103, 43] that takes into account views, downloads, and likes. Inclusion on YouTube’s top 50 list provides an unranked measure of video consumption.

3.2 Cultural Values

Culture has been defined as a set of values maintained across generations through the socialization process [62, 55]. Although individual attitudes and beliefs may be in constant flux, cultural values are thought to be stable attributes of societies [11]. Cultural values are defined as enduring beliefs that “a specific mode of conduct or end-state of existence is personally or socially preferable” (p.5) [110]. These cultural values influence user decisions about what to view, download, or like, which suggests that YouTube video consumption can be expected to vary across cultures [54, 34].

We operationalized cultural values using Hofstede’s four-dimensional model [62], based on aggregated survey responses from IBM employees in 76 countries. Hofstede’s approach has been criticized by culture scholars who argue that culture is too subtle to quantify, especially in multi-cultural coun-

tries like the United States [10, 40]. Nevertheless, Hofstede’s measures have been widely applied in prominent studies showing how cultural values influence cross-cultural communication behaviors such as media selection and adoption [4, 22], political discussion engagement [31], and use of emoticons on Twitter [92].

The four dimensions in Hofstede’s model are individualism (IDV), uncertainty avoidance (UAI), power distance (PDI), and masculinity (MAS). Each of these dimensions has implications for cross-cultural media consumption that we operationalize in turn below.

3.2.1 Individualism-collectivism (IDV)

Countries with high IDV are more inclined to emphasize “I” rather than “we” and to privilege individual interests over collective welfare (p.130) [62]. Individualistic cultures do not demand conformity around shared opinions, beliefs, or attitudes and are therefore more likely to embrace cultural diversity and to show “respect for other cultures” (p.99) [62]. Hofstede’s argument has been supported by studies [11] showing that people in individualistic cultures tend to be more tolerant of diversity and appreciative of cultural differences. Other studies have found that people in high IDV countries consume more cross-national products [29], adopt global platforms like B2C e-commerce [48] and SNS [22], and purchase newly launched brands [27]. Using a large international hyperlink network, Barnett and Sung [9] found that high IDV countries occupied more central positions in the international information-sharing network. We hypothesize that this pattern will extend to cross-cultural video consumption

on YouTube:

H1: People in individualistic countries will be more likely to consume videos that are popular in culturally different countries, compared to those in collectivistic countries.

3.2.2 Uncertainty Avoidance Index (UAI)

People in high UAI countries are more likely to “feel threatened by ambiguous or unknown situations” (p.191) [62]. For example, opinion surveys have found that people in European countries with high UAI scored higher in aggressive nationalism, ethnocentrism, and xenophobia, including beliefs that immigrants should be sent back to their countries of origin [74]. People living in high UAI countries are reluctant to purchase newly launched products or adopt technological innovations, including the Internet [80], mobile phones [73], SNS [70], and B2C e-commerce [48]. This attitude may extend to consumption of foreign videos:

H2: Countries with high uncertainty avoidance will be less likely to consume videos that are popular in culturally different countries, compared to those in low uncertainty avoidance countries.

3.2.3 Power Distance Index (PDI)

People in high PDI countries are more likely to “expect and accept that power is distributed unequally” (p.61) [62] in groups or organizations. Because Hofst-

ede's measure is based on survey responses of employees, PDI applies most directly to the relationship between bosses and subordinates in organizations [62]. Thus, few studies have examined how PDI influences cross-cultural behavior. However, PDI has implications for beliefs about status inequality that imply cultural preferences for products that signal cultural superiority and it has been demonstrated that people in cultures with high PDI tend to consume products that help them establish and express their status [84]. Bourdieu's classic study [14] shows how cultural products are used to construct and define social class hierarchies. People with high status are believed to have more cultural sophistication, including more extensive and detailed knowledge about foreign cultures. Foreign products provide symbolic benefits such as modernity, prestige, and associations with foreign lifestyles [141] in a similar manner that products with recognized, exclusive, and relatively expensive brand names tend to have higher levels of social status attached to them compared to more generic and less exclusive brands [27]. In high PDI countries, these symbolic benefits constitute a primary motivation for foreign product consumption [141]. This suggests the possibility that people in high PDI countries (including elites as well as those with elite pretensions) are more likely to regard xenophilia as a signal for cultural sophistication [27, 84]. More formally, we expect:

H3: People in high PDI countries will be more likely to consume videos that are popular in culturally different countries than those in low PDI countries.

3.2.4 Masculinity (MAS)

People in high MAS countries are more likely to conform to gender role stereotypes that “men are supposed to be assertive, tough, and focused on material success, whereas women are supposed to be more modest, tender, and concerned with quality of life” (p.140) [62]. Even more than with PDI, MAS does not have straightforward implications for cross-cultural media consumption. On the one hand, it might be argued that masculinity encourages cultural boldness, which implies a greater likelihood to consume unfamiliar cultural content. On the other, traditional gender roles may be associated with parochial cultural tendencies, which implies the opposite association. Moreover, these opposing effects may cancel each other out. We therefore do not hypothesize an association in either direction but instead test to see if high and low MAS countries differ in video consumption.

3.3 Materials and Methods

3.3.1 YouTube Data Collection

YouTube only provides aggregate country-level measures of popularity and we therefore do not have individual user-level data. For each country, YouTube lists daily the “most popular” videos, accessible through the YouTube Application Programming Interface (API). We collected the 50 most popular videos for each of 74 countries over 6 months from November 15, 2014 to April 5, 2015 (approximately 40,700 observations per day) for a total of 4,979,077 observations and

561,931 unique videos. Each observation contains the date, category, title, tags, video duration, average view duration, comments, and popularity metrics, including the number of views, likes, dislikes, and shares for that day.

3.3.2 Bipartite Co-consumption Network

We used the pairwise co-listings of popular videos to construct a bipartite projected network of countries. We first built a bipartite network as proposed by [103, 111] with two types of nodes: (1) 74 countries, each with a list of popular videos collected from YouTube and (2) 561,931 videos on those countries' popular video lists. In the projection of this bipartite network, each country was regarded as a node and an edge was assigned if a pair of countries shared one or more videos on their popular video lists. As it happens, all countries were connected, that is, all had at least one overlapping video with another country. Following the method suggested by Newman [86], an edge weighting was applied that privileges videos that appear less frequently across all 74 lists. Thus, a pair of countries that co-lists a set of videos that are universally popular has relatively low weight compared to a pair of countries that has in common a set of videos that appear on no other lists. This weighting method mitigates the effect of overly popular outliers (potentially due to an artificial increase of viewers) on the co-consumption patterns. This edge weighting combines two components: (1) the number of videos co-listed by a pair of countries and (2) the global popularity (or out-degree) of each video in the co-list, defined as the number of countries in which the video was listed. The weighting metric was formalized as follows:

$$W_{i,j} = \sum_K \frac{\delta_i^k \delta_j^k}{n_k - 1}$$

where $W_{i,j}$ is the weight of the edge between countries i and j ; k is a unique YouTube video in the set of videos co-listed by i and j ; n_k is the number of countries that listed k ; and δ_i^k is 1 if video k is co-listed on popular video lists including country i and 0 otherwise.

Fig 3.1A illustrates the co-consumption pattern on a stylized bipartite network, and 3.1B shows how the edge weight is computed between a pair of countries, the U.S. and Germany. The first video is popular only in the U.S. and Germany, giving it a weight of 1. The weight of the last video is 1/4 because it appears on the most popular lists of five different countries (including the U.S. and Germany). We then derive the edge weight as the sum the weights over all the co-listed videos. Thus, the edge weight reflects the number of co-listed videos weighted by the inverse of the video out-degree (the number of countries that list that video). In Fig 3.1, the edge weight between the U.S. and Germany is 7/4. Using this weighting method, we computed edge weights between all possible pairs of countries.

The outcome of the final step is shown in Fig 3.1C, in which the edge weights are filtered to preserve only those edges that deviate from the expected weight in a null model iteratively produced by a random assignment from a uniform distribution. By imposing a significance level of $p < 0.05$, the links whose weights exceed a randomly expected value are preserved. The remaining links constitute the “backbone” structure of the network [115].

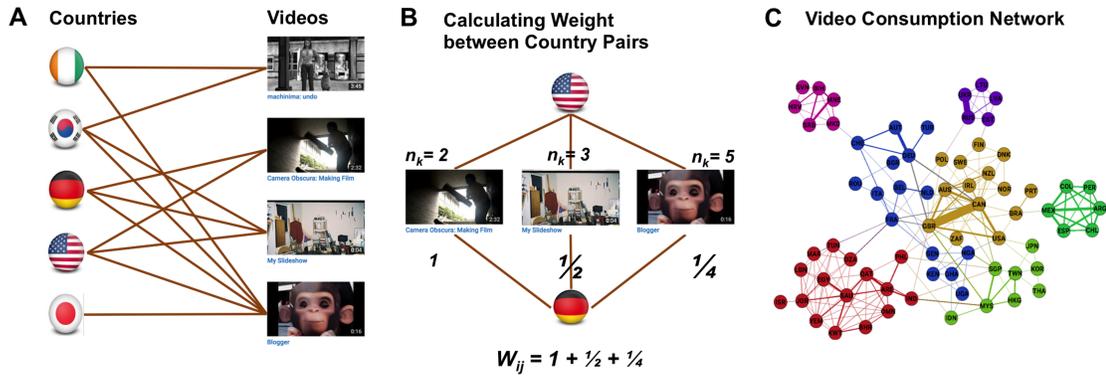


Figure 3.1: Construction of the bipartite network of video co-consumption on YouTube.

3.3.3 Measure of Cultural Openness

We refer to “cultural openness” as the conceptual outcome of interest. A country with high cultural openness can achieve an “optimal blend of novelty and familiarity” by creating cultural bridges [5]. At the same time, a country with high cultural openness can co-consume cultural products with many other countries across different cultural clusters such that they are close to most other countries in the network, which can be perceived as “openness to diversity.” Instead of using a single measure that combines these aspects, we operationalize cultural openness as having two distinct dimensions. Betweenness centrality was used as an indicator of bridging between cultures, measured as video overlap with other countries that do not overlap with one another. Closeness centrality, measured as a country’s level of video overlap with all other countries, provides an indicator of cultural diversity.

We measured cultural betweenness and closeness using Opsahl et al.’s [91] centrality measures in weighted networks to take both the number of ties and the tie weights into account. Those weighted centrality measures are variants

of Dijkstra's algorithm, a well-known method for finding and computing the shortest paths among nodes in a network. Using this approach, the shortest path d between two nodes (i, j) can be defined as follows:

$$d^{w\alpha}(i, j) = \min\left(\frac{1}{(W_{ih})^\alpha} + \dots + \frac{1}{(W_{jh})^\alpha}\right)$$

where w is the weight of the tie between nodes; h are intermediary nodes on paths between node i and j ; and α is a tuning parameter that reflects the influence of edge weights. When $\alpha = 0$, Opsahl's algorithm reduces to the familiar binary measure in which a network edge either exists or does not (i.e., the level of similarity or affiliation between countries can not be captured at all because tie weights are ignored). When $\alpha = 1$, the algorithm is identical to Dijkstra's (i.e., the original feature of the measures, particularly the number of ties, is ignored because tie weights are the sole determinant). A value for $\alpha < 1$ assigns the path with the greatest number of intermediary nodes the longest distance whereas the impact of additional intermediary nodes is relatively unimportant compared to the strength of the ties when $\alpha > 1$. Hence, for $\alpha < 1$, a shorter path composed of weak ties is favored over a longer path with strong ties. Conversely, for $\alpha > 1$, paths with more intermediaries connected by strong ties are favored. The tuning parameter is used to operationalize the extent to which openness reflects a more balanced weight distribution in a node's local network along with its degree. We set $\alpha = 0.5$, although results are fairly robust across other values of the tuning parameter smaller than or equal to 1. Formally, cultural betweenness is given by:

$$C_B^{w\alpha}(i) = \frac{g_{jk}^{w\alpha}(i)}{g_{jk}^{w\alpha}}$$

where g is the sum of shortest paths that pass through node i as a proportion of all shortest paths in the network. Cultural closeness, as the inverse sum of shortest distances to all other nodes from a focal node, is given by:

$$C_C^{w\alpha}(i) = \left[\sum_{j=1}^N d^{w\alpha}(i, j) \right]^{-1}$$

We limited the analysis to the 58 countries for which we had Hofstede scores. The list of countries included and excluded in the analysis is provided in B.1 Table. Descriptive statistics of the four scores on the 58 countries are: (1) individualism ($M = 41.00$, $SD = 23.13$), (2) uncertainty avoidance ($M = 66.72$, $SD = 22.81$), (3) power distance ($M = 61.95$, $SD = 21.59$), and (4) masculinity ($M = 49.55$, $SD = 17.01$).

3.3.4 Economic, Linguistic, and Technological Measures

To disentangle cultural influence from other factors that have been found to affect cultural openness, we included economic, linguistic, and technological measures. Per capita GDP has been shown to be strongly associated with cross-cultural communication on Twitter [39] and in international transactions and communication flows [9, 8]. Previous research also shows strong correlations between GDP per capita and Hofstede's cultural values [62, 79]. In short, GDP per capita is associated with both cultural openness and cultural values. We used the GDP per capita data archived by the World Bank in 2013 [136]. Since the average GDP per capita across 58 countries showed a right-skewed distribution, the base 10 log-transformed GDP per capita ($M = 4.06$, $SD = .56$, $Median = 4.17$) was used in the analysis.

Language is an obvious barrier to any global communication [122, 123] and social media interaction in particular [111, 38]. As a consequence, English as a lingua franca allows greater access to cultural diversity [111, 123], compared to local languages such as Korean or Japanese. Following Ronen et al.'s [111] algorithm, we computed eigenvector centrality of a country's language. The average eigenvector centrality of language across 58 countries is $M = 0.18$ ($SD = .32$, $Median = .025$). The higher the centrality, the lower the linguistic barriers to global communication.

Internet penetration is strongly correlated with Hofstede's cultural values [80] and also limits access to online cultural content. We used the World Bank measure of Internet penetration as the number of Internet users per 100 people [137]. The distribution of Internet penetration is normally distributed ($M = 64.90$, $SD = 21.67$, $Median = 69.48$, ranging from 12.30 to 95.05) and thus does not require log transformation.

3.4 Results

Table 3.1 reports descriptive statistics for the video co-consumption networks across different categories that were automatically classified by YouTube. Each network is based on the edge weights derived from the co-listing of videos in a particular category. The number of nodes is not identical across categories because of data sparsity. Countries were deleted for which there were too few videos listed in that category to obtain statistically significant links in the "backbone" network.

These network characteristics reveal interesting differences across video cat-

Category	Nodes	Edges	Degree	Weighted Degree	Modularity	CC	APL
Combined	72	195	5.417	324.864	0.736	2	3.067
News	73	263	7.205	628.543	0.667	2	2.887
Music	68	159	4.676	145.903	0.410	3	3.332
Games	72	168	4.667	574.939	0.767	5	3.812
Sports	70	173	4.943	391.373	0.695	3	4.406
Entertainment	72	217	6.028	461.752	0.675	4	3.036
Film	73	205	5.616	401.584	0.637	1	3.480
People	73	257	7.041	469.417	0.624	1	2.917
Tech	68	179	5.265	293.704	0.619	2	3.383
Comedy	72	179	4.972	324.729	0.619	2	3.232
Travel	72	228	6.333	300.295	0.598	2	2.857

Note: CC = Connected Components; APL = Average Path Length

<https://doi.org/10.1371/journal.pone.0177865.t001>

Table 3.1: Descriptive statistics of network structure by video category.

egories. For example, the co-consumption news network has low average path length (APL) but high average degree, indicating that news videos were more likely to be consumed among a broader global audience than other types of videos. In contrast, the co-consumption music network has fewer nodes and edges and has low average degree and high APL, which indicates that a country's video list contained less globally popular music and more locally popular music. Interestingly, the gaming network has high modularity and many connected components (CCs), indicating more clustered video preferences.

Although these patterns invite category-specific analyses, the theoretical motivation for this study is focused on differences between countries, not differences between cultural categories. We therefore report results for the combined network based on all videos, regardless of category. However, we also checked the robustness of the overall pattern by examining each category-specific network and found no important differences.

The relationships between cultural values and cultural openness Tables 3.2 and 3.3 report results for regression analyses of cultural openness, operationalized as cultural betweenness and closeness. The cultural model consists of Hof-

stede’s four cultural values and the non-cultural model includes economic, linguistic, and technological measures. The combined model reports the effects of cultural values net of non-cultural.

