American Sociological Review
1–35

DOI: 10.1177/0003122417728662
journals.sagepub.com/home/asr

What Makes Popular Culture Popular? Product Features and Optimal Differentiation in Music

Noah Askin[a] and Michael Mauskapf[b]

Abstract

In this article, we propose a new explanation for why certain cultural products outperform their peers to achieve widespread success. We argue that products' position in feature space significantly predicts their popular success. Using tools from computer science, we construct a novel dataset allowing us to examine whether the musical features of nearly 27,000 songs from *Billboard*'s Hot 100 charts predict their levels of success in this cultural market. We find that, in addition to artist familiarity, genre affiliation, and institutional support, a song's perceived proximity to its peers influences its position on the charts. Contrary to the claim that all popular music sounds the same, we find that songs sounding too much like previous and contemporaneous productions—those that are highly typical—are less likely to succeed. Songs exhibiting some degree of optimal differentiation are more likely to rise to the top of the charts. These findings offer a new perspective on success in cultural markets by specifying how content organizes product competition and audience consumption behavior.

Keywords

consumption, music, optimal differentiation, popular culture, product features, typicality

What makes popular culture popular? Scholars across the humanities and social sciences have spilled considerable ink trying to answer this question. However, our understanding of why certain cultural products succeed over others remains incomplete. Popular culture tends to reflect, or is intentionally aimed toward, the tastes of the public, yet there exists wide variation in the relative popularity of these products (Rosen 1981; Storey 2006). Extant research in sociology and related disciplines suggests that audiences seek and use a wide range of information as signals of the quality and value of new products (Keuschnigg 2015). This includes the characteristics of and relations between cultural producers (Peterson 1997; Uzzi and Spiro 2005; Yogev 2009), peer preferences and related social influence dynamics (Lizardo 2006; Mark 1998; Salganik, Dodds, and Watts 2006), and various elements in the institutional environment (Hirsch 1972; Peterson 1990).

Each of these signals plays an important role in determining which products audiences

[a]INSEAD
[b]Columbia Business School

Corresponding Author:
Noah Askin, INSEAD, Boulevard de Constance, 77305 Fontainebleau, France
E-mail: noah.askin@insead.edu

select, evaluate, and recommend to others. These choices and the preferences they express vary widely over time and across individuals. Nevertheless, research suggests that the inherent quality of cultural products also affects how audiences classify and evaluate them (Goldberg, Hannan, and Kovács 2016; Jones et al. 2012; Lena 2006; Rubio 2012; Salganik et al. 2006). Certain product features may independently signal quality and attract audience attention (Hamlen 1991), however, it is more likely that these features matter most as an ensemble. They work both by creating a multidimensional representation of products and by positioning those products across the plane of possible feature combinations.[1] Rather than existing in a vacuum, cultural products are perceived in relation to one another in feature space, and these relationships shape how consumers organize and discern the art worlds around them (Becker 1982).

One way to think about how product position shapes performance outcomes is through the lens of categories research. This work highlights how social classification systems organize audiences' expectations and preferences (Hsu 2006; Zuckerman 1999), helping them draw connections between products. We agree that producer categories play a significant role in structuring taste and consumption behavior (Bourdieu 1993). However, much of the work in this area makes the implicit assumption that category *labels* remain tightly coupled with a set of underlying product *features*. Recent research shows that product features need not necessarily cluster or align with prevailing classification schemes (Anderson 1991; Kovacs and Hannan 2015; Pontikes and Hannan 2014).[2]

Category labels (e.g., "country" in the case of musical genres) work well when navigating stable product markets with clearly defined category boundaries. These labels, however, do not always reflect how audiences actually make sense of the world in which they are embedded. This is especially the case in contexts where products are complex and tastes are idiosyncratic and dynamic (Lena 2015). In these domains, extant category labels may not provide adequate or accurate information to consumers, who must instead rely on products' underlying features to draw comparisons and make decisions.

We build on these insights to propose a new explanation for why certain cultural products outperform their competitors to achieve success. In the context of popular music, we argue that audiences use musical features as signals to draw latent associations between songs. These associations exist in partial independence from traditional categories. As such, feature-associations help organize the choice set from which audiences select and evaluate products, positioning certain songs more advantageously.