	Full model	Non-culture Model	Culture Model
Intercept	-0.425* (0.172)	0.014 (0.073)	-0.339* (0.147)
Non-cultural factors			
GDP per capita (log-transformed)	0.092 (0.219)	0.087 (0.228)	
Language eigenvector centrality	0.092 (0.085)	0.173* (0.077)	
Number of Internet users	0.028 (0.240)	0.026 (0.242)	
Cultural values			
Individualism (IDV)	0.354* (0.138)		0.426** (0.123)
Uncertainty avoidance (UAI)	-0.031 (0.122)		-0.061 (0.113)
Power distance (PDI)	0.410* (0.156)		0.373* (0.147)
Masculinity (MAS)	0.214 (0.125)		0.250* (0.124)
Sample size (number of countries)	58	58	58
Model-fit indices			
R^2	0.290	0.110	0.263
Adjusted R^2	0.191	0.061	0.208

Note

* $p < .05$

** $p < .01$. Unstandardized coefficients are reported with standard errors in parentheses. In order to compare coefficients, variables included in the analyses were rescaled to the unit interval.

<https://doi.org/10.1371/journal.pone.0177865.t002>

Table 3.2: OLS regression model of cultural betweenness among 58 countries.

We tested both models for heteroscedasticity and multicollinearity. For cultural closeness, we could not reject the null hypothesis that the variance of the residuals is constant, i.e., heteroscedasticity is not present, using the test of non-constant variance score. For the model of cultural betweenness, in contrast, we inferred that the residuals are heteroscedastic. However, as described earlier, results are robust across different tuning parameters and a model constructed with a composite measure of betweenness and closeness designed using Rao-Stirling diversity [121, 96] (see S1 Text in [94] for more details). The variance inflation factor on each variable of the full model is smaller than two except for

	Full model	Non-culture Model	Culture Model
Intercept	0.138 (0.109)	0.368*** (0.051)	0.285** (0.100)
Non-cultural factors			
GDP per capita (log-transformed)	0.206 (0.139)	0.193 (0.161)	
Language eigenvector centrality	0.074 (0.54)	0.190** (0.055)	
Number of Internet users	0.014 (0.153)	0.001 (0.171)	
Cultural values			
Individualism (IDV)	0.229* (0.088)		0.328*** (0.084)
Uncertainty avoidance (UAI)	-0.229** (0.077)		-0.234** (0.077)
Power distance (PDI)	0.314** (0.099)		0.236* (0.100)
Masculinity (MAS)	0.245** (0.086)		0.270** (0.084)
Sample size (number of countries)	58	58	58
Model-fit indices			
R^2	0.540	0.284	0.452
Adjusted R^2	0.476	0.244	0.411

Note

* $p < .05$

** $p < .01$

*** $p < .001$. Unstandardized coefficients are reported with standard errors in parentheses. In order to compare coefficients, variables included in the analyses were rescaled to the unit interval.

<https://doi.org/10.1371/journal.pone.0177865.t003>

Table 3.3: OLS regression model of cultural closeness among 58 countries.

GDP per capita (2.41) and Internet diffusion (2.49) that are strongly correlated with each other but neither contributes significantly to model predictions.

As shown in Tables 3.2 and 3.3, a country's cultural openness is much better explained by cultural values (adjusted $R^2 = .208$ for cultural betweenness; adjusted $R^2 = .411$ for cultural closeness) than non-cultural measures (adjusted $R^2 = .061$ for cultural betweenness; adjusted $R^2 = .244$ for cultural closeness). Indeed, non-cultural measures do not make a significant contribution to the full models' explanatory power, and removing these measures even improves the adjusted R^2 (.191) of the model on cultural betweenness.

The coefficients in Tables 3.2 and 3.3 provide more detailed results. The eigenvector centrality of language [111] indicates that countries using more

global languages (e.g., English) have greater cultural openness, consistent with the findings in previous studies ($b = .173, p < .05$ for cultural betweenness; $b = .190, p < .01$ for cultural closeness) [122, 123]. However, this effect largely disappears when cultural values are included in the model ($b = .092, p = .29$ for cultural betweenness; $b = .075, p = .18$ for cultural closeness), indicating that cultural values are stronger predictors of cultural openness and capture most of the linguistic effect.

The results in Table 3.2 (cultural betweenness) support H1 and H3 but not H2. As hypothesized, YouTube users consume more videos in common with other countries that do not overlap with one another if those users are located in countries that are more individualistic ($b = .354, p < .05$) and with greater power distance ($b = .410, p < .05$). However, a country's cultural openness is not predicted by uncertainty avoidance ($b = -.031, p = .80$) or masculinity ($b = .214, p = .12$). In short, individualism and acceptance of power inequality are associated with an optimal blend of novelty and familiarity, as indicated by greater cultural betweenness.

Results in Table 3.3 (cultural closeness) support H1, H2, and H3. YouTube users consume more videos in common with a larger number of culturally diverse countries (i.e., higher cultural closeness) if those users are located in countries that are more individualistic ($b = .229, p < .05$), with less uncertainty avoidance ($b = .229, p < .01$), greater power distance ($b = .314, p < .01$), and higher conformity to gender role stereotypes ($b = .245, p < .01$). In short, cultural closeness is associated with more cultural values than is cultural betweenness, and both are more important than the non-cultural factors that have been the focus of previous research.

3.5 Discussion

Our findings are consistent with the view that cross-cultural convergence, especially cultural closeness, is more advanced in cosmopolitan countries with cultural values that favor individualism, power inequality, and tolerance for uncertainty. Online social media facilitate global access to cultural products, yet this technological capability does not result in cultural convergence [119, 111, 87]. Instead, consumption of popular videos in culturally different countries appears to be constrained by cultural values.

These findings contrast with studies showing that shared language, common economic system, and geographical proximity are associated with cross-cultural consumption of tangible products [18, 122] and flows of information [9, 119, 123]. The difference with previous results may reflect fewer linguistic, economic, and geographic constraints on video consumption, as well as less need for active interaction with people of different cultures compared to exchanges of e-mail or Tweets, making it easier and more comfortable for YouTube users to encounter and enjoy videos from diverse cultures.

Our findings have implications for the recent upsurge of nationalist movements opposing open borders and free trade. On the one hand, contact theory [1] and “soft power” research [89] suggest the possibility that cross-cultural exposure could promote cultural innovation and mutual understanding. On the other hand, cultural openness may erode a country’s unique cultural identity, leading to a nationalist backlash.

Additionally, this study has substantial implications for the distinction between cultural betweenness and closeness in cross-cultural experience. Bail [5]

highlighted the implications of cultural betweenness as an indicator of cultural bridgesnetwork positions that can “achieve an optimal blend of novelty and familiarity.” However, cultural closeness has not received theoretical attention or empirical inquiry. Our findings show that cultural closeness can be an indicator of multicultural identity: countries with low closeness (e.g., Kenya) have a narrow range of video preferences that forms a cultural niche, possibly associated with national identity. In contrast, countries with high closeness (e.g., Canada) have a wide range of video preferences that spans cultural niches and might be associated with a multicultural national identity.

An important limitation of this study is that the units of analysis are countries, not individual users. This poses the possibility that the results we report are susceptible to the ecological fallacy. For example, in some cases, we found that there are significant overlaps in popular video consumption between countries of migration destination and origin. It is possible that individual members of each immigrant group have parochial cultural preferences, but because the groups differ in their preferences, the country appears to be culturally open. We therefore tested for the spurious effects of migration and found significant correlations between cultural openness and a countrys degree in the international migration network, where edges correspond to migrations from the country of origin to the country of destination, derived from 2015 UN migration stock data ($r = 0.33, p < .05$ for cultural betweenness; $r = 0.44, p < .001$ for cultural closeness; models including this migration degree as an additional independent variable show identical results with original models, but individualism is no longer significant; more details are provided in S2 Text in [94]). Individual user data is needed so that this possibility might be tested more fully in future research. Future research with individual data might also explore possible asso-

ciations between cultural openness and the incidence of cultural “omnivores” in the population [132].

3.6 Conclusion

Our study makes two important contributions. First, we found that cultural values are significantly associated with the cultural openness of a country, as measured by the consumption of YouTube videos that are popular across diverse cultures. Moreover, this association with cultural values appears to account for effects that previous research has attributed to non-cultural factors. Second, we provide a new angle from which to view the cultural proximity hypothesis in the era of social media on the globalized Web.

CHAPTER 4
UNDERSTANDING DIVERSITY OF MUSICAL TASTES VIA ONLINE
SOCIAL MEDIA

Originally published in ICWSM 2015 [96]

Abstract Musicologists and sociologists have long been interested in patterns of music consumption and their relation to socioeconomic status. In particular, the Omnivore Thesis examines the relationship between these variables and the diversity of music a person consumes. Using data from social media users of Last.fm and Twitter, we design and evaluate a measure that reasonably captures diversity of musical tastes. We use that measure to explore associations between musical diversity and variables that capture socioeconomic status, demographics, and personal traits such as openness and degree of interest in music (intensity). Our musical diversity measure can provide a useful means for studies of musical preferences and consumption. Also, our study of the Omnivore Thesis provides insights that extend previous survey and interview-based studies.

4.1 Introduction

The cultural and social significance of music is universal; music is found in every known human culture, and plays a role in rituals, wars, ceremonies, work, and everyday life [133]. Tia DeNora [28] noted that “Music is not merely a meaningful or communicative medium. It does much more than convey signification through non-verbal means. At the level of daily life, music has power. It is implicated in every dimension of social agency.” As social media become more ingrained in our lives, it follows that connections between social media use, and habits and norms regarding music consumption, will occur. In this pa-

per, we present an empirical analysis of social media data as they relate to and reveal details of users' musical tastes.

A person's musical consumption can reveal a lot about their personality, preferences, and sense of self. One can have limited tastes; they may listen to a single genre like pop or rap, and not diverge into other genres. On the other hand, another individual may be eclectic in their musical choices and have a playlist filled with jazz, hip-hop, indie rock, classical, and so forth. We often think of such differences as a matter of individual choice and expression; however, to a great degree, it is hypothesized and tested that the diversity of musical tastes can be explained by external factors. For example, previous research has identified a relationship between musical tastes and social factors, and produced the *cultural omnivore thesis*. This thesis describes "a shift in the orientation of high-status individuals toward an inclusive range of musical preferences that traverses the traditional boundaries between *highbrow*, *middlebrow*, and *lowbrow* genres [97, 98, 99]." However, symbolic boundaries between musical genres have been eroding [45] in recent years, which provides an opportunity to rethink the high-to-lowbrow cultural categories in relation to musical diversity. This can lead to a better understanding of the impact of social conditioning on diverse musical tastes, and by proxy, a better understanding of the connection between socioeconomic status, demographics, and the diversity of musical preferences.

To date, the social computing community has examined online listening activity as source of information and recommendations for music [16, 140, 33, 129]. However, computational tools and online outlets such as social media can make further contributions toward understanding human behavior related to musi-

cal consumption and help to elaborate user-centered music retrieval systems by analyzing personal characteristics. We focus on exploring a new means of measuring the diversity of individual musical tastes by using data collected from social media, and examine the relationship between musical diversity and various individual factors including socioeconomic and demographic information, as well as social and individual information that can be collected from social media.

Through a multi-platform analysis of a dataset of U.S. Last.fm¹ users and their corresponding Twitter accounts, we examine music consumption together with demographics (e.g., age and gender) and other descriptive variables for a community music fans who have an online presence. Using Twitter-derived information for these users, we inferred their socioeconomic information (e.g., income, education level, and area of their residence) as well as other social and personal variables (e.g., how diverse their friends and interests are, and how ‘open’ and ‘into music’ they are). We then defined a measure for musical diversity by applying the notion of shared understanding as socially perceived distances between genres. We suggest that designing a diversity measure can provide a useful means for studies in recommendation systems. Moving from designing a measure to analysis of associations between diversity and individual factors, we suggest this type of analysis can provide meaningful insights that are complementary to those provided by previous survey and interview-based studies regarding the musical omnivore thesis. Our main contributions therefore are as follows:

¹Last.fm is a music recommendation service. The site builds a detailed profile of each user’s musical consumption by recording details of the tracks the user listens to, either from Internet radio stations, or the user’s computer or many portable music devices. It also offers some social networking features such as recommending and playing artists to Last.fm friends [135].

- We propose and validate a novel diversity measure that borrows the concept of Rao-Stirling diversity for music consumption. While recent studies [63, 33] define diversity (as it relates to music consumption) as the total number of unique genres associated with all artists listened to, we go into more detail, and define diversity as a multidimensional property that has three main attributes: *variety* (the number of unique genres one listened to), *balance* (the listening frequency distribution across these genres), and *disparity* (the degree of distance between musical categories).
- We investigate the relation between musical diversity and various other variables including socioeconomic factors. In particular, we find that followers of high-profile news media are more likely to have diverse musical tastes. We also consistently find a weak, but robust trend for people who are more ‘into’ music to have less diverse tastes. Along with these findings, our results also show that demographic factors such as age and gender are associated with musical diversity rather than conventional socioeconomic status such as income and education level.

We begin by reviewing the primary key research around the diversity of musical tastes, and then identify possible challenges for developing better measures of diversity.

4.2 Related Literature

Disciplines such as sociology and social computing addressed the notion of *cultural omnivorism* and the importance of understanding the musical diversity. Given the wealth of related work on these topics, our review focuses on what

could be tested by complementing the limitations of previous studies through social media data and how we can design a meaningful measure for the diversity of musical tastes.

4.2.1 Changing Status of the Omnivore Thesis

Since the publication of Bourdieu's seminal work *Distinction* [13], in which he explains the notion of cultural capital and exhibits how access to education, knowledge of the arts, and familiarity with other highly regarded aspects of western culture lead to a 'highbrow' status, copious research has investigated the relationship between socioeconomic position and musical tastes [26]. The majority of the current studies on the *omnivore thesis* in relation to musical tastes, proposed by Richard Peterson [97] show that people with a higher socioeconomic status have broader (omnivorous) musical tastes than those with a lower socioeconomic status who have limited (univorous) musical preferences in low-brow music. There are generally two definitions of omnivorousness, referred to as the *volume* and the *compositional* definitions [134]. The first refers to higher socioeconomic status people favoring more musical genres than those of lower socioeconomic status. The second refers to the situation that people with higher socioeconomic status tend to have more eclectic tastes across the spectrum of high-to-lowbrow music than people with lower socioeconomic status.

More recently, however, Peterson [99] conducted comparative research and noted that "despite the attention paid to the concept by numerous scholars, the subtypes of omnivorousness suggested by them were diverse and fall into no recurrent patterns due to changes in the socio-cultural world." Indeed, though

there is a little disagreement that the contemporary era has witnessed shifts in the ways cultural preferences and practices are mapped onto social locations, the extent to which this implies changes in the functioning of cultural capital remains unclear [109]. In addition, Peterson [99] raised a question regarding the traditional measurement of omnivorousness, and recent qualitative studies identified a number of limitations in conventional survey-based studies [134, 109]: First, the simple or compositional volume of genres preferred by an individual is insufficient to show the full picture of one's form of engagement and social status since different conceptual frameworks may provide different understandings. Second, there is a tendency to discriminate genres within preferred genres (i.e., even though one answers 'rock' as a preferred genre, it does not mean that one likes *all* kinds of rock; therefore, it is possible that someone who likes a Heavy Metal, a subgenre of rock, says "I like rock," and someone who likes the same subgenre says "I don't like rock"). This inability to discriminate genres, or lack of knowledge regarding how to best express what genres one prefers, can create confusion [108]. This gap may bring inconsistency in the preference scoring across survey participants. Finally, the high-to-lowbrow scheme should be reconsidered in contemporary social contexts as Peterson (2005) argues that there is no consensus. In addition, a lot of research has used inconsistent levels of genres, e.g., a questionnaire of preferences for opera, jazz, rock, and heavy metal may be used in these types of surveys, even though heavy metal is often considered a subgenre of rock.

We believe online social media data can help rectify some of these limitations and provide a unique and useful perspective on the musical omnivore thesis: data collected from social media sites can provide a unique capacity to (i) reduce the inconsistency of preference scoring (which may differ across people

due to their inability to discriminate) by systematically classifying the genres consumed by users, (ii) explore a different level of relationship between social status and musical tastes by accessing the subgenres of choice among users, which are more fine-grained than higher-level genres, and (iii) analyze data on a consistent level of genre-hierarchy. Further, social media data can provide users with open-ended spaces [76] in which to list their favorite music, concert attendance, and direct/indirect musical information sources, which offers an unprecedented opportunity to examine how tastes are associated with various individual factors. Up to now, the majority of research on musical tastes has relied on closed-ended surveys typically measuring preferences in terms of genres, and our aim is to contribute a new way to look at the relationship between musical preference and various social and individual factors.