We hypothesize that hit songs are able to manage a similarity–differentiation tradeoff. Successful songs invoke conventional feature combinations associated with previous hits while at the same time displaying some degree of novelty distinguishing them from their peers. This prediction speaks to the competitive benefits of optimal differentiation, a finding that reoccurs across multiple studies and areas in sociology and beyond (Goldberg et al. 2016; Lounsbury and Glynn 2001; Uzzi et al. 2013; Zuckerman 2016).

In this article, we test this prediction with the aim of better understanding the relationship between product features and success in music. To that end, we constructed a novel dataset consisting of nearly 27,000 songs that appear on the *Billboard* Hot 100 charts between 1958 and 2016. The data include algorithmically derived features that describe a song's sonic qualities. Sonic features range from relatively objective musical characteristics, such as "key," "mode," and "tempo," to perceptual features that quantify a song's "acousticness," "energy," and "danceability," among others.

First, we establish the baseline validity of individual features in predicting a song's peak position and longevity on the charts. We then use these features to construct a measure of sonic similarity or typicality and examine its effect on chart performance. Popular opinion suggests that songs are most likely to succeed when they adhere to a conventional and reproducible template (Dhanaraj and Logan 2005; Thompson 2014). However, we find that the most successful songs in our dataset are optimally differentiated from their peers.

Our results provide strong evidence that, net of other factors such as artist familiarity and genre affiliation, product features matter, particularly in the way they structure songs' relationships to each other. Using new, micro-level feature data to specify how cultural content organizes how audiences distinguish products compels us to rethink some of the basic mechanisms behind consumption and taste formation. These findings, and the data and methods we use, make important contributions to cultural and economic sociology by offering a new perspective on success in cultural markets.

CULTURAL PREFERENCES AND THE SIMILARITY–DIFFERENTIATION TRADEOFF

Predicting how well a new product will fare in the marketplace for audience attention presents a difficult challenge. This is primarily due to the countless variables and contingencies that may influence performance outcomes. This challenge is particularly pronounced in the realm of the cultural or "creative" industries (Caves 2000; Hadida 2015). The reason for this is that these industries tend to generate products and experiences whose evaluation involves a subjective component (Krueger 2005). Even after a cultural product—a painting, film, or song—has been anointed a "success," it can be difficult to explain *ex post* why certain products enjoy more success than others (Bielby and Bielby 1994; Lieberson 2000).[3]

The relative popularity of a cultural product is usually ascribed to prevailing tastes, which are largely considered to be a function of individuals' idiosyncratic preferences, past experiences, and exposure patterns, along with the prevailing opinions of others. Moreover, different types of performance outcomes (e.g., mass appeal versus critical acclaim) beget different modes of explanation, and they require audiences to consider distinct dimensions of evaluation that are often context specific. Thus, our ability to explain what constitutes a hit versus a flop remains limited.

Producer Characteristics and Professional Networks

Scholars interested in this question traditionally take one of several approaches to explain the determinants of cultural preferences and product performance. The first set of explanations focuses on the characteristics of cultural producers. These include artist reputation (Bourdieu 1993), past performance outcomes (Peterson 1997), and the structure of artistic professional networks (Godart, Shipilov, and Claes 2014; Yogev 2009). Indeed, just as cultural products are perceived by audiences in relation to one another, they are also created by producers who form collaborative relationships and draw inspiration from each other's work.

In the context of Broadway musicals, Uzzi and Spiro (2005) find that when the network of collaborations between artists and producers displays small-world properties, cultural productions are more likely to achieve critical and commercial success. Phillips (2011, 2013) finds that the artists who are most likely to re-record and release jazz standards come, surprisingly, from structurally disconnected cities. Research on sampling in rap music (Lena and Pachucki 2013), innovations in video game production (de Vaan, Vedres, and Stark 2015), and the creative success of inventors (Fleming, Mingo, and Chen 2007) provides ample evidence that certain types of producer networks are more likely to generate new and successful products via the recombination of diverse ideas. Thus, the interconnectedness of producers and of the production process more generally plays an important role in shaping product performance and consumer taste.[4]

Social Influence

The second set of variables used to explain the success of cultural products pertains to audience or demand-side characteristics. Variables of this sort include individual and collective trends in demand, as well as other related consumer dynamics, such as homophily (Mark 1998) and endogenous diffusion patterns (Rossman 2012). These explanations speak to the significant role of social influence, which is often responsible for wide variation in product adoption and taste formation (DellaPosta, Shi, and Macy 2015). In a series of online experiments, Salganik and colleagues investigated how product quality and social influence affect success in an artificial music market (Salganik et al. 2006; Salganik and Watts 2008, 2009). Despite the outsized role of social influence, they found compelling evidence that the likelihood of a song being downloaded by participants is determined in part by its inherent quality—but the exact nature of such "quality" remains obscure.

Category Labels

The categories literature provides a third class of explanations for the variable success of cultural products (Hsu 2006; Jones et al. 2012). Product categories and the labels attached to them reflect largely agreed-upon conventions that audiences attribute to certain groups of products. In this sense, "products are cultural objects imbued with meaning based on shared understandings, and are themselves symbols or representations of those meanings" (Fligstein and Dauter 2007:115).

Much of the research on social classification explores the role of categories in organizing product markets and consumer choice. This process is particularly salient in cultural markets (Caves 2000; DiMaggio 1987). In these settings, classification systems provide the context through which producers and consumers structure their tastes, preferences, and identities (Bourdieu 1993; Peterson 1997). Classification systems also determine how people search and evaluate the art worlds around them (Becker 1982). Indeed, the emergence and institutionalization of genre categories features prominently in explanations of market competition across a number of cultural domains. These include film (Hsu 2006), painting (Wijnberg and Gemser 2000), literature (Frow 1995, 2006), and music (Frith 1996; Holt 2007; Lena and Peterson 2008; Negus 1992).

Category researchers have made considerable contributions to our understanding of when and why certain kinds of organizations or products succeed (Hsu, Negro, and Perretti 2012; Zuckerman 1999). However, this work suffers from two important limitations. First, while it is true that categories play an important role in shaping how audiences search, select, and evaluate products, they often provide a relatively coarse and static picture of "the market," assuming a nested hierarchical structure that is more or less agreed-upon by market actors. But we know that categories and their boundaries are dynamic and contested; pointing to different meanings for members of different communities (Lena 2012; Sonnett 2004).

Second, most research in this area highlights the *symbolic* labels attached to categories, often ignoring the *features* of the products that occupy them. Labels constitute socially constructed and symbolically ascribed descriptors for a given category. Features, on the other hand, provide more fine-grained information about a focal product's underlying composition and position in "conceptual space" (Kovács and Hannan 2015). Recent research indicates that individuals classify products and other entities across a number of different dimensions, including shared cultural frames or world views (Goldberg 2011), overlapping cognitive interpretations (de Vaan et al. 2015), and interpersonal connections between producers or consumers (Lena 2015). The classification structures that emerge from these processes may or may not align with explicit categorical prescriptions such as musicological genre, suggesting an alternative criterion by which audiences position and compare similar producers and their products in the marketplace.

Product Features and Audience Associations

Category labels are usually coupled with a set of underlying features or attributes, but the degree of coupling between features and labels is highly variable (Anderson 1991; Pontikes and Hannan 2014). For example, Bob Dylan's version of "Like a Rolling Stone" might be tagged with labels like "folk," "Americana," or even "rock-n-roll," but it also exhibits a large number of features, including its duration (6:09), key (C Major), instrumentation (vocals, guitar, bass, electric organ, harmonica, tambourine), and thematic message (love, resentment). These features—the high-dimensional space of attributes that constitute the DNA of individual products—are culturally determined, grounding cultural products in material reality and granting them structural autonomy (Alexander and Smith 2002).