4.2.2 Technology and Music Listening Practice

Exploring musical diversity is an interesting challenge in social computing, as well as music information retrieval (MIR); it also has many applications in real-life scenarios. In MIR, some researchers have explored to achieve the optimal balance between the two objectives on recommendation, similarity and diversity, because it has been recognized that being accurate with similarity metric alone is not enough to judge the effectiveness of a recommendation system [82, 21]. In addition, recent studies [21, 33] suggest that one's personality might have a role in the formation and maintenance of music preferences, and diversity of musical tastes could serve as a proxy of the level of openness of one's personality. These studies show that looking at musical diversity as an indicator of openness can have an impact on the performance of a collaborative

filtering recommender system. In social computing, diversity has been considered in studying phenomena such as peer influence and music consuming mechanism. Some of this research confirms that informational influence is the key underlying mechanism of music listening practices [138] and systematic recommendations affect users' choices of music tracks and listening behaviors [17].

4.2.3 Research Questions

We believe associations between musical categories (e.g., genre-to-genre and subgenre-to-subgenre) can be reasonably derived from the perception of crowds by analyzing their musical consumption, and these distances may help design better measures of musical diversity. The existing measures, *volume* or *entropy*, are different from diversity, and thus cannot accurately capture its essence. Volume, which is defined as the number of musical categories one listens to, does not consider whether a person listens with balance. A 99%–1% split between two genres would be treated the same as a 50%–50% split. Entropy, on the other hand, takes the distribution into account, so a more skewed distribution would be considered less balanced. However, entropy does not look at the similarities of the musical categories and implicitly assumes all categories to be equidistant to each other (e.g., listening to three different styles of metal music would be the same as listening to classical music, death metal, and salsa). People, however, do consider certain types of music as similar or dissimilar [85]. To define and to quantify this notion of similarity we use *co-consumption* behavior. For example, if both rap and hip-hop are consumed by many people we assume that these two genres are similar. Having musical consumption data for a large user set can reveal the distance between musical categories.

The challenges and opportunities in studying musical diversity lead us to introduce two research questions that guide the remainder of this paper:

RQ1 *Can a novel diversity measure using variety, balance, and distance between musical categories capture the diversity of musical tastes better than existing methods?*

RQ2 *What variables are associated with diversity in music consumption? Is socioeconomic status a factor or are other factors also associated?*

4.3 Method

The literature referenced in the previous section points to three major dimensions of explanatory variables: socioeconomic status, demographic information, and ‘openness’ (degree of appreciation for novelty and variety of experience). With these dimensions and the additional dimension of ‘into-ness’ (degree of self-disclosed interest in music) as a guide, we identified 15 variables. We inferred socioeconomic status including income, education level, ethnic diversity of area of residence, and urbanness of area of residence by using geocoded tweets. Into-ness (i.e., degree of music-related topics of interest in Twitter) and openness including number of friends, timezone diversity of friends, and interest diversity was inferred by using tweets, profile descriptions, and friendship information in Twitter. We directly downloaded demographic information (e.g., gender and age) and other types of into-ness (e.g., number of event attendance in the past, number of loved tracks, period after registration, and number of friends in Last.fm) through the Last.fm API.

4.3.1 Initial Data Collection

To identify and obtain a sample of Last.fm users in the U.S. who share gender, age, and Twitter user names in their Last.fm profiles, we used the Google Custom Search API and the Bing Search API. We created a custom query containing parameters that returned only Last.fm user pages which contained this particular information. To augment the sample size, we collected U.S. Twitter users who share their Last.fm accounts in their Twitter profiles by using the ‘Search Bio’ feature in Followerwonk². This allowed us to obtain 23,294 unique users. Then, we collected all publicly available tweets from that user population. During this process 4,392 unique users were screened out since some of them did not allow public access to their tweets or had removed their accounts in the meantime. This left us with 18,902 unique users. To infer socioeconomic status by using geocodes in tweets, we limited our remaining sample to those users who posted at least ten tweets with geocodes, which resulted in 3,548 users. Along with Twitter data, we collected Last.fm data including ‘Top artists’ list (i.e., the 50 musicians a user listened to the most; listening frequency for each artist is included) as well as demographic and some into-ness information directly through the Last.fm API.

4.3.2 Socioeconomic Status

We used home location derived from Twitter as an index to approximate socioeconomic data, and news interests, expressed via Twitter’s following network, as another proxy for socioeconomic status.

²<https://followerwonk.com>

A user’s home location can be a marker of their socioeconomic status. In particular, the socioeconomic status of social media users can be estimated by extracting the users’ hometown ZIP codes and matching that to the median ZIP code household income according to the Census Bureau [76]. In addition, using the inferred home location we can check whether a user lives in an urban or rural area [60].

To obtain the home location for a user, we followed a procedure that involved three different methods of identifying a user’s possible home ZIP code. We first reverse-geocoded all the latitude and longitude tags for the user into the ZIP codes, using the Nominatim API³. We also extracted Federal Information Processing Standard (FIPS) codes, which represent specific regions in counties, using the Coordinates to Political Areas API in Data Science Toolkit⁴. Using the ZIP code data for the user, we inferred a probable home location of a user when we found an intersection between the sets of potential ZIP codes for the user computed by three different methods, the *plurality* and *n-days* methods summarized in [60] and the *plurality with time limitation* described in [19].

The plurality approach [60] assumes that the single region in which a user was the most active is the user’s home location. Using this approach, we find the user’s mode ZIP code(s) from which tweets were most frequently posted. The *plurality with time limitation* method is based on the finding in [19], that people are most likely home between 10pm – 6am. Using these parameters, we identify the user’s mode ZIP code(s) from which tweets were most frequently posted during that time period. Since the plurality approaches may not be appropriate for users who travel frequently, the final method we used identified the ZIP

³<http://www.nominatim.org>

⁴<http://www.datasciencetoolkit.org>

code(s) in which a user posted over a period of at least 10 days, considering them ‘local’ to that area if they did.

We selected a single home ZIP code (and FIPS code) for each user by intersecting the ZIP code sets resulting from the three methods mentioned above. The final set of users with non-empty intersection had 1,306 users (there were 3,451, 3,258, and 1,822 users with non-empty sets for each of plurality, plurality with time constraint, and n-days methods respectively). All other users for which we could not robustly estimate a location were removed from the data.

Finally, to extract socioeconomic data, we used each ZIP code to query the 2010 US Census data to determine income, education level, and ethnic diversity in the area. We matched each FIPS code to NCHS data for urban–city classification of the area which places every U.S. county on a discrete scale from 1 (a large central metro area) to 6 (a sparse rural area). For each user we thus have values for median household income, percentage of bachelor degrees, proportion of white people⁵, and urbanness: these are our socioeconomic proxy measures. This process resulted in 1,306 users for whom we have self-declared gender and age, as well as inferred income, education level, and characteristics of the area of residence⁶.

In addition to location-derived socioeconomic data, we used news interest as

⁵We tested relation between white ratio and ‘racial and ethnic diversity’ by using the Ethnic/Racial Diversity Index which defines racial and ethnic diversity as $1 - \sum_{r \in G} P(r)^2$ where $P(r)$ is proportion of a race population r and G is represented race groups (in our case: white, black, Native American, Asian, Hispanic, Pacific Islander, two or more races, and other races by following ethnicity distribution in the 2010 Census). A higher index number denotes more diversity. However, there is confusion among the general population about the designation of the Hispanic identity since ‘Hispanic’ in the census refers to any ‘race,’ both black and white. So, we decided to use the simple metric, $1 - \text{white ratio}$, as ‘Racial Diversity’ since it is clearer. The Pearson correlation between the white ratio and ethnic diversity was 0.667 ($p < 0.001$).

⁶We ignored 97 users due to various ZIP code issues, such as ZIP code that were invalid, not available from the census data, or too small to have socioeconomic statistics.

a proxy for socioeconomic variables. According to Pew Research [101], regular news audiences often are more formally educated and have higher household incomes. In particular, readers of *The New Yorker* and *The Economist* news media tend to be highly educated and high earners [101]. We therefore created a variable that indicates whether each of our users follows *The New Yorker* (@NewYorker) or *The Economist* (@TheEconomist) on Twitter.

4.3.3 Genre and Subgenre Information Collection

For each user, we extracted the categories of music they listen to at both genre and subgenre ('style') levels. For each user we retrieved the top 50 artists the user listened to via the Last.fm API. We collected genre and subgenre information for each artist using the API for *Allmusic*⁷, a well-known music database (DB). Unlike other music content databases, *Allmusic*'s metadata is professionally edited and thus is likely to be more consistent when assigning genres or subgenres to artists. Many high-profile music sources like iTunes and Spotify currently use *Allmusic* to handle relevant artist information.

We matched each artist name collected from Last.fm to an artist entry on the *Allmusic* DB only if the result exactly matched the queried artist name. When multiple musicians with the same name were matched, we used the *Allmusic* engine's relevance ranking which is based on usage data and editorial weighting. We manually validated the *Allmusic* ranking for a random selection of 100 artists that had multiple entries. We examined the Last.fm page for the artist (as linked from the user's Top 50 list, i.e. uniquely identified) and the *Allmusic* page for the top-ranked artist by the same name as retrieved by the API. We

⁷<http://www.allmusic.com/>

found that the top-ranked artist matches with the Last.fm artist for all cases in this sample of 100.

A single artist could be classified into multiple genres and subgenres, in which case we distributed the artist’s ‘weight’ equally between the respective genres or subgenres. During this data processing, we dropped 292 users who did not have full set of 50 artists that were classified by Allmusic and listened to more than 100 times by the user. As a result, data for 1,014 users were analyzed. There were 8,490 unique artists among the Top 50 artists of 1,014 users, and 987 artists among the unique artists were matched with more than one exact name in Allmusic DB (e.g., Nirvana and Spoon).

4.3.4 Measuring Diversity

We calculated the diversity of music consumption for each user using both genre- and subgenre-level data derived from their Last.fm activity. We previously argued that in order to explore diversity, we need to investigate multiple factors, namely: the number of genres listened to (variety), the distribution of playing frequency among genres (balance), and, crucially, how related these different genres are (measured via some distance or similarity). These assumptions align well with the concept of Rao-Stirling diversity [120, 121, 104, 77].

To operationalize the concept of diversity, following Rao-Stirling, we computed the diversity of musical tastes of a user u as $\sum_{i,j \in N} p_{u,i} \times p_{u,j} \times d(i, j)$. In this formulation, $p_{u,i}$ is the fraction of user u ’s preference for genre i (we performed separate and equivalent calculations for genres and subgenre information; the description here focuses on genre information). To compute $d(i, j)$, we com-

puted the pairwise co-consumption between musical categories as a proxy of closeness. Using an $M \times N$ genre proportion matrix of $p_{u,i}$ values (for each row u , $\sum_i p_{u,i} = 1$), we computed every possible pair of genre-to-genre cosine distances between the matrix columns, representing closeness between genres. The distance $d(i, j)$ is the cosine distance, i.e., $1 - \text{cosine similarity}$, between the genres. As mentioned above, we repeated the same process with subgenre information. For illustration, the resulting distances for *genres*, embedded in two dimensions using multidimensional scaling [71] (MDS), are shown in Figure 4.1⁸.

This approach to computing diversity of music consumption has a number of useful qualities. A user who equally (balance) consumes many types of music (variety) that are pairwise highly dissimilar (distance) will have a large diversity score, whereas a user disproportionately consuming a few pairwise similar types of music will have a low diversity score. We evaluate this approach and its robustness below.

4.3.5 Into-ness and Openness

For each user, we calculated several variables that capture openness (preference for novelty and variety) and into-ness (degree of interest in music) using Twitter and Last.fm data. To help inferring into-ness and openness regarding each user's interests, we first inferred the user's general interests by using a method proposed in [12]. For a given Twitter user u (whose interests are to be inferred), the method first checks which other users u is following, i.e., users from whom u

⁸Interestingly, highbrow and middlebrow genres (e.g., classical, easy listening, and jazz) are close to each other rather than being close to lowbrow genres (e.g., pop&rock, folk, country, rap) even though we used an inductive approach to identify the distance between musical categories rather than assuming that musical tastes are shaped by certain schemes.



Figure 4.1: Multidimensional scaling for distance between genres.

is interested in receiving information. It then identifies the topics of expertise of those users (whom u is following) to infer u 's interests, i.e., the topics on which u is interested in receiving information. Expertise is defined by the users bio or tweets via the Lists feature in Twitter [42].

Using the interest topics for each user, we computed openness and into-ness measures. As a proxy of openness, we computed the diversity of the user's interests using the same method we calculated music consumption diversity above. In this case, for example, similarity of interests can be derived from the cosine distance between interest in a matrix that captures users' interest breakdown. As other measures of openness, we counted for each user in our dataset the number of people they are following on Twitter and also the number of unique timezone in 100 randomly sampled people from whom they are following. We collected these openness variables inspired by [113, 106]⁹.

⁹We did not consider lexical features of tweets as variables since previous efforts [44, 105, 114] showed a disagreement regarding predicting features for openness.

As a proxy of music into-ness, we used the proportion of music-related interests (any interest topic that included the term ‘music’) among the entire set of user interests along with other types of into-ness that were directly collected via the Last.fm API: number of event attendance in the past, number of loved tracks, period after Last.fm registration, and number of friends in Last.fm.

Table 4.1 presents 15 variables we identified and diversity on genre and sub-genre along with their distributions.

4.3.6 Data Validation and Preparation

Given that some of our variables were indirectly derived from social media data, we performed validation tests for our key variables.

Reverse Geocoding

To validate our geocoding framework, we matched the inferred ZIP code to the self-reported home location of the user on their Twitter profile. Out of 100 randomly sampled users, eight users did not disclose their location on their Twitter profile or did not properly disclose their location like “not in a cornfield but...close” and “up in the air.” Among the rest of them (92% of users), only eight users’ locations did not overlap with the inferred zip code location. In other words, more than 90% of inferred locations were well-matched to the self-reported home locations at town/city/state levels.

Note that it is unusual to have as much as 92% of users with a valid location field [59]. Our dataset, though, includes Twitter users who are also heavy users

Socioeconomic Variables	Distribution	Max
Income		192,250
Education		100
Racial Diversity		0.98
High-profile News Reader	High-profile News Media Followers: 12.4% Non-followers: 87.6%	
Urbanness		1-6 (Scale)
Demographic Variables		
Age		52
Gender	Female: 30.7% Male: 69.3%	
Into-ness Variables		
Musical Event Attendance		1,504
# of Loved Tracks		12,619
Days from Registration		4,305
# of Last.fm Friends		2,036
Interest in Music		2,456
Openness Variables		
# of Twitter Friends		10,954
Timezone Diversity of Friends		31
Interest Diversity		0.76
Diversity		
Diversity on Genre		0.67
Diversity on Subgenre		0.80

Table 4.1: Fifteen variables used to explain the measured musical diversity scores and genre- and subgenre-level of diversity scores. The distributions accompanying each variable begin at zero and end at the adjacent maximum. Many variables are not normally distributed.

of the geo-tagged tweets feature; it is conceivable that the same group more readily exposes location in their profile data.

Socioeconomic Status

Even if we get the user's location right, the derivation of their socioeconomic information may be wrong as the user may not be *representative* of where they live. For example, it is possible that people who use both Twitter and Last.fm have similar socioeconomic status, regardless of what sort of neighborhood they live in. However, if the inferred socioeconomic information are correct, they should correlate with our other proxy for socioeconomic status: following the New Yorker or Economist. We thus validate our socioeconomic measures by examining whether our inferred income and education level are associated with following the New Yorker (@NewYorker) or The Economist (@TheEconomist) Twitter accounts. Indeed, compared to other users, New Yorker and Economist followers had higher status for all inferred income and education values, including adjusted gross income (AGI), household income, and level of post-secondary degree (both bachelor's and graduate). These differences were statistically significant as determined by a one-way ANOVA (New Yorker followers AGI: $p < 0.01$; median household income: $p < 0.05$; bachelor degree: $p < 0.001$; graduate degree: $p < 0.001$; Economist followers AGI: $p < 0.001$; median household income: $p < 0.05$; bachelor degree: $p < 0.001$; graduate degree: $p < 0.001$).