Recent research suggests that the features of cultural products also shape classification processes and performance outcomes (Jones et al. 2012; Lena 2006; Rossman and Schilke 2014). Like category labels, features can be used to position products that seem more or less similar to each other (see Cerulo 1988), shaping consumers' perceptions and sense-making in distinct ways (Tversky 1977). Furthermore, empirical evidence from popular music studies suggests that certain features (e.g., instrumentation) shape listening preferences and play an important role in determining why some products succeed and others fail (Nunes and Ordanini 2014).

Our reading of these literatures suggests there is a gap in the way product features have been conceptualized in previous work and their role in positioning products for success. Rather than influencing consumption in isolation from one another, features cohere in particular combinations to generate holistic, *gestalt* representations of products. Recent work at the vanguard of network neuroscience is beginning to explore how individuals make sense of these representations (Brashears and Quintane 2015; Zerubavel et al. 2015). Yet we still know little about how this process unfolds.[5]

In the context of cultural consumption, we argue that consumers position products across a multidimensional feature space. In this space, certain objects are perceived to be more (or less) similar depending on the features they do (or do not) share with other products. These latent associations represent the world of products in which consumers are embedded, and they exhibit a social life all their own (Carroll, Khessina, and McKendrick 2010; Douglas and Isherwood 1996).[6] These associations also organize the relevant comparison set from which consumers select and evaluate cultural products.

This argument goes beyond previous treatments of the determinants of success in cultural markets in two important ways. First, we highlight the significance of the implicit relationships formed within product space. In this way, we refrain from making success purely about dynamics external to the cultural product itself, such as producer networks and interpersonal consumer relationships. We argue instead that audience evaluations of products are shaped not only by producer and consumer characteristics, or social influence pressures, but also by a product's position within a broader ecosystem of cultural production and consumption. Thus, the choices consumers make are shaped by their individual preferences, relationships, and various other factors, but they are also influenced by the feature-based similarity space within which products are embedded (Kovács and Hannan 2015). Put another way, consumers' direct and indirect exposure to a given set of related products plays a critical role in shaping their future selection decisions and preferences.

Second, we argue that the structure and effect of these feature-based associations are conceptually and analytically distinct from those usually attributed to traditional categories. Research on category emergence suggests that labels and features operate in separate planes, which may or may not align with one other (Pontikes and Hannan 2014). We already know that consumers refer to established categories to make sense of the products they encounter (Zuckerman 1999). However, recent work at the intersection of culture, cognition, and strategy identifies the distinctive role of "product concepts"—loose relational structures that shape consumer cognition beyond purely categorical classification (Kahl 2015). These insights justify our focus on feature-based associations, suggesting that consumers in certain contexts are likely to use an amalgamation of features rather than (or in addition to) labels to position, select, and evaluate products. In the analysis that follows, we account for both of these dimensions to explain why certain songs attract audience attention and outperform their competition in the market for popular music.

The Similarity–Differentiation Tradeoff

We have already reviewed a number of plausible explanations for the variable success of cultural products, including producer reputation and category membership. However, the study of product features and the associations they generate highlights a new set of mechanisms to explain why certain products achieve popularity while others do not. One common way to examine the effects of product positioning on market performance is to measure crowding and differentiation dynamics (e.g., Bothner, Kang, and Stuart 2007). This strategy has been particularly useful in the organizational ecology literature (Barroso et al. 2014; Podolny, Stuart, and Hannan 1996). In this line of work, the presence of competitors can saturate a consumer or product space (e.g., niche), making it increasingly difficult for new entrants to survive. Research across a variety of empirical contexts shows that the ability to differentiate oneself and develop a distinctive identity can help products, organizations, and other entities compete within or across niches (Deephouse 1999; Hannan and Freeman 1977; Hsu and Hannan 2005; Swaminathan and Delacroix 1991; Zhao et al. 2017).

Alternatively, work in cognitive and social psychology argues that conformity is the recipe for success. For example, research on liking (Zajonc 1968) suggests that the more people are exposed to a stimulus, the more they enjoy it, regardless of whether they recognize having been previously exposed. In music, this means the more a song sounds like something listeners have heard before, the more likely they are to evaluate it positively and listen to it again (see Huron 2013). This argument lies at the heart of "hit song science," which claims that, with enough marketing support, artists can produce a hit song simply by imitating past successes (Dhanaraj and Logan 2005; Thompson 2014).