Data Imputation and Standardization

In our final dataset, 189 out of 1,014 subjects had missing values in one or more variables. According to [56], if the missing data level is under 10% in each variable, any imputation method can be used to augment the missing values. We used multiple imputation methods in our dataset: we applied Bayesian linear regression for continuous variables, and linear discriminant analysis for factor

variables. We also standardized all the variables for the final analysis.

4.4 Results

Our primary purposes for this study were (i) to design a measure that reasonably captures the notion of ‘diversity of musical tastes’ and (ii) to explore associations between musical diversity and various individual factors regarding dimensions of socioeconomic status, demographics, and personal traits including openness and into-ness in music.

4.4.1 Diversity Measure

To answer RQ1, we estimated the reliability of our diversity measure. We asked three independent annotators to assign a diversity level to the musical consumption of 25 randomly chosen users. The annotators ranged in their music knowledge; we had an expert (musicologist), a music fan, and a casual listener. We provided the annotators two sets of tables of genre- and subgenre-based listening proportion of the 25 users. We asked the annotators to carefully examine each user’s listening pattern and apply a 6-point diversity Likert scale where ‘5’ meant very diverse musical taste, ‘1’ meant very low diversity, and ‘0’ meant no diversity at all (it is possible that a user listened only to one genre). We did not provide the annotators with any other information or instructions (such as “consider the relationship between genres”) as we wanted to know their natural impressions and interpretations of diversity based on their own experiences. Fleiss’s Kappa and average pairwise Cohen’s Kappa were used

to assess the inter-rater reliability for the evaluation. For genre-level the Fleiss Kappa score was 0.411 ($p < 0.001$) indicating moderate agreement, and the Cohen's Kappa score was 0.819 ($p < 0.001$) indicating almost perfect agreement. For subgenre-level, the respective scores were 0.011 ($p > 0.1$) indicating slight agreement and 0.415 ($p < 0.05$) indicating moderate agreement. We averaged the rater responses for each user and used that below as the raters' diversity score.

To evaluate our diversity measure, we calculated the Pearson correlation between the raters' average score and our computed diversity score. For genre-level diversity, the correlation between our measure and the raters' diversity was 0.94 ($p < 0.001$). For the subgenre-level diversity, the average correlation was 0.87 ($p < 0.05$). Interestingly, looking at correlations between individual raters' and our diversity score, the expert annotator had the highest correlation with our diversity score in both settings.

Other commonly used diversity measures were more sensitive to the level of analysis. We correlated the raters diversity scores with the diversity scores computed by Shannon entropy and by the count of musical categories a user listened to ('volume'). In the genre-level analysis, both the entropy and volume methods showed significant correlation with the raters. The Pearson correlation between the raters' average scores and the entropy values was 0.95 ($p < 0.001$). The average correlation between raters and the volume measure was 0.86 ($p < 0.001$). However, in subgenre-level analysis we found more notable differences between the raters' and our diversity scores. The Pearson correlations between the entropy and the rater scores was 0.79 ($p < 0.05$). With volume, the average correlation was 0.46 ($p < 0.05$).

This result initially indicates that our diversity measure is promising as it captures human rater evaluations of diversity more robustly than traditional measures—it is less dependent on changes in categorical hierarchies. The distance between musical categories can be an important factor for understanding musical diversity, especially in highly complex musical classifications.

4.4.2 Correlates of Musical Diversity

To address RQ2, we used multiple regression analyses to examine factors associated with the diversity of musical consumption. We examined socioeconomic status variables as well as demographics, openness, and into-ness measures.

Table 4.2 presents the standardized coefficients of the explanatory variables¹⁰. The model (1) in Table 4.2 estimates the effects of socioeconomic, demographic, and other individual variables on the diversity of musical consumptions on genres. Among the ‘socioeconomic status’ variables, *High-profile News Reader* variable had a high coefficient due to users who follow *The Economist* or *The New Yorker* having higher musical diversity than those who do not (one-way ANOVA confirmed the significance; $p < 0.001$). Even though we exclude this variable to check whether income and education variables are associated with diversity of music consumption, we could not find any change regarding significance level and direction of correlation. The readers of high-profile news reports may have indirect or subtle difference in terms of socioeconomic status.

Racial Diversity positively associates with the diversity of music consumption. This may imply that people in our sample who live in more ethnically

¹⁰All variance inflation factors are below 1.64 ($\mu = 1.28$ and $\sigma = 0.16$); Pearson correlation between genre and subgenre diversities is 0.68 ($p < 0.001$).

	<i>Dependent variable:</i>	
	Genre (1)	Subgenre (2)
Income	-0.047 (0.037)	0.007 (0.037)
Education	0.027 (0.039)	-0.020 (0.039)
Racial Diversity	0.108** (0.036)	0.089* (0.036)
Urbanness	-0.040 (0.035)	0.052 (0.035)
High-profile News Reader	0.366*** (0.095)	0.301** (0.096)
Age	0.121*** (0.033)	0.161*** (0.033)
Gender (Male)	0.111· (0.067)	0.153* (0.067)
Music Event Attendance	-0.145*** (0.034)	-0.042 (0.034)
# of Loved Tracks	0.079* (0.033)	0.089** (0.033)
Days from Registration	-0.102** (0.033)	-0.029 (0.033)
# of Last.fm Friends	0.023 (0.036)	-0.081* (0.036)
Interest in Music	-0.143*** (0.034)	-0.113*** (0.034)
# of Twitter Friends	0.085* (0.033)	0.050 (0.034)
Friends' Timezone Diversity	0.026 (0.032)	0.074* (0.032)
Interest Diversity	-0.027 (0.032)	-0.017 (0.031)
Constant	-0.123* (0.057)	-0.143* (0.057)
Observations	1,014	1,014
R ²	0.101	0.087
Adjusted R ²	0.088	0.073
Residual Std. Error (df = 998)	0.955	0.963
F Statistic (df = 15; 998)	7.487***	6.322***

Note: · p<0.1; *p<0.05; **p<0.01; ***p<0.001

Table 4.2: Multiple regression coefficients of individual factors on the musical diversity of genre and subgenre.

diverse area are more likely to have higher musical diversity. By considering the relationship between white ratio and ethnic diversity, this result might be related to the effect of residential segregation. Both of *Age* and *Gender* in the 'demographic' variables have positive effect on diversity: being older or male is more likely to have more diverse musical tastes.

Among variables about 'into-ness,' *Musical Event Attendance* and *Days from Registration* appear to be negatively associated with diversity, whereas *Number of Last.fm Friends* does not show a significant relationship and *Number of Loved Tracks* appears to positively associated with diversity. *Number of Twitter Friends* as a 'openness' variable appears to be positively associated with diversity while *Timezone Diversity of Friends* and *Interest Diversity* shows no effect. On this basis, one could speculate that few variables within the same set of variables correlate with musical diversity in different directions. We discuss these trends below.

Model (2) in Table 4.2 estimates the effects of the same variables on the diversity of musical consumption of subgenres; it shows very similar trends with model (1). However, *Gender* is more significantly associated with diversity. Among the 'into-ness' variables, *Number of Last.fm Friends* is significantly associated with diversity rather than *Days from Registration*. But, the general trends of 'into-ness' are in common. Among the 'openness' variables *Timezone Diversity of Friends* is significantly associated with diversity rather than *Number of Twitter Friends* while the general trends of the 'openness' are in common.

4.5 Discussion

Our results provide initial evidence for the value of our ‘music diversity measure’ which aims to balance three qualities: variety, balance, and distance. Our diversity measure has shown to be more robust than other conventional measures such as volume and entropy.

Differences between Pearson correlation coefficients at the genre- and subgenre-levels computed by our measure, as well as the average rates assigned by independent coders on a 6-point Likert scale, were not significantly different. For the other measures of diversity, when moving between genre and subgenre levels, the average correlation coefficients dropped more steeply, especially the volume measure. Musical diversity can be computed by simple methods, but it may underestimate or overestimate diversity depending on the complexity of musical categories and the disparity between musical categories that people perceive. Our results show that volume and entropy might not be the best solution for computing the musical diversity of people on a highly complex map of musical categories such as subgenres.

We only considered the genre and subgenre categories, but new methods for music classification may result in categories that are even more complex, making a robust diversity measure even more important. For example, research efforts have developed novel methods for music classification using various data sources such as audio features and song metadata [61, 36].

In addition, diversity of music consumption was correlated with interest in high-profile news media; users who follow high-profile news media are much more likely to have a higher level of musical diversity. When we think about

whether one consumes high-profile news media, it is not necessarily a variable that is as straightforward as income or education level. To understand news reports, readers need more than a basic grasp of word order and word meaning; a particular 'knowledge of the world' is also necessary. Van Dijk [130] explains this when he writes: "Readers of a news report first of all need to understand its words, sentences, or the structural properties. This does not only mean they must know the language and its grammar and lexicon, possibly including rather technical words such as those of modern politics, management, science, or the professions. Users of the media need to know something about the specific organization and functions of news reports in the press, including the functions of headlines, leads, background information, or quotations. Besides such grammatical and textual knowledge, media users need vast amounts of properly organized knowledge of the world." Van Dijk's point alludes to the possibility that if one has access to particular understandings of 'the world,' then they are better equipped to seek out and benefit from high profile news sources. If this is the case, then we can begin to think about level of music diversity as a potential variable vis-à-vis knowledge.

Our results also confirm a number of previous findings about demographic variables associated with the diversity of music consumption. They show that male users are more likely to have diverse musical tastes, which confirms prior research showing that males tend to consider mainstream music as *unhip* while females consider it in another way of saying *popular* music [23]; such perceptions might affect musical consumption. Males are also more likely to prefer more unique styles of music than females [107]. In addition, people who are older in our sample are more likely to have diverse musical tastes. This result closely echoes the analyses of [134]: young people may identify strongly with one or

only a few genres and styles of music, which reveals the significance of their representational dimensions.

A potentially surprising finding is that people who attended more musical events are less likely to have diverse listening habits. [134] also argues that there is a tendency, or an openness, towards unfamiliar musical forms and evidence of relatively diverse tastes in people who are limited in how they can engage in musical activities. The development of a broad palette of musical tastes was not valued by people for whom music is more accessible. We note that urban dwellers may have better access to musical activities, but we could not find a significant association with urbanness in our results.

Users with diverse patterns of music consumption are less likely to follow music-related accounts on Twitter. This finding can be due to a different set of music-related interests between diverse listeners—who care more about the music itself—and more casual music fans who may care more about the celebrity factor. If this were shown to be true, we may refer to it as the Justin Bieber effect (no offense to his fans should they be reading this paper).

4.6 Final Remarks

In this paper, we have designed a reasonable measure that quantifies the diversity of musical tastes. In addition, we provide an analysis of diversity as it relates to the cultural omnivore thesis. Based on well-known individual factors which relate diversity of musical preferences across various theoretical work and empirical studies, we identified key factors for designing a diversity measure, and located individual-level variables for exploring correlations of musical

diversity.

We acknowledge that the manner in which we inferred the socioeconomic status variable could produce significant inaccuracies. For example, users' home locations were inferred in ZIP code resolution and using geocoded Twitter data. These methods are prone to error. Other methods for collecting more direct or fine-grained location data, or maybe even a direct collection of socioeconomic variables, might give us a better opportunity to study this correlation with music consumption. Second, our user population and the music they listen to are both potentially highly biased. Our population is comprised of users who make an explicit connection between their Twitter and Last.fm accounts, which may indicate search biases on our behalf. In addition, the tracks and artists displayed for each user are based on their public listening behavior, which may or may not be reflective of their overall listening habits. Finally, we could see rating differences among coders due to knowledge differences. The measurement validations can be improved by better systematic investigations using more listening history samples and annotators with different levels of knowledge background. At the same time, it would be interesting to see if the way of rating changes when the music listeners themselves are asked about their diversity.

Future research along this vein will provide a richer and more complex picture of musical preferences. This picture will in turn contribute to a greater understanding of the changing face of the cultural omnivore, as it manifests through analyses of social media data, and also contribute to a empirical recommendation system aiming to provide contents based on tastes and aesthetics preferences.

CHAPTER 5

CONCLUSIONS

Instead of reviewing what has already been covered in the previous chapters, we conclude by looking forward and describing some possible next research projects. These projects are all in some way to the ideas presented in this dissertation and are all very preliminary.

5.1 Mood Management through Cultural Consumption

We have observed the surprisingly consistent patterns of affective preference over time of the day and day of the week and across seasons (Chapter 2): people prefer relaxing music over the sleeping time and arousing music during working hours. By comparing the diurnal and seasonal patterns of affective preference to affective expression (i.e., emotional status), we could also speculate people use music both for reflecting (or reinforcing) and shaping (or altering) mood. However, it is still unclear when people use music or other external stimuli to influence emotions.

The congruence and disparity between affective preference and expression in Chapter 2 leads to a hypothesis that people use music to reinforce their **good** moods and to alter their **bad** moods. One way to address this question might be comparing the dynamics of physiological arousal and affective preference. Using smartphone apps or wearable sensors, one can continuously monitor a marker of physiological arousal (e.g., body temperature, heart rate, and heart rate variability) and this physiological arousal can be matched to music con-

sumption patterns. These two dynamics should allow a statistical test on the causal directions of music consumption for mood management using a regression discontinuity at or around the point where a certain mood drops under a specific threshold or continuously decreases over a particular period of time.

Also, previous research has shown that people often use other media (e.g., TV and books) for mood management [125]. Thus, it should also be interesting to see whether people have a similar strategy to use other media and those media have similar effects on emotions.

5.2 Diurnal and Seasonal Patterns of Variety-seeking

Do other consumption behaviors vary by time of the day and across seasons? People may have a different preference (or at least make different choices) at a different time of the day and different season of the year due to varying levels of affective preference as we have seen in Chapter 2.

A type of behavior that affective preference may operate on might be variety-seeking (or even novelty-seeking). Like prior work in affective preference, previous research in variety-seeking has mainly focused on individual differences [2] and situational factors [75] rather than within-individual changes [116]. Given the link between stimulation and variety (e.g., choosing various items feels stimulating [30, 65]), we can easily speculate that people may also show particular patterns of diurnal and seasonal patterns of variety-seeking as a response to both psychophysiological and social needs like the case in music consumption in Chapter 2.

It might be even more interesting if we can decompose when people make choices as a response to psychophysiology process and when people do it as a social process.

5.3 Diffusion of Cultural Interests

Cultural consumption operates at the transnational level. However, current efforts are often criticized for being overly simplistic in presenting cultural flows as a simple one-way flow from the core to peripheral countries (reflecting world-systems theory) [78, 122, 94]. In light of this, a more generalizable analytical model is required to provide a useful way of seeing how cultural interests flow across different cultures and are interwoven with everyday life. To approach this problem, I propose to collect a set of longitudinal lists of popular media contents (e.g., videos from YouTube and popular topics from Google Trends) across many countries. Since existing tree-based and dyadic-level of network approaches are limited to draw a macro-level structure of the flow due to high connectivity and temporal sparsity of country-to-country relationships derived from such real-world diffusion data, to model trajectories of cultural interests across countries representation learning methods (e.g., Word2Vec) might be very useful.

One challenge might be addressing mechanisms beyond a descriptive outcome based on a visualization of embeddings. With such new techniques, however, we can still ask some important questions not only like whether cultural interests flow from more prestigious countries to less prestigious countries but also like whether such flow interacts with a network position of a country.

5.4 Effects of Collective Traces and Recommendations

Although the consumer industry enables people to benefit from recommender systems and collective traces like comments and view counts in some domains (e.g., choosing a video to play), neither algorithms nor collective preferences are available for many everyday decisions. It remains unclear how individuals can best leverage the experience of others when they consume a product and share prior experience about the available options with others (from a relatively small community of peers to a large pool of unknown people). Should they use strategies that aggregate the opinions of many individuals, or is it better to rely on the opinions of just a few similar network neighbors? Can we also apply insights from this kind of work in an observational study? These questions are particularly relevant in the era where human-data interaction is nearly ubiquitous in practice.

5.5 Data Science Tools for Large-scale and High-dimensional Observational Data

Throughout this dissertation, we have constantly encountered instances of analyzing large-scale activity and relational datasets to extract insights on human behaviors. Due to the observational nature of the data, we had to pay special attention to rule out key confounding factors and alternative explanations. Robustness checks involve various tasks and techniques, from manual validation and experiment (Chapter 4) to matching techniques and leveraging natural experiments (Chapter 2).

Due to the high cost and limited scope of randomized controlled trials, large-scale observational studies should be continuously necessary. However, since the standards of validity are going to be higher, we need tools or guidelines that help researchers produce high-quality predictions and inferences from observational data. For example, such tools should inform potential confounding factors whenever possible and help establish causal relationships based on data and variables. Specific challenges may include identifying generalizable patterns and handling different types of variables (from binary to high-dimensional data like language).

5.6 Inference of Sociodemographic Attributes of Online Users

While extensive behavioral and relational data are available through various online sources such as social media and streaming services, previous research has an important limitation: lack of information about the sociodemographic characteristics of individual users, often covariates for behavioral and relational measures.

In this dissertation (Chapter 4 in particular), we have addressed this challenge by using geocoded traces and connecting an individual user's multiple social media profiles and activities. For example, we could infer a user's income level by identifying a plausible home location as a neighborhood level and matching the home location to Census records. Also, we could identify the socioeconomic status from one's geocoded tweets while preserving self-disclosed demographic information from their corresponding Last.fm profiles.

However, such an approach can be further improved to collect more gran-

ular data at a greater scale. For example, we can use image-processing based applications (e.g., Face++ [32]) to infer demographic information such as age, gender, and ethnicity. Since the performance of such an application can be substandard in a certain context (e.g., it may work well for Asian vs. Black but may not work well for White vs. Hispanic), we may need to incorporate the application's output with other sociodemographic markers like language usages and names to improve the performance of the classifications.

Also, inferring socioeconomic status can be more challenging. A notable example of a potential solution is identifying a home location based on the most frequent area of activity and then take a housing price as a proxy of income level [7]. An issue in this approach is that it is hard to distinguish whether the home is owned or rented. Also, since this approach pinpoints an individual's plausible home location, it may raise some ethical concerns. Despite these limitations, this approach provides a reasonable estimate of the tercile of the income distribution each individual belongs to.

5.7 Final Thoughts

We need to work out forms of observations appropriate for our new cultural conditions where substantial quantitative and qualitative changes in cultural production and participation happened worldwide over a couple of decades. Because cultural technologies now change with faster speed, we need to keep looking for new forms and not be satisfied with what we found yesterday.

As can be seen, by the proceeding discussion, this dissertation can lead to many different research directions. We hope this dissertation can be an inspira-

tion to some people and we can see more studies in this field.

APPENDIX A

SUPPLEMENTARY INFORMATION OF CHAPTER 2

(GLOBAL MUSIC STREAMING DATA REVEAL DIURNAL AND SEASONAL PATTERNS OF AFFECTIVE PREFERENCE)

Country Name	Musical Intensity Baseline	Female				Male				Total
		Age 13-24	Age 25-54	Age 55-64	Age Over 65	Age 13-24	Age 25-54	Age 55-64	Age Over 65	
United States	0.82831512	3.17	4.00	1.36	1.76	3.32	4.01	1.27	1.41	20.31
Brazil	1.013215721	1.72	2.85	0.53	0.37	1.49	2.80	0.56	0.46	10.78
Mexico	0.999635427	1.18	1.63	0.32	0.30	1.06	1.53	0.27	0.25	6.54
Indonesia	0.744305192	1.50	1.01	0.01	0.05	1.49	2.07	0.02	0.06	6.20
Germany	0.857832623	0.56	1.00	0.37	0.56	0.59	1.02	0.36	0.49	4.95
Philippines	0.680104135	0.89	1.18	0.12	0.13	0.83	1.22	0.14	0.12	4.63
United Kingdom	0.771448609	0.58	0.80	0.24	0.37	0.61	0.83	0.24	0.33	4.00
France	0.767174886	0.62	0.78	0.27	0.14	0.65	0.79	0.25	0.31	3.81
Italy	0.789425956	0.44	0.82	0.19	0.14	0.46	0.80	0.24	0.28	3.39
Spain	0.875204246	0.37	0.67	0.19	0.31	0.39	0.70	0.18	0.23	3.04
Turkey	0.59022227	0.46	0.67	0.01	0.01	0.48	1.10	0.03	0.03	2.80
Argentina	1.107293222	0.52	0.54	0.13	0.19	0.38	0.54	0.12	0.13	2.55
Colombia	1.026238443	0.35	0.63	0.08	0.07	0.33	0.62	0.13	0.09	2.30
Canada	0.778978523	0.29	0.44	0.16	0.23	0.31	0.45	0.15	0.18	2.21
Poland	0.750765111	0.30	0.51	0.03	0.04	0.31	0.53	0.09	0.09	1.90
Japan	0.868370632	0.29	0.37	0.01	0.00	0.38	0.76	0.05	0.02	1.88
Peru	1.12687843	0.28	0.40	0.06	0.04	0.24	0.37	0.07	0.07	1.54
Australia	0.784043711	0.22	0.29	0.09	0.13	0.23	0.30	0.08	0.11	1.44
Malaysia	0.643098991	0.23	0.40	0.02	0.03	0.22	0.40	0.04	0.04	1.36
Chile	1.088352036	0.19	0.24	0.06	0.07	0.18	0.24	0.06	0.05	1.09
Netherlands	0.678077343	0.15	0.21	0.07	0.11	0.15	0.21	0.07	0.09	1.06
Taiwan	0.354642144	0.10	0.34	0.01	0.01	0.11	0.34	0.02	0.01	0.94
Ecuador	1.083445667	0.13	0.21	0.02	0.02	0.12	0.20	0.04	0.03	0.75
Belgium	0.636552148	0.10	0.14	0.05	0.07	0.10	0.14	0.05	0.06	0.71
Guatemala	1.069761795	0.13	0.17	0.01	0.01	0.12	0.16	0.02	0.01	0.63
Sweden	0.831712202	0.09	0.12	0.04	0.07	0.09	0.12	0.04	0.06	0.62
Czech Republic	0.787927426	0.08	0.14	0.01	0.02	0.08	0.15	0.03	0.04	0.56
Portugal	0.741288681	0.09	0.14	0.01	0.02	0.08	0.14	0.03	0.03	0.54
Austria	0.793222967	0.07	0.12	0.02	0.04	0.07	0.12	0.04	0.05	0.51
Switzerland	0.710010378	0.06	0.11	0.03	0.05	0.07	0.11	0.03	0.04	0.51
Dominican Republic	1.058202697	0.09	0.13	0.01	0.01	0.08	0.14	0.02	0.02	0.49
Hungary	0.838234468	0.07	0.13	0.01	0.01	0.06	0.13	0.03	0.02	0.46
Bolivia	1.086961984	0.08	0.13	0.01	0.01	0.08	0.13	0.01	0.01	0.45
Greece	0.5697184	0.05	0.14	0.01	0.01	0.05	0.14	0.02	0.02	0.44
Honduras	1.10194196	0.08	0.10	0.01	0.01	0.07	0.10	0.01	0.01	0.38
Singapore	0.633801547	0.05	0.09	0.02	0.02	0.05	0.09	0.02	0.02	0.37
Denmark	0.840295757	0.05	0.07	0.02	0.04	0.05	0.07	0.02	0.03	0.35
Norway	0.860553202	0.05	0.07	0.02	0.03	0.05	0.07	0.02	0.03	0.33
Paraguay	1.181182803	0.06	0.09	0.01	0.00	0.06	0.09	0.01	0.01	0.32
Ireland	0.726392631	0.05	0.07	0.02	0.02	0.05	0.07	0.02	0.02	0.31
Costa Rica	1.026703782	0.06	0.07	0.01	0.01	0.06	0.07	0.01	0.01	0.31
Finland	0.898012516	0.05	0.06	0.02	0.01	0.05	0.07	0.02	0.02	0.30
El Salvador	1.039898858	0.06	0.08	0.01	0.01	0.05	0.07	0.01	0.01	0.29
New Zealand	0.858512316	0.04	0.06	0.02	0.02	0.05	0.06	0.02	0.02	0.28
Bulgaria	0.835160736	0.03	0.09	0.00	0.01	0.03	0.10	0.00	0.01	0.27
Nicaragua	1.039331207	0.05	0.08	0.00	0.00	0.04	0.07	0.01	0.00	0.26
Slovakia	0.757090092	0.04	0.08	0.00	0.00	0.03	0.08	0.01	0.01	0.24
Panama	1.106890677	0.04	0.05	0.01	0.01	0.03	0.05	0.01	0.01	0.20
Hong Kong	0.50869091	0.02	0.05	0.00	0.00	0.03	0.07	0.00	0.00	0.18
Lithuania	0.691073248	0.02	0.04	0.00	0.01	0.02	0.03	0.01	0.01	0.13
Latvia	0.727272948	0.01	0.03	0.00	0.00	0.02	0.03	0.01	0.01	0.10

Table A.1: Detailed statistics related to baseline comparisons in musical intensity. We report additional information including exact probability values, degrees of freedom, confidence intervals, and effect sizes for baseline comparisons between groups to support our results related to the diurnal and seasonal patterns of affective preference in musical intensity based on temporal music consumption of one million Spotify users over a year. We performed all the tests using Welch's two-sample t-test (two-sided) to correct for unequal size and variance between paired samples.

Country Name	Musical Intensity Baseline	Female				Male				Total
		Age 13-24	Age 25-54	Age 55-64	Age Over 65	Age 13-24	Age 25-54	Age 55-64	Age Over 65	
United States	0.82831512	3.17	4.00	1.36	1.76	3.32	4.01	1.27	1.41	20.31
Brazil	1.013215721	1.72	2.85	0.53	0.37	1.49	2.80	0.56	0.46	10.78
Mexico	0.999635427	1.18	1.63	0.32	0.30	1.06	1.53	0.27	0.25	6.54
Indonesia	0.744305192	1.50	1.01	0.01	0.05	1.49	2.07	0.02	0.06	6.20
Germany	0.857832623	0.56	1.00	0.37	0.56	0.59	1.02	0.36	0.49	4.95
Philippines	0.680104135	0.89	1.18	0.12	0.13	0.83	1.22	0.14	0.12	4.63
United Kingdom	0.771448609	0.58	0.80	0.24	0.37	0.61	0.83	0.24	0.33	4.00
France	0.767174886	0.62	0.78	0.27	0.14	0.65	0.79	0.25	0.31	3.81
Italy	0.789425956	0.44	0.82	0.19	0.14	0.46	0.80	0.24	0.28	3.39
Spain	0.875204246	0.37	0.67	0.19	0.31	0.39	0.70	0.18	0.23	3.04
Turkey	0.59092227	0.46	0.67	0.01	0.01	0.48	1.10	0.03	0.03	2.80
Argentina	1.107293222	0.52	0.54	0.13	0.19	0.38	0.54	0.12	0.13	2.55
Colombia	1.026238443	0.35	0.63	0.08	0.07	0.33	0.62	0.13	0.09	2.30
Canada	0.778978523	0.29	0.44	0.16	0.23	0.31	0.45	0.15	0.18	2.21
Poland	0.750765111	0.30	0.51	0.03	0.04	0.31	0.53	0.09	0.09	1.90
Japan	0.868370632	0.29	0.37	0.01	0.00	0.38	0.76	0.05	0.02	1.88
Peru	1.126667843	0.28	0.40	0.06	0.04	0.24	0.37	0.07	0.07	1.54
Australia	0.784043711	0.22	0.29	0.09	0.13	0.23	0.30	0.08	0.11	1.44
Malaysia	0.643309891	0.23	0.40	0.02	0.03	0.22	0.40	0.04	0.04	1.36
Chile	1.088352036	0.19	0.24	0.06	0.07	0.18	0.24	0.06	0.05	1.09
Netherlands	0.678077343	0.15	0.21	0.07	0.11	0.15	0.21	0.07	0.09	1.06
Taiwan	0.354642144	0.10	0.34	0.01	0.01	0.11	0.34	0.02	0.01	0.94
Ecuador	1.083445667	0.13	0.21	0.02	0.02	0.12	0.20	0.04	0.03	0.75
Belgium	0.636552148	0.10	0.14	0.05	0.07	0.10	0.14	0.05	0.06	0.71
Guatemala	1.069761795	0.13	0.17	0.01	0.01	0.12	0.16	0.02	0.01	0.63
Sweden	0.831712202	0.09	0.12	0.04	0.07	0.09	0.12	0.04	0.06	0.62
Czech Republic	0.787927426	0.08	0.14	0.01	0.02	0.08	0.15	0.03	0.04	0.56
Portugal	0.741288681	0.09	0.14	0.01	0.02	0.08	0.14	0.03	0.03	0.54
Austria	0.793222967	0.07	0.12	0.02	0.04	0.07	0.12	0.04	0.05	0.51
Switzerland	0.710010378	0.06	0.11	0.03	0.05	0.07	0.11	0.03	0.04	0.51
Dominican Republic	1.058202697	0.09	0.13	0.01	0.01	0.08	0.14	0.02	0.02	0.49
Hungary	0.838234468	0.07	0.13	0.01	0.01	0.06	0.13	0.03	0.02	0.46
Bolivia	1.086961984	0.08	0.13	0.01	0.01	0.08	0.13	0.01	0.01	0.45
Greece	0.5697184	0.05	0.14	0.01	0.01	0.05	0.14	0.02	0.02	0.44
Honduras	1.101194196	0.08	0.10	0.01	0.01	0.07	0.10	0.01	0.01	0.38
Singapore	0.633801547	0.05	0.09	0.02	0.02	0.05	0.09	0.02	0.02	0.37
Denmark	0.840295757	0.05	0.07	0.02	0.04	0.05	0.07	0.02	0.03	0.35
Norway	0.860553202	0.05	0.07	0.02	0.03	0.05	0.07	0.02	0.03	0.33
Paraguay	1.181182803	0.06	0.09	0.01	0.00	0.06	0.09	0.01	0.01	0.32
Ireland	0.726392631	0.05	0.07	0.02	0.02	0.05	0.07	0.02	0.02	0.31
Costa Rica	1.026703782	0.06	0.07	0.01	0.01	0.06	0.07	0.01	0.01	0.31
Finland	0.898012516	0.05	0.06	0.02	0.01	0.05	0.07	0.02	0.02	0.30
El Salvador	1.039898858	0.06	0.08	0.01	0.01	0.05	0.07	0.01	0.01	0.29
New Zealand	0.858512316	0.04	0.06	0.02	0.02	0.05	0.06	0.02	0.02	0.28
Bulgaria	0.835160736	0.03	0.09	0.00	0.01	0.03	0.10	0.00	0.01	0.27
Nicaragua	1.039331207	0.05	0.08	0.00	0.00	0.04	0.07	0.01	0.00	0.26
Slovakia	0.757090092	0.04	0.08	0.00	0.00	0.03	0.08	0.01	0.01	0.24
Panama	1.106890677	0.04	0.05	0.01	0.01	0.03	0.05	0.01	0.01	0.20
Hong Kong	0.508569091	0.02	0.05	0.00	0.00	0.03	0.07	0.00	0.00	0.18
Lithuania	0.691073248	0.02	0.04	0.00	0.01	0.02	0.03	0.01	0.01	0.13
Latvia	0.727272948	0.01	0.03	0.00	0.00	0.02	0.03	0.01	0.01	0.10

Table A.2: Proportion of sample in each demographic group from each of the 51 countries based on the World Factbook. Countries are ordered by the proportion of sampled users, which reflects the country's relative demographic distributions compared to the world population distribution, not Spotify's user distribution over the globe. Although the population distribution in the World Factbook breaks populations under age of 25 in age 0–14 and age 15–24, we merged the two age groups (0–14 and 15–24) into one (13–24) as users need to be 13 or older to sign up for Spotify.