Rather than test these competing predictions individually, we hypothesize that the pressures toward conformity and differentiation act in concert. Products must differentiate themselves from the competition to avoid crowding, yet they cannot differentiate to such an extent as to make themselves unrecognizable (Kaufman 2004). Research on consumer behavior suggests that audiences engage in this tradeoff as well. When choosing a product, audiences conform on certain identity-signaling attributes (e.g., a product's brand or category), while distinguishing themselves on others (e.g., color or instrumentation; see Chan, Berger, and Van Boven 2012). This tension between conformity and differentiation is central to our understanding of social identities (Brewer 1991), category spanning (Hsu 2006; Zuckerman 1999), storytelling (Lounsbury and Glynn 2001), consumer products (Lancaster 1975), and taste (Lieberson 2000). Taken together, this work signals a common trope across the social sciences: the path to success requires some degree of both conventionality and novelty (Uzzi et al. 2013).

In the context of popular music, songs must strike a balance between being recognizable and being different. Those that best manage this similarity–differentiation tradeoff will attract more audience attention and experience more success. Stated more formally, we predict an inverted U-shaped relationship between a song's relative typicality and its performance on the *Billboard* Hot 100 charts. Our analysis highlights the opposing pressures of crowding and differentiation by constructing a summary measure of song typicality. This index accounts for how features position a song relative to its musical neighbors. Adjusting for a host of other factors, including an artist's previous success and genre affiliation, we expect that songs exhibiting *optimal differentiation* across the feature space are more likely to achieve widespread popularity, whereas songs that are deemed too similar to—or dissimilar from—their peers will struggle to reach the top of the charts (cf. Zuckerman 2016).

DATA AND METHODS

Studying the relative typicality of products can shed light on how audience preferences are shaped across a number of empirical contexts. We believe music represents an ideal setting in which to test these dynamics, due in part to its reliance on an internally consistent grammar. While songs can be quite different from one another, they follow the same set of basic principles or "rules" based on melody, harmony, and rhythm. Listeners' tastes, on the other hand, do not have such concrete bounds.

Salganik and colleagues (2006) showed that consumer choice in an artificial music market is driven by both social influence and a song's inherent quality. However, their measure of "quality" was audience preference in the absence of experimental manipulation. Measuring quality "objectively," on the other hand, requires a comprehensive technical understanding of music's form and features. Due to the specialized skills needed to identify, categorize, and evaluate such features reliably, work that meets these demands is limited. The research that has been conducted uses musicological techniques to construct systems of comparable musical codes that may be more or less present in a particular musical work (Cerulo 1988; La Rue 2001; Nunes and Ordanini 2014). Yet even if social scientists learned these techniques, or collaborated more often with musicologists, it would be very difficult to apply and automate such complex codes at scale. These difficulties have been partially attenuated by the application of digital data sources and new computational methods to the study of cultural objects. These techniques were first developed by computer scientists. More recently, these technologies have begun to filter into the toolkits of cultural sociologists (Bail 2014), who have traditionally been criticized for being "methods poor" (DiMaggio, Nag, and Blei 2013:2).

Most relevant for our purposes are advances in music information retrieval (MIR) and machine learning (e.g., Friberg et al. 2014; Serrà et al. 2012). These fields have developed new methods to reduce the high dimensionality of musical compositions to a set of discrete features, much like what topic modeling has done for the study of large text corpora (Blei, Ng, and Jordan 2003). These developments have generated new research possibilities that were previously considered impractical. We use a novel dataset that includes discrete representations of musical features in the form of sonic features (a song's "acoustic footprint"). These data allow us to investigate how popular success is contingent, in part, on a song's relative position within feature space.

Our primary data come from the weekly *Billboard* Hot 100 charts, which we reconstructed from their inception on August 4, 1958 through March 26, 2016. The Hot 100 charts are published by *Billboard Magazine*, but the data we use for our analysis come from an online repository known as "The Whitburn Project." Joel Whitburn collected and published anthologies of the charts (Whitburn 1986, 1991) and, beginning in 1998, a dedicated fan base started to collect, digitize, and add to the information contained in those guides. This augmented existing chart data and provides additional details about the songs and albums on the charts, including metadata and week-to-week rankings for more than 26,800 songs spanning almost 60 years.