Category	Group x	Group y	Mean x	Mean y	Statistic	P-value	N x	N y	DF	Conf. Low	Conf. High	Method	Alternative	Cohen's D	
Day of Week	Tue	Mon	0.835435109	0.827878454	39.44044937	0	105868141	103532851	209247088.4	0.007181132	0.007932177	Welch Two Sample t-test	two_sided	0.005451711	
	Wed	Mon	0.842948079	0.827878454	79.12743446	0	107681991	103532851	210731017	0.014696355	0.015442895	Welch Two Sample t-test	two_sided	0.010893119	
	Thu	Mon	0.852416879	0.827878454	129.6478357	0	109588622	103532851	212095776.5	0.024167463	0.024909387	Welch Two Sample t-test	two_sided	0.01777517	
	Mon	Fri	0.827878454	0.879186303	-276.8819033	0	103532851	115918607	215066876	-0.051671041	-0.050944655	Welch Two Sample t-test	two_sided	0.037504131	
	Mon	Sat	0.827878454	0.88253005	-296.8924007	0	103532851	118640317	216001423.8	-0.056012383	-0.054290807	Welch Two Sample t-test	two_sided	0.040018461	
	Mon	Sun	0.827878454	0.82015569	39.99497199	0	103532851	103762231	207294683.7	0.007344309	0.008101221	Welch Two Sample t-test	two_sided	0.00555725	
	Tue	Wed	0.835435109	0.842948079	-39.77168249	0	105868141	107681991	213457812.4	-0.007883213	-0.007142728	Welch Two Sample t-test	two_sided	0.005443569	
	Tue	Thu	0.835435109	0.852416879	-90.46565261	0	105868141	109588622	215096995.8	-0.017349685	-0.016613855	Welch Two Sample t-test	two_sided	0.012329843	
	Tue	Sat	0.835435109	0.88253005	-258.0829664	0	105868141	118640317	220074976.1	-0.047452594	-0.046737287	Welch Two Sample t-test	two_sided	0.034559718	
	Tue	Sun	0.835435109	0.82015569	79.75905997	0	105868141	103762231	209491600.4	0.014903949	0.015654889	Welch Two Sample t-test	two_sided	0.011018677	
	Tue	Fri	0.835435109	0.879186303	-238.1436906	0	105868141	115918607	218857960.8	-0.044111274	-0.043391114	Welch Two Sample t-test	two_sided	0.032050873	
	Wed	Thu	0.842948079	0.852416879	-50.75956349	0	107681991	109588622	217169032.1	-0.009834416	-0.009103184	Welch Two Sample t-test	two_sided	0.006887786	
	Wed	Sat	0.842948079	0.88253005	-218.3555197	0	107681991	118640317	223048550.1	-0.039937259	-0.039226681	Welch Two Sample t-test	two_sided	0.029097133	
	Wed	Sun	0.842948079	0.82015569	119.6950333	0	107681991	103762231	210987372.2	0.022419172	0.023165607	Welch Two Sample t-test	two_sided	0.016468866	
	Wed	Fri	0.842948079	0.879186303	-198.5449269	0	107681991	115918607	221605010.2	-0.036595954	-0.035880402	Welch Two Sample t-test	two_sided	0.026594166	
	Thu	Sat	0.852416879	0.88253005	-167.2625504	0	109588622	118640317	225990122.7	-0.030466033	-0.029760308	Welch Two Sample t-test	two_sided	0.022178532	
	Thu	Sun	0.852416879	0.82015569	170.4752612	0	109588622	103762231	212364326.8	0.031890281	0.032632008	Welch Two Sample t-test	two_sided	0.023390666	
	Thu	Fri	0.852416879	0.879186303	-147.6609376	0	109588622	115918607	224302855.6	-0.027124745	-0.026414102	Welch Two Sample t-test	two_sided	0.019683184	
	Sat	Fri	0.88253005	0.879186303	191.0332177	1.32E-80	118640317	115918607	234397272	0.002999061	0.003688443	Welch Two Sample t-test	two_sided	0.002483171	
	Sun	Fri	0.82015569	0.879186303	-318.6054193	0	103762231	115918607	215374929.8	-0.059939375	-0.058667474	Welch Two Sample t-test	two_sided	0.043131011	
	Sat	Sun	0.88253005	0.82015569	338.8974547	0	118640317	103762231	216322887.8	0.062013627	0.062735093	Welch Two Sample t-test	two_sided	0.04565248	
	Region	Europe	North America	0.804062884	0.829946942	-12.58603664	2.55164E-36	282710	223258	511069.129	-0.029914869	-0.021853246	Welch Two Sample t-test	two_sided	0.034671667
		Europe	Asia	0.804062884	0.697934455	50.0626565	0	282710	181845	442405.1465	0.102022674	0.110234184	Welch Two Sample t-test	two_sided	0.144901421
		Europe	Latin America	0.804062884	1.053349994	-134.9161023	0	282710	286146	574690.2806	-0.252908581	-0.245665638	Welch Two Sample t-test	two_sided	0.350555599
Europe		Oceania	0.804062884	0.806838996	-0.477686961	0.63287825	282710	17076	1998.46569	-0.014167271	0.008615048	Welch Two Sample t-test	two_sided	0.003617257	
North America		Asia	0.829946942	0.697934455	61.5257519	0	223258	181845	414322.599	0.127807085	0.136217888	Welch Two Sample t-test	two_sided	0.189194313	
North America		Latin America	0.829946942	1.053349994	-117.2618022	0	223258	286146	472013.5272	-0.227137115	-0.219688989	Welch Two Sample t-test	two_sided	0.328400013	
North America		Oceania	0.829946942	0.806838996	3.963566368	7.40918E-05	223258	17076	20239.72624	0.011680498	0.034535395	Welch Two Sample t-test	two_sided	0.032095317	
Asia		Latin America	0.697934455	1.053349994	-182.6035716	0	181845	286146	393897.3981	-0.359230361	-0.351600696	Welch Two Sample t-test	two_sided	0.539446448	
Asia		Oceania	0.697934455	0.806838996	-18.63627255	7.08663E-77	181845	17076	20426.15723	-0.120358635	-0.097450446	Welch Two Sample t-test	two_sided	0.160254452	
Latin America		Oceania	1.053349994	0.806838996	42.79902738	0	286146	17076	19285.2554	0.235221419	0.257800577	Welch Two Sample t-test	two_sided	0.375845241	
Gender		F	M	0.84409257	0.880868431	-26.04153712	1.8755E-149	487251	503784	1033791.987	-0.039532973	-0.033998749	Welch Two Sample t-test	two_sided	0.051209544
		20-29	10-19	0.969895324	1.162041457	-129.3459251	0	314448	216634	517818.7169	-0.195057708	-0.189234558	Welch Two Sample t-test	two_sided	0.347446522
	30-39	10-19	0.840987933	1.162041457	-172.3417831	0	192220	216634	342105.8799	-0.324704731	-0.317402316	Welch Two Sample t-test	two_sided	0.550737547	
	10-19	40-49	1.162041457	0.786824678	145.4267506	0	216634	82349	111843.6606	0.388114515	0.398719044	Welch Two Sample t-test	two_sided	0.706396603	
	10-19	Over 50	1.162041457	0.483620946	316.6248121	0	216634	228303	354180.8696	0.674220953	0.682620069	Welch Two Sample t-test	two_sided	0.93054685	
	20-29	30-39	0.969895324	0.840987933	68.61353999	0	314448	192220	366814.2859	0.125255105	0.132589677	Welch Two Sample t-test	two_sided	0.204875392	
	20-29	40-49	0.969895324	0.786824678	74.10005128	0	314448	82349	113959.3449	0.195946934	0.20659436	Welch Two Sample t-test	two_sided	0.322678883	
	20-29	Over 50	0.969895324	0.483620946	225.4961055	0	314448	228303	371121.5471	0.482047771	0.490509895	Welch Two Sample t-test	two_sided	0.659180175	
	30-39	40-49	0.840987933	0.786824678	24.61658685	1.559E-133	192220	82349	148325.8722	0.066601672	0.07812484	Welch Two Sample t-test	two_sided	0.104803258	
	30-39	Over 50	0.840987933	0.483620946	146.9524136	0	192220	228303	416015.0847	0.352606025	0.36213335	Welch Two Sample t-test	two_sided	0.444448678	
	40-49	Over 50	0.786824678	0.483620946	91.21820439	0	82349	228303	180573.541	0.278879945	0.291127518	Welch Two Sample t-test	two_sided	0.334383127	
	Chronotype	Evening	Afternoon	0.902819466	0.861249247	26.69846012	5.77E-157	347369	444035	784687.5074	0.038518498	0.04462194	Welch Two Sample t-test	two_sided	0.059063892
Evening		Morning	0.902819466	0.833943147	31.99292133	3.39E-224	347369	150011	288846.9766	0.064656771	0.073095867	Welch Two Sample t-test	two_sided	0.099042958	
Evening		Night Owl	0.902819466	0.683628582	53.60016987	0	347369	49620	60904.8013	0.211175709	0.227209606	Welch Two Sample t-test	two_sided	0.30329535	
Afternoon		Morning	0.861249247	0.833943147	13.03924483	7.52E-39	444035	150011	267198.8545	0.022301628	0.031410572	Welch Two Sample t-test	two_sided	0.038234742	
Afternoon		Night Owl	0.861249247	0.683628582	43.76223225	0	444035	49620	59132.83571	0.169665468	0.185575864	Welch Two Sample t-test	two_sided	0.243160234	
Morning		Night Owl	0.833943147	0.683628582	34.77486272	7.33E-263	150011	49620	74974.20983	0.141842472	0.158786659	Welch Two Sample t-test	two_sided	0.198457238	
Hemisphere	Northern	Southern	0.824294572	0.975070307	-99.22659784	0	739251	251784	505996.9196	-0.154365778	-0.148385694	Welch Two Sample t-test	two_sided	0.211668974	

Table A.3: Country names included and excluded for analyses. Countries collected from YouTube in this study.

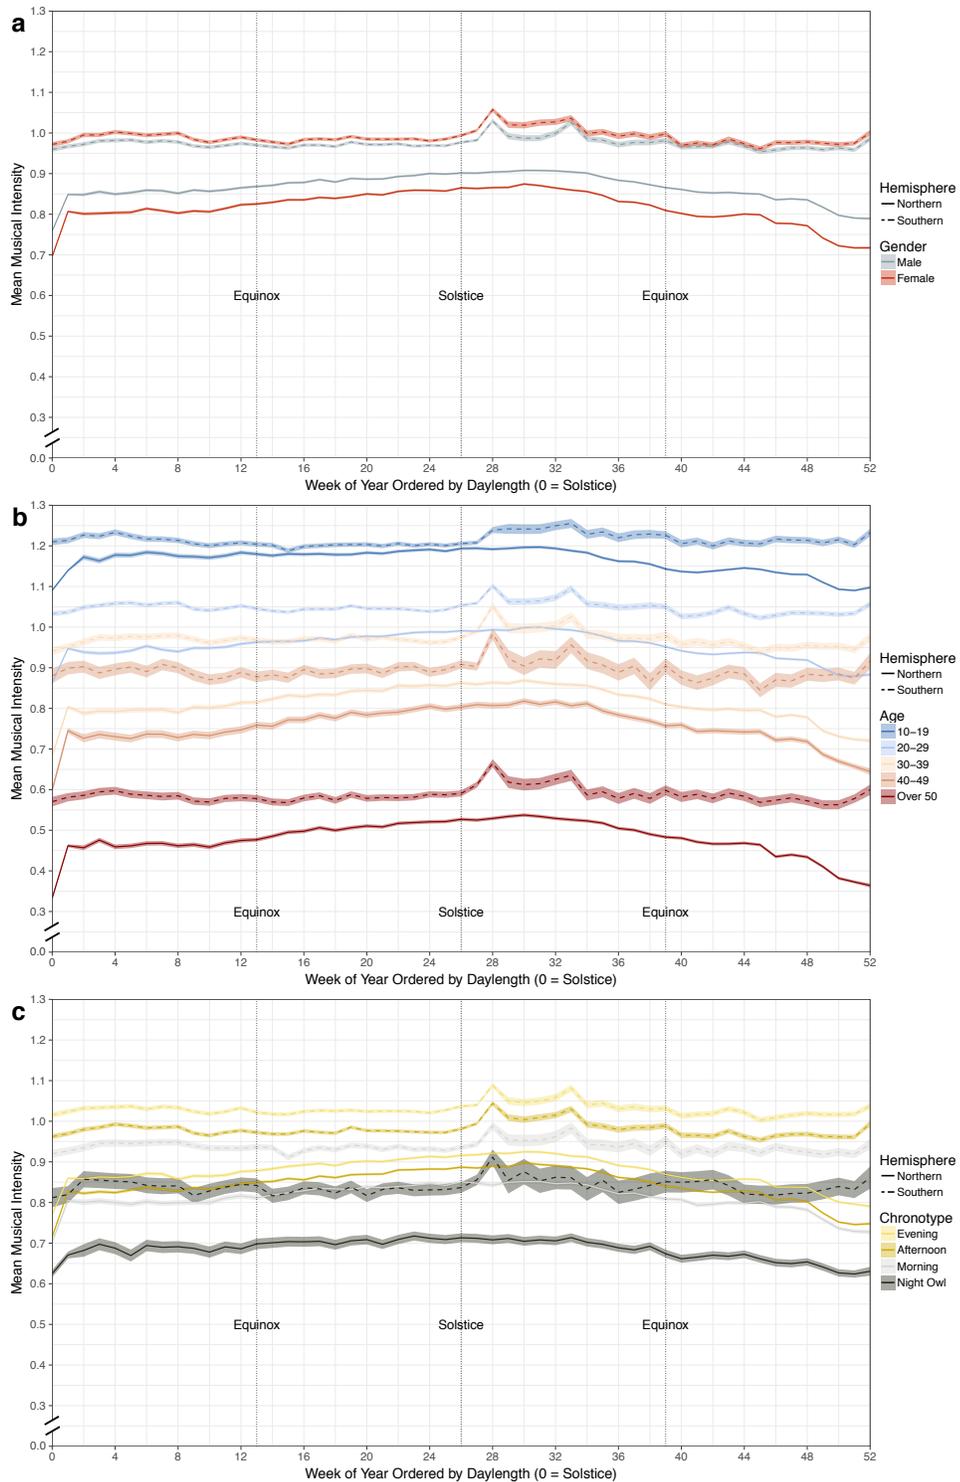


Figure A.1: Seasonal variations in affective preferences revealed by music streaming exhibit robust patterns across different geographic regions and user groups.

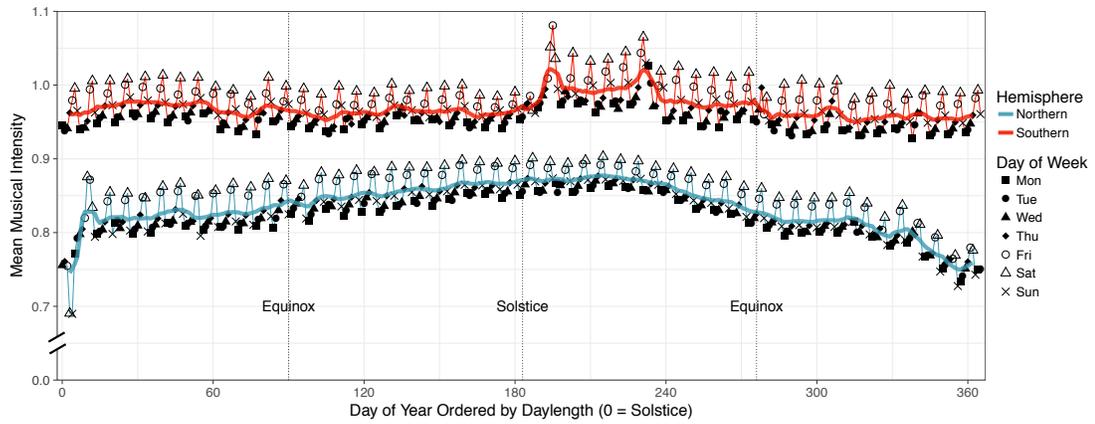


Figure A.2: Affective preferences revealed through music selection varies from month to month, but daily differences are seasonally robust.

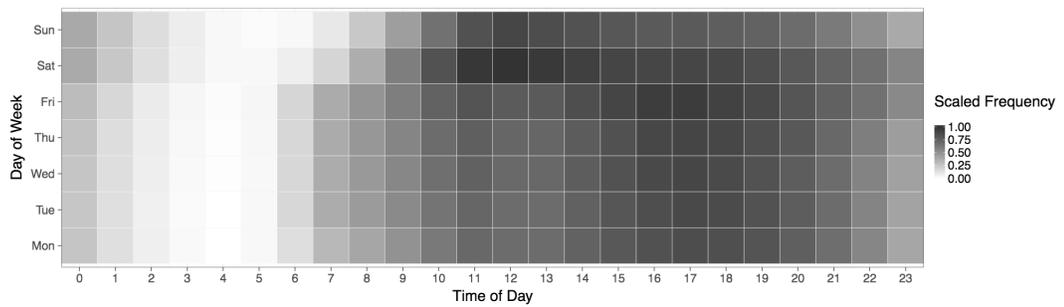


Figure A.3: The hourly distribution of plays from global music streaming data.

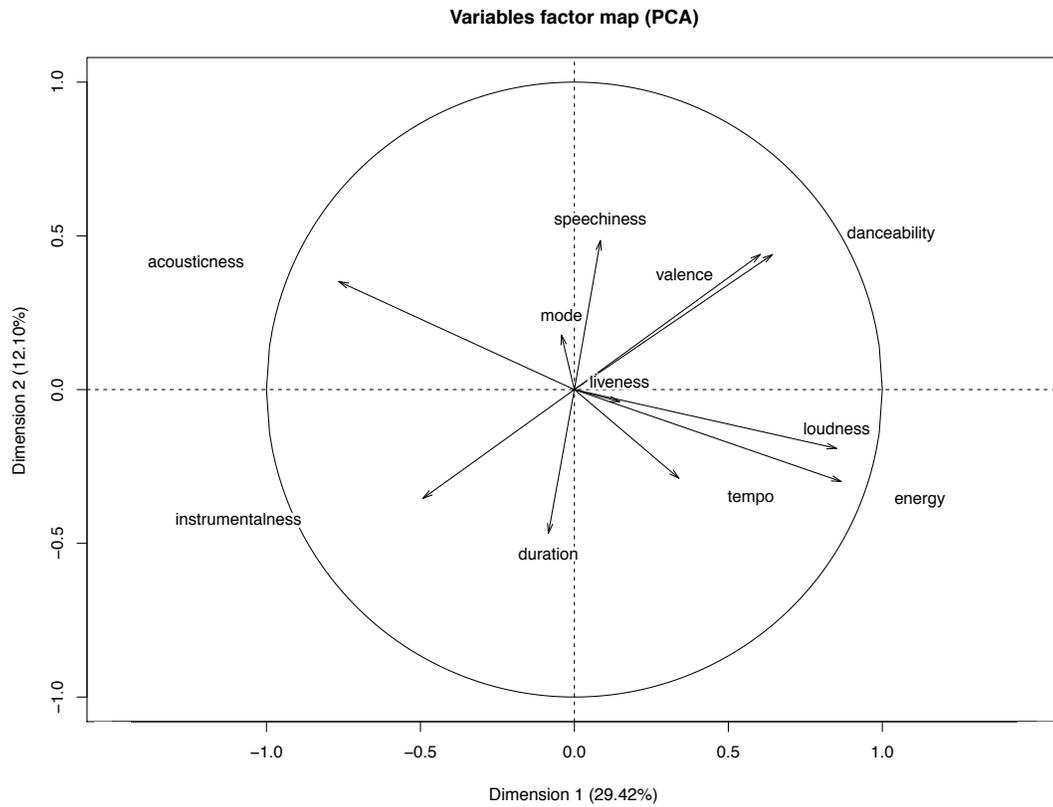


Figure A.4: Principal component analysis (PCA) of 11 musical attributes identified a first principal component corresponding to musical intensity.

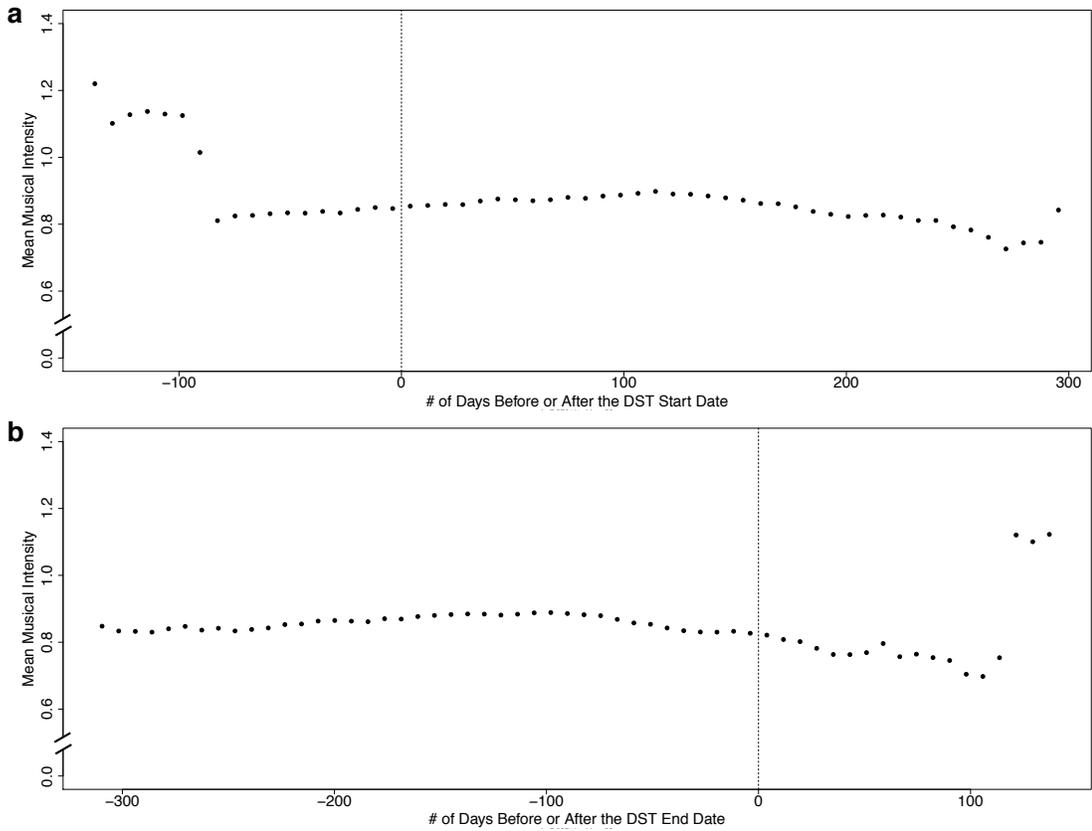


Figure A.5: Regression discontinuity analysis reveals no impact of daylight saving time (DST) transitions on musical intensity.

APPENDIX B

SUPPLEMENTARY INFORMATION OF CHAPTER 3

(CULTURAL VALUES AND CROSS-CULTURAL VIDEO CONSUMPTION

ON YOUTUBE)

Countries excluded in the regression			
BHR	Bahrain	OMN	Oman
BIH	Bosnia and Herzegovina	QAT	Qatar
CZE	Czech Republic	ROU	Romania
DZA	Algeria	SVK	Slovakia
GRC	Greece	TUN	Tunisia
HUN	Hungary	UGA	Uganda
MKD	Macedonia, Republic of	UKR	Ukraine
MNE	Montenegro	YEM	Yemen
Countries included in the regression			
ARE	United Arab Emirates	KEN	Kenya
ARG	Argentina	KOR	Korea, Republic of
AUS	Australia	KWT	Kuwait
AUT	Austria	LBN	Lebanon
BEL	Belgium	LTU	Lithuania
BGR	Bulgaria	LVA	Latvia
BRA	Brazil	MAR	Morocco
CAN	Canada	MEX	Mexico
CHE	Switzerland	MYS	Malaysia
CHL	Chile	NGA	Nigeria
COL	Colombia	NLD	Netherlands
DEU	Germany	NOR	Norway
DNK	Denmark	NZL	New Zealand
EGY	Egypt	PER	Peru
ESP	Spain	PHL	Philippines
EST	Estonia	POL	Poland
FIN	Finland	PRT	Portugal
FRA	France	RUS	Russian Federation
GBR	United Kingdom	SAU	Saudi Arabia
GHA	Ghana	SEN	Senegal
HKG	Hong Kong	SGP	Singapore
HRV	Croatia	SRB	Serbia
IDN	Indonesia	SVN	Slovenia
IND	India	SWE	Sweden
IRL	Ireland	THA	Thailand
ISR	Israel	TUR	Turkey
ITA	Italy	TWN	Taiwan, Republic of China
JOR	Jordan	USA	United States of America
JPN	Japan	ZAF	South Africa

Table B.1: Country names included and excluded for analyses.

BIBLIOGRAPHY

- [1] Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. The nature of prejudice. 1954.
- [2] Dan Ariely and Jonathan Levav. Sequential choice in group settings: Taking the road less traveled and less enjoyed. *Journal of consumer Research*, 27(3):279–290, 2000.
- [3] Noah Askin and Michael Mauskopf. What makes popular culture popular? product features and optimal differentiation in music. *American Sociological Review*, 82(5):910–944, 2017.
- [4] Young Min Baek. Relationship between cultural distance and cross-cultural music video consumption on youtube. *Social Science Computer Review*, 33(6):730–748, 2015.
- [5] Christopher Andrew Bail. Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media. *Proceedings of the National Academy of Sciences*, 113(42):11823–11828, 2016.
- [6] Laura-Lee Balkwill and William Forde Thompson. A cross-cultural investigation of the perception of emotion in music: Psychophysical and cultural cues. *Music Perception: An Interdisciplinary Journal*, 17(1):43–64, 1999.
- [7] Pablo Barberá. Less is more? how demographic sample weights can improve public opinion estimates based on twitter data. *Work Pap NYU*, 2016.
- [8] George A Barnett, Thomas Jacobson, Young Choi, and Sulien Sun-Miller. An examination of the international telecommunication network. *Journal of International Communication*, 3(2):19–43, 1996.
- [9] George A Barnett and Eunjung Sung. Culture and the structure of the international hyperlink network. *Journal of Computer-Mediated Communication*, 11(1):217–238, 2005.
- [10] Rachel F Baskerville. Hofstede never studied culture. *Accounting, organizations and society*, 28(1):1–14, 2003.

- [11] John W Berry, John Widdup Berry, Ype H Poortinga, Marshall H Segall, and Pierre R Dasen. *Cross-cultural psychology: Research and applications*. Cambridge University Press, 2002.
- [12] Parantapa Bhattacharya, Muhammad Bilal Zafar, Niloy Ganguly, Saptarshi Ghosh, and Krishna P Gummadi. Inferring user interests in the twitter social network. In *ACM Conference on Recommender systems (RecSys)*, pages 357–360, 2014.
- [13] Pierre Bourdieu. *Distinction: A social critique of the judgement of taste*. Harvard University Press, 1984.
- [14] Pierre Bourdieu. *Distinction: A social critique of the judgement of taste*. Routledge, 2013.
- [15] Thomas D Brock. Calculating solar radiation for ecological studies. *Ecological Modelling*, 14(1-2):1–19, 1981.
- [16] Jiajun Bu, Shulong Tan, Chun Chen, Can Wang, Hao Wu, Lijun Zhang, and Xiaofei He. Music recommendation by unified hypergraph: Combining social media information and music content. In *ACM International Conference on Multimedia (ACMMM)*, pages 391–400, 2010.
- [17] Javier M Buldú, P Cano, M Koppenberger, Juan A Almendral, and S Boccaletti. The complex network of musical tastes. *New Journal of Physics*, 9(6):172, 2007.
- [18] Robert Burnett. *The global jukebox: The international music industry*. Routledge, 2002.
- [19] Gabriella Castelli, Marco Mamei, Alberto Rosi, and Franco Zambonelli. Extracting high-level information from location data: The w4 diary example. *Mobile Networks and Applications*, 14(1):107–119, 2009.
- [20] Central Intelligence Agency. The world factbook. <https://www.cia.gov/library/publications/the-world-factbook/fields/2010.html>, 2017. [Online; accessed 15-October-2018].
- [21] Li Chen, Wen Wu, and Liang He. How personality influences users’ needs for recommendation diversity? In *CHI’13 Extended Abstracts on Human Factors in Computing Systems*, pages 829–834, 2013.

- [22] Junho Choi, Jaemin Jung, and Sang-Woo Lee. What causes users to switch from a local to a global social network site? the cultural, social, economic, and motivational factors of facebook's globalization. *Computers in Human Behavior*, 29(6):2665–2673, 2013.
- [23] Peter G Christenson and Jon Brian Peterson. Genre and gender in the structure of music preferences. *Communication Research*, 15(3):282–301, 1988.
- [24] Cisco. Global ip traffic forecast. <http://tinyurl.com/lhnyfto>, 2015. [Online; accessed 19-September-2015].
- [25] Connect, Ipsos. Music consumer insight report.(2016). <http://www.ifpi.org/downloads/Music-Consumer-Insight-Report-2016.pdf>, 2016. [Online; accessed 15-October-2018].
- [26] Philippe Coulangeon and Yannick Lemel. Is 'distinction' really outdated? questioning the meaning of the omnivorization of musical taste in contemporary france. *Poetics*, 35(2):93–111, 2007.
- [27] Marieke De Mooij and Geert Hofstede. Cross-cultural consumer behavior: A review of research findings. *Journal of International Consumer Marketing*, 23(3-4):181–192, 2011.
- [28] Tia DeNora. *Music in everyday life*. Cambridge University Press, 2000.
- [29] Sean Dwyer, Hani Mesak, and Maxwell Hsu. An exploratory examination of the influence of national culture on cross-national product diffusion. *Journal of International Marketing*, 13(2):1–27, 2005.
- [30] Jordan Etkin and Cassie Mogilner. Does variety among activities increase happiness? *Journal of Consumer Research*, 43(2):210–229, 2016.
- [31] William P Eveland Jr, Hyunjin Song, and Paul A Beck. Cultural variations in the relationships among network political agreement, political discussion frequency, and voting turnout. *International Journal of Public Opinion Research*, 27(4):461–480, 2015.
- [32] Face++. Attributes. <https://www.faceplusplus.com/>, 2019. [Online; accessed 15-May-2019].
- [33] Katayoun Farrahi, Markus Schedl, Andreu Vall, David Hauger, and

- Marko Tkalcić. Impact of listening behavior on music recommendation. In *International Society for Music Information Retrieval Conference (ISMIR)*, 2014.
- [34] Norman T Feather. Bridging the gap between values and actions: Recent applications of the expectancy-value model. *Handbook of motivation and cognition: Foundations of social behavior*, 2:151–192, 1990.
- [35] William C Forsythe, Edward J Rykiel, Randal S Stahl, Hsin-i Wu, and Robert M Schoolfield. A model comparison for daylength as a function of latitude and day of year. *Ecological Modelling*, 80(1):87–95, 1995.
- [36] Rémi Foucard, Slim Essid, Gaël Richard, and Mathieu Lagrange. Exploring new features for music classification. In *IEEE International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)*, pages 1–4, 2013.
- [37] Kai R Fricke, David M Greenberg, Peter J Rentfrow, and Philipp Yorck Herzberg. Computer-based music feature analysis mirrors human perception and can be used to measure individual music preference. *Journal of Research in Personality*, 75:94–102, 2018.
- [38] Ruth García-Gavilanes, Yelena Mejova, and Daniele Quercia. Twitter ain’t without frontiers: economic, social, and cultural boundaries in international communication. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 1511–1522. ACM, 2014.
- [39] Ruth Garcia-Gavilanes, Daniele Quercia, and Alejandro Jaimes. Cultural dimensions in twitter: Time, individualism and power. In *AAAI Conference on Weblogs and Social Media (ICWSM)*, 2013.
- [40] Clifford Geertz. *The interpretation of cultures*, volume 5019. Basic books, 1973.
- [41] Andrew Gelman and Guido Imbens. Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, pages 0–0, 2017.
- [42] Saptarshi Ghosh, Naveen Sharma, Fabricio Benevenuto, Niloy Ganguly, and Krishna Gummadi. Cognos: Crowdsourcing search for topic experts in microblogs. In *International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 575–590, 2012.

- [43] Tarleton Gillespie. 9 the relevance of algorithms. *Media technologies: Essays on communication, materiality, and society*, page 167, 2014.
- [44] Jennifer Golbeck, Cristina Robles, Michon Edmondson, and Karen Turner. Predicting personality from twitter. In *IEEE International Conference on Social Computing (SocialCom)*, pages 149–156, 2011.
- [45] Amir Goldberg. Mapping shared understandings using relational class analysis: The case of the cultural omnivore reexamined. *American Journal of Sociology*, 116(5):1397–1436, 2011.
- [46] Scott A Golder and Michael W Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, 2011.
- [47] Scott A Golder and Michael W Macy. Digital footprints: Opportunities and challenges for online social research. *Annual Review of Sociology*, 40:129–152, 2014.
- [48] Wen Gong. National culture and global diffusion of business-to-consumer e-commerce. *Cross cultural management: an international journal*, 16(1):83–101, 2009.
- [49] Google Trends. Year in search 2018. <https://trends.google.com/trends/yis/2018/GLOBAL/>, 2018. [Online; accessed 15-May-2019].
- [50] David M Greenberg, Michal Kosinski, David J Stillwell, Brian L Monteiro, Daniel J Levitin, and Peter J Rentfrow. The song is you: Preferences for musical attribute dimensions reflect personality. *Social Psychological and Personality Science*, 7(6):597–605, 2016.
- [51] David M Greenberg and Peter J Rentfrow. Music and big data: a new frontier. *Current opinion in behavioral sciences*, 18:50–56, 2017.
- [52] James J Gross. The emerging field of emotion regulation: An integrative review. *Review of general psychology*, 2(3):271–299, 1998.
- [53] James J Gross and Oliver P John. Individual differences in two emotion regulation processes: implications for affect, relationships, and well-being. *Journal of personality and social psychology*, 85(2):348, 2003.