A descriptive comparison of these songs with others that did not appear on the Hot 100 charts suggests that, although the observations included in our analysis constitute a slightly more homogenous or "typical" sample than is represented in music broadly, *the distribution of song typicality across these samples is nearly identical*. This makes the charts an appropriate proxy for studying consumer evaluation and product performance in the field of popular music (see Part A of the online supplement for a more formal comparison).[7] Furthermore, although the algorithm used to create the charts has evolved over the years, these charts remain the industry standard.[8] They have been used extensively in social science research on popular music (Alexander 1996; Anand and Peterson 2000; Bradlow and Fader 2001; Dowd 2004; Lena 2006; Lena and Pachucki 2013; Peterson and Berger 1975). They are also noted for their reliability as indicators of popular taste (e.g., Eastman and Pettijohn 2015).

Outcome variables. The weekly *Billboard* charts provide us with a real-world performance outcome that reflects the general popularity of a song and can be tracked and compared over time. Unlike movie box-office results or television show ratings, music's content-owners closely guard sales data. This leaves songs' diffusion across radio stations (Rossman 2012) or their chart position as the most reliable and readily available performance outcome. In their examination of fads in baby naming, Berger and Le Mens (2009) use peak popularity and longevity as key variables in the measurement of cultural diffusion; we adapt these as our outcome variables, *peak position* and *weeks on charts*. Although these two outcomes are related to one another (i.e., songs that reach a higher peak chart position tend to remain on the charts longer, $R \approx .72$), we use both outcomes in our analysis. We also reverse-code peak chart position (101 – chart position) so that positive coefficients on our independent variables indicate a positive relationship with a song's success on the charts.

To account for the competitive dynamics between songs appearing on the same chart, we include a set of models that use a third indicator of success based on week-to-week change in chart position. We subtracted each song's (reverse-coded) position during the previous week (t) from its current position ($t + 1$) to determine the effect of song typicality on weekly changes in chart position. We realize that a third dependent variable adds complexity to the analysis. However, we believe this approach is appropriate because it (1) better captures the dynamic nature of the charts, which can change considerably from week to week, while allowing us to include fixed effects for songs; (2) does not penalize the relatively short shelf-life of song popularity; and (3) accounts for the fact that songs appearing near the bottom of the charts have greater opportunity for improvement than do those at the top.

Genre data. The *Billboard* data require augmentation to capture the multifaceted social and compositional elements of songs and artists. Genre categories evolve and are potentially contentious (Lena and Pachucki 2013). In spite of this, they provide an important form of symbolic standing, organizing the listening patterns and evaluations of producers, consumers, and critics (Bourdieu 1993; Holt 2007; Lena 2012). Moreover, genres play a significant role in defining and shaping artists' identities (Peterson 1997; Phillips and Kim 2008). In turn, artists' identities help determine the type of listeners who end up seeking out and being exposed to new music. Audiences consequently reinforce artist identities and genre structures (Frith 1996; Negus 1992), setting expectations for both producers and their products (Peterson 1997).

To account for the effect of traditional category labels, we collected musicological genre data from Discogs.com, an encyclopedic music site and marketplace containing extensive information on music recordings, specifically singles and albums (see Montauti and Wezel 2016). Like other music websites, particularly those with user-generated and curated data, Discogs includes multiple genre and style (or subgenre) attributions for each release (i.e., single, album, EP, or LP). Up to three genre and six style attributions are possible. However, we created dummy variables for the *primary genre* affiliated with each release in our analysis (see "crossovers" below for an exception). Many songs on the Hot 100 were released as singles, allowing us to obtain fine-grained, song-level genre classification data. For songs that were not released as singles, we use the primary genre attributed to the album on which the song appears.[9] Based on these data, our sample covers 15 genre categories—including pop, rock, blues, electronic, jazz, and hip hop.[10]

The Echo Nest sonic feature data. Although genre represents an important means of symbolic classification in music, our interest in more fine-grained, feature-based associations required the collection of data summarizing the sonic attributes of each song. For these data we turned to The Echo Nest, an online music intelligence provider that offered access to much of their data via a suite of Application Programming Interfaces (APIs). This organization represents the current gold standard in MIR, having been purchased by music streaming leader Spotify in 2014 to power its analytics and recommendation engines. Using web crawling and audio encoding technology, The Echo Nest has collected and continuously updates information on more than 30 million songs and 3 million artists. Their data contain objective and derived qualities of audio recordings, as well as qualitative information about artists based on text analyses of artist mentions in digital articles and blog posts.

We accessed The Echo Nest API to collect complete data on 94 percent of the songs (25,102 of 26,846 total songs) that appeared on the charts between 1958 and 2016. This information includes several objective musical features (e.g., tempo, mode, and key), as well as some of the company's own creations (e.g., valence, danceability, and acousticness). Songs are assigned a quantitative value for each feature, which are measured using various scales. Table 1 briefly describes the 11 features used in our analysis. We recognize that there are limitations associated with distilling complex cultural products into a handful of discrete features. However, we believe these features represent the best available approximation of what people hear when they listen to music. Nearly 20 years of research and advancements in MIR techniques have produced both high- and low-level audio features that provide an increasingly robust representation of how listeners perceive music (Friberg et al. 2014). Personal communications with leading MIR researchers support our belief that these measures provide the most systematic attempt to capture songs' material and sensory composition at scale. Moreover, these features were created specifically for song-to-song comparisons to inform algorithmically-generated recommendations for listeners.

Predictor variable: song typicality. In an effort to provide a more nuanced explanation of how a song's relative position within feature space affects performance, we constructed a dynamic measure of song typicality. For this variable, *genre-weighted typicality (yearly)*, we measure the cosine similarity between songs using the sonic features provided by The Echo Nest—normalizing each to a 0 to 1 scale so as to not allow any individual attribute undue influence over our similarity calculation, and then collapsing them into a single vector $\boldsymbol{V}_i$ for each song in our dataset.[11] For each song i, we pulled every other song that appeared on the charts during the year prior to song i's debut, and we calculated the cosine similarity between each song-pair's vector of features. The resulting vector $\mathbf{V}_{it}$ includes the cosine similarity between song i and every other song j from the previous 52 weeks' charts, which we consider the boundary of the relevant comparison set against which each song is competing.

After thoughtful consideration, we determined that simply taking the average of each song's row of similarities in $\mathbf{V}_{it}$—in essence, creating a summary typicality score for each song in our dataset—left open the possibility that two songs that looked similar (in terms of their constitutive features) might actually sound different, thus biasing our analysis. Furthermore, research suggests that consumers tend to be split into segments defined by the type of music they consume. These segments, or communities, may or may not align with traditional musicological genre categories, which have their own distinct traditions and histories (cf. Lena 2012, 2015). Although

Table 1. The Echo Nest Sonic Features

Attribute	Scale	Definition
Acousticness	0–1	Represents the likelihood that the song was recorded solely by acoustic means (as opposed to more electronic/electric means).
Danceability	0–1	Describes how suitable a track is for dancing. This measure includes tempo, regularity of beat, and beat strength.
Energy	0–1	A perceptual measure of intensity throughout the track. Think fast, loud, and noisy (i.e., hard rock) more than dance tracks.
Instrumentalness	0–1	The likelihood that a track is predominantly instrumental. Not necessarily the inverse of speechiness.
Key	0–11 (integers only)	The estimated, overall key of the track, from C through B. We enter key as a series of dummy variables.
Liveness	0–1	Detects the presence of a live audience during the recording. Heavily studio-produced tracks score low on this measure.
Mode	0 or 1	Whether the song is in a minor (0) or major (1) key.
Speechiness	0–1	Detects the presence of spoken word throughout the track. Sung vocals are not considered spoken word.
Tempo	Beats per minute (BPM)	The overall average tempo of a track.
Time Signature	Beats per bar/measure	Estimated, overall time signature of the track. 4/4 is the most common time signature by far and is entered as a dummy variable in our analyses.
Valence	0–1	The musical positiveness of the track.

Note: This list of features includes all but one of the attributes provided by The Echo Nest's suite of algorithms: loudness. We cut this variable from our final analysis at the suggestion of the company's senior engineer, who explained that loudness is primarily determined by the mastering technology used to make a particular recording, a characteristic that is confounded through radio play and other forms of distribution.

research shows that omnivorous consumption behavior is now prevalent among certain consumer segments (Lizardo 2014), we believe that the perceived sonic similarity between two songs will decrease if those songs are associated with different genres (e.g., a country song and a reggae song may have similar beat and chord structures, thereby appearing similar when seen as a vector of features, but be perceived to sound quite different by listeners). Thus, we weight each song-pair's raw cosine similarity by the average similarity of those songs' primary genres over the preceding 52 weeks.

We chose to use a genre-weighted cosine similarity measure for two reasons. First, we wanted to generate a fine-grained, feature-based measure, rather than one based purely on shared symbolic classification. Although the latter represents an important signal of how listeners identify and process music, we focus on the former because we believe it provides a more objective representation of a song's sonic fingerprint. Moreover, cosine similarity is a common measure for clustering multidimensional vectors (Evans 2010). Second, we believed it was important to include all songs in a given year, rather than only songs from within a particular genre, when constructing a relevant comparison set to measure typicality. Listeners may be more likely to listen to and compare songs from within the same genre—this is why we chose to incorporate a genre weighting scheme in the first place—but we recognize that for many listeners these genres and their boundaries are not absolute, particularly when it comes to the most mainstream

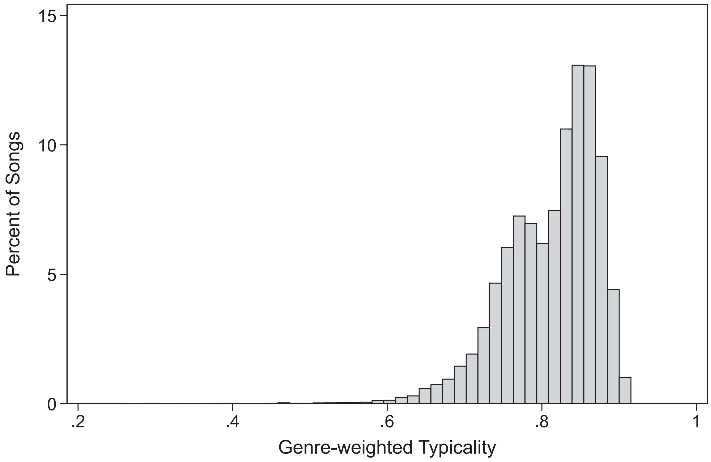

Figure 1. Distribution of Genre-Weighted Song Typicality (Yearly)
Note: The slight dip in this distribution around .80 reflects the binary (0, 1) nature of one of the sonic features included in our typicality measure: mode. Songs written in major and minor keys are equally typical, on average, but the sonic distance between a pair of major and minor songs is likely to be greater than a pair pulled at random, producing the bimodal tendency visualized here.

music captured on the Hot 100 charts. We therefore decided to include songs from all genres when defining the relevant comparison set for our main typicality measure.

To construct the weights for our *genre-weighted typicality (yearly)* measure, we calculated yearly within-genre averages for each sonic feature, and then again used a cosine similarity measure to index the average proximity of each pair of genres in feature space. We then applied the resulting similarities to the raw similarity measure summarized above for each song pair. For example, if a rock song and a folk song had a raw similarity of .75, and the average similarity between rock and folk in year x is .8, then the genre-weighted similarity between those two songs would be .75 × .8 = .6.[12] If both songs were categorized as rock, then the weight would equal 1, and the genre-weighted similarity between songs would be .75. We then calculated the weighted average of each cell in $\mathbf{V}_{it}$ to create the variable used in our main models: a weighted average of each song's distance from all other songs that appeared on the charts in a given year. A simple frequency histogram of this measure provides