- [54] William B Gudykunst, Yuko Matsumoto, STELLA Ting-Toomey, Tsukasa Nishida, Kwangsu Kim, and Sam Heyman. The influence of cultural individualism-collectivism, self construals, and individual values on communication styles across cultures. *Human communication research*, 22(4):510–543, 1996.
- [55] William B Gudykunst and Tsukasa Nishida. The influence of culture and strength of cultural identity on individual values in japan and the united states. *Intercultural Communication Studies*, 9(1):1–18, 2000.
- [56] Joseph F Hair, Ronald L Tatham, Rolph E Anderson, and William Black. *Multivariate data analysis*, volume 6. Pearson Prentice Hall Upper Saddle River, NJ, 2006.
- [57] David J Hargreaves and Adrian C North. The functions of music in everyday life: Redefining the social in music psychology. *Psychology of music*, 27(1):71–83, 1999.
- [58] Morton G Harmatz, Arnold D Well, Christopher E Overtree, Kathleen Y Kawamura, Milagros Rosal, and Ira S Ockene. Seasonal variation of depression and other moods: a longitudinal approach. *Journal of Biological Rhythms*, 15(4):344–350, 2000.
- [59] Brent Hecht, Lichan Hong, Bongwon Suh, and Ed H Chi. Tweets from justin bieber’s heart: The dynamics of the location field in user profiles. In *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 237–246, 2011.
- [60] Brent Hecht and Monica Stephens. A tale of cities: Urban biases in volunteered geographic information. In *AAAI conference on Weblogs and Social Media (ICWSM)*, 2014.
- [61] Mikael Henaff, Kevin Jarrett, Koray Kavukcuoglu, and Yann LeCun. Unsupervised learning of sparse features for scalable audio classification. In *International Society of Music Information Retrieval (ISMIR)*, pages 681–686, 2011.
- [62] Geert Hofstede, Geert Jan Hofstede, and Michael Minkov. *Cultures and organizations: software of the mind, Third Edition*. McGraw-Hill, 2010.
- [63] Neil Hurley and Mi Zhang. Novelty and diversity in top-n recommendation-analysis and evaluation. *ACM Transactions on Internet Technology (TOIT)*, 10(4):14, 2011.

- [64] Guido Imbens and Karthik Kalyanaraman. Optimal bandwidth choice for the regression discontinuity estimator. *The Review of Economic Studies*, 79(3):933–959, 2012.
- [65] Barbara E Kahn. Consumer variety-seeking among goods and services: An integrative review. *Journal of retailing and consumer services*, 2(3):139–148, 1995.
- [66] Matthew C Keller, Barbara L Fredrickson, Oscar Ybarra, Stéphane Côté, Kareem Johnson, Joe Mikels, Anne Conway, and Tor Wager. A warm heart and a clear head: The contingent effects of weather on mood and cognition. *Psychological Science*, 16(9):724–731, 2005.
- [67] Silvia Knobloch and Dolf Zillmann. Mood management via the digital jukebox. *Journal of communication*, 52(2):351–366, 2002.
- [68] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.
- [69] Amanda E Krause and Adrian C North. tis the season: Music-playlist preferences for the seasons. *Psychology of Aesthetics, Creativity, and the Arts*, 12(1):89, 2018.
- [70] Satish Krishnan and John Lymm. Determinants of virtual social networks diffusion: Insights from cross-country data. *Computers in Human Behavior*, 54:691–700, 2016.
- [71] Joseph B Kruskal and Myron Wish. *Multidimensional scaling*, volume 11. Sage, 1978.
- [72] David Lazer, Alex Pentland, Lada Adamic, Sinan Aral, Albert-László Barabási, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, et al. Computational social science. *Science*, 323(5915):721–723, 2009.
- [73] Sang-Gun Lee, Silvana Trimi, and Changsoo Kim. The impact of cultural differences on technology adoption. *Journal of world business*, 48(1):20–29, 2013.
- [74] Chan-Hoong Leong and Colleen Ward. Cultural values and attitudes toward immigrants and multiculturalism: The case of the eurobarometer

- survey on racism and xenophobia. *International Journal of Intercultural Relations*, 30(6):799–810, 2006.
- [75] Jonathan Levav and Rui Zhu. Seeking freedom through variety. *Journal of Consumer Research*, 36(4):600–610, 2009.
- [76] Kevin Lewis, Marco Gonzalez, and Jason Kaufman. Social selection and peer influence in an online social network. *Proceedings of the National Academy of Sciences*, 109(1):68–72, 2012.
- [77] Loet Leydesdorff and Ismael Rafols. Indicators of the interdisciplinarity of journals: Diversity, centrality, and citations. *Journal of Informetrics*, 5(1):87–100, 2011.
- [78] Omar Lizardo. Understanding the flow of symbolic goods in the global cultural economy. *International Journal of Contemporary Sociology*, 45(1):13, 2008.
- [79] Walter J Lonner, John W Berry, and Geert H Hofstede. Culture’s consequences: International differences in work-related values. *University of Illinois at Urbana-Champaign’s Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*, 1980.
- [80] Carleen F Maitland and Johannes M Bauer. National level culture and global diffusion: The case of the internet. *Culture, technology, communication: Towards an intercultural global village*, pages 87–128, 2001.
- [81] Justin McCrary. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714, 2008.
- [82] Sean M McNee, John Riedl, and Joseph A Konstan. Being accurate is not enough: How accuracy metrics have hurt recommender systems. In *CHI’06 Extended Abstracts on Human Factors in Computing Systems*, pages 1097–1101, 2006.
- [83] Alan Mislove, Sune Lehmann, Yong-Yeol Ahn, Jukka-Pekka Onnela, and J Niels Rosenquist. Pulse of the nation: U.S. mood throughout the day inferred from twitter. *Northeastern University*, 2010.
- [84] Marieke De Mooij. Convergence and divergence in consumer behaviour:

- implications for global advertising. *International Journal of advertising*, 22(2):183–202, 2003.
- [85] Fabian Mörchen, Alfred Ultsch, Mario Nöcker, and Christian Stamm. Databionic visualization of music collections according to perceptual distance. In *International Society of Music Information Retrieval (ISMIR)*, pages 396–403, 2005.
- [86] Mark EJ Newman. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical review E*, 64(1):016132, 2001.
- [87] Pippa Norris and Ronald Inglehart. *Cosmopolitan communications: Cultural diversity in a globalized world*. Cambridge University Press, 2009.
- [88] Adrian C North, David J Hargreaves, and Jon J Hargreaves. Uses of music in everyday life. *Music Perception: An Interdisciplinary Journal*, 22(1):41–77, 2004.
- [89] Joseph S Nye Jr. *Soft power: The means to success in world politics*. Public affairs, 2004.
- [90] Ofcom. 2017 communications market report: UK. <https://www.ofcom.org.uk/research-and-data/multi-sector-research/cmr/cmr-2017>, 2017. [Online; accessed 15-October-2018].
- [91] Tore Opsahl, Filip Agneessens, and John Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social networks*, 32(3):245–251, 2010.
- [92] Jaram Park, Young Min Baek, and Meeyoung Cha. Cross-cultural comparison of nonverbal cues in emoticons on twitter: Evidence from big data analysis. *Journal of Communication*, 64(2):333–354, 2012.
- [93] Minsu Park, Mor Naaman, and Jonah Berger. A data-driven study of view duration on youtube. In *AAAI Conference on Web and Social Media (ICWSM)*, 2016.
- [94] Minsu Park, Jaram Park, Young Min Baek, and Michael Macy. Cultural values and cross-cultural video consumption on youtube. *PLoS one*, 12(5):e0177865, 2017.
- [95] Minsu Park, Jennifer Thom, Sarah Mennicken, Henriette Cramer, and

- Michael Macy. Global music streaming data reveal diurnal and seasonal patterns of affective preference. *Nature Human Behaviour*, 3(3):230, 2019.
- [96] Minsu Park, Ingmar Weber, Mor Naaman, and Sarah Vieweg. Understanding musical diversity via online social media. In *AAAI Conference on Web and Social Media (ICWSM)*, 2015.
- [97] Richard A Peterson. Understanding audience segmentation: From elite and mass to omnivore and univore. *Poetics*, 21(4):243–258, 1992.
- [98] Richard A Peterson. The rise and fall of highbrow snobbery as a status marker. *Poetics*, 25(2):75–92, 1997.
- [99] Richard A Peterson. Problems in comparative research: The example of omnivorousness. *poetics*, 33(5):257–282, 2005.
- [100] Terry F Pettijohn, Greg M Williams, and Tiffany C Carter. Music for the seasons: seasonal music preferences in college students. *Current Psychology*, 29(4):328–345, 2010.
- [101] Pew Research Center. Demographics and political views of news audiences. <http://tinyurl.com/kbb6n9m>, 2012. [Online; accessed 30-December-2014].
- [102] Pew Research Center. 2016 U.S. smartphone use. <http://www.pewresearch.org/fact-tank/2016/01/29/us-smartphone-use>, 2015. [Online; accessed 15-October-2018].
- [103] Edward L Platt, Rahul Bhargava, and Ethan Zuckerman. The international affiliation network of youtube trends. In *AAAI Conference on Web and Social Media (ICWSM)*, 2015.
- [104] Alan L Porter and Ismael Rafols. Is science becoming more interdisciplinary? measuring and mapping six research fields over time. *Scientometrics*, 81(3):719–745, 2009.
- [105] Lin Qiu, Han Lin, Jonathan Ramsay, and Fang Yang. You are what you tweet: Personality expression and perception on twitter. *Journal of Research in Personality*, 46(6):710–718, 2012.
- [106] Daniele Quercia, Michal Kosinski, David Stillwell, and Jon Crowcroft. Our twitter profiles, our selves: Predicting personality with twitter. In

IEEE International Conference on Social Computing (SocialCom), pages 180–185, 2011.

- [107] David Rawlings and Vera Ciancarelli. Music preference and the five-factor model of the neo personality inventory. *Psychology of Music*, 25(2):120–132, 1997.
- [108] Peter J Rentfrow and Samuel D Gosling. The do re mi’s of everyday life: The structure and personality correlates of music preferences. *Journal of personality and social psychology*, 84(6):1236, 2003.
- [109] Mark Rimmer. Beyond omnivores and univores: The promise of a concept of musical habitus. *Cultural Sociology*, 6(3):299–318, 2012.
- [110] Milton Rokeach. *The nature of human values*. Free press, 1973.
- [111] Shahar Ronen, Bruno Gonçalves, Kevin Z Hu, Alessandro Vespignani, Steven Pinker, and César A Hidalgo. Links that speak: The global language network and its association with global fame. *Proceedings of the National Academy of Sciences*, 111(52):E5616–E5622, 2014.
- [112] Paul A Samuelson. Consumption theory in terms of revealed preference. *Economica*, 15(60):243–253, 1948.
- [113] Johann Schrammel, Christina Köffel, and Manfred Tscheligi. Personality traits, usage patterns and information disclosure in online communities. In *British HCI Group Annual Conference on People and Computers: Celebrating People and Technology*, pages 169–174, 2009.
- [114] H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS One*, 8(9):e73791, 2013.
- [115] M Ángeles Serrano, Marián Boguná, and Alessandro Vespignani. Extracting the multiscale backbone of complex weighted networks. *Proceedings of the national academy of sciences*, 106(16):6483–6488, 2009.
- [116] Itamar Simonson. The effect of purchase quantity and timing on variety-seeking behavior. *Journal of Marketing Research*, 27(2):150–162, 1990.

- [117] John A Sloboda and Patrik N Juslin. *Music and emotion: Theory and research*. Oxford University Press, 2001.
- [118] Spotify. Spotify technology s.a. direct listing prospectus., 2018. [Online; accessed 15-October-2018].
- [119] Bogdan State, Patrick Park, Ingmar Weber, Michael Macy, et al. The mesh of civilizations in the global network of digital communication. *PloS one*, 10(5):e0122543, 2015.
- [120] Andrew Stirling. On the economics and analysis of diversity. *Science Policy Research Unit (SPRU), Electronic Working Papers Series*, 28:1–156, 1998.
- [121] Andy Stirling. A general framework for analysing diversity in science, technology and society. *Journal of the Royal Society Interface*, 4(15):707–719, 2007.
- [122] Joseph D Straubhaar. Beyond media imperialism: Assymetrical interdependence and cultural proximity. *Critical Studies in media communication*, 8(1):39–59, 1991.
- [123] Harsh Taneja and James G Webster. How do global audiences take shape? the role of institutions and culture in patterns of web use. *Journal of Communication*, 66(1):161–182, 2015.
- [124] Robert E Thayer. *The origin of everyday moods: Managing energy, tension, and stress*. Oxford University Press, USA, 1997.
- [125] Robert E Thayer, J Robert Newman, and Tracey M McClain. Self-regulation of mood: Strategies for changing a bad mood, raising energy, and reducing tension. *Journal of personality and social psychology*, 67(5):910, 1994.
- [126] Mike Thelwall, Kevan Buckley, and Georgios Paltoglou. Sentiment in twitter events. *Journal of the Association for Information Science and Technology*, 62(2):406–418, 2011.
- [127] Donald L Thistlethwaite and Donald T Campbell. Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational psychology*, 51(6):309, 1960.

- [128] Jeanne L Tsai. Ideal affect: Cultural causes and behavioral consequences. *Perspectives on Psychological Science*, 2(3):242–259, 2007.
- [129] Douglas R Turnbull, Justin A Zupnick, Kristofer B Stensland, Andrew R Horwitz, Alexander J Wolf, Alexander E Spigel, Stephen P Meyerhofer, and Thorsten Joachims. Using personalized radio to enhance local music discovery. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*, pages 2023–2028, 2014.
- [130] Teun A Van Dijk. Power and the news media. *Political communication in action*, pages 9–36, 1996.
- [131] Koen Van Eijck. Social differentiation in musical taste patterns. *Social forces*, 79(3):1163–1185, 2001.
- [132] Alexander Vander Stichele and Rudi Laermans. Cultural participation in flanders: Testing the cultural omnivore thesis with population data. *Poetics*, 34(1):45–64, 2006.
- [133] Nils Lennart Wallin, Björn Merker, and Steven Brown. *The origins of music*. MIT press, 2001.
- [134] Alan Warde, David Wright, and Modesto Gayo-Cal. Understanding cultural omnivorousness: Or, the myth of the cultural omnivore. *Cultural sociology*, 1(2):143–164, 2007.
- [135] Wikipedia. Last.fm. <http://en.wikipedia.org/wiki/Last.fm>, 2015. [Online; accessed 19-January-2015].
- [136] World Bank Open Data. Gdp per capita (current us\$). <http://data.worldbank.org/>, 2014. [Online; accessed 19-September-2016].
- [137] World Bank Open Data. Internet users (per 100 people). <http://data.worldbank.org/>, 2014. [Online; accessed 19-September-2016].
- [138] Zhiyong Yang, Jingguo Wang, and Mehdi Murali. Effect of peer influence on unauthorized music downloading and sharing: The moderating role of self-construal. *Journal of Business Research*, 2014.

- [139] YouTube. Statistics. <http://tinyurl.com/lkbdgfx>, 2015. [Online; accessed 19-September-2015].
- [140] Elena Zheleva, John Guiver, Eduarda Mendes Rodrigues, and Nataša Milić-Frayling. Statistical models of music-listening sessions in social media. In *International World Wide Web Conference (WWW)*, pages 1019–1028, 2010.
- [141] Lianxi Zhou and Michael K Hui. Symbolic value of foreign products in the people’s republic of china. *Journal of international marketing*, 11(2):36–58, 2003.
- [142] Dolf Zillmann. Mood management: Using entertainment to full advantage. In *Communication, Social Cognition, and Affect (PLE: Emotion)*, pages 163–188. Psychology Press, 2015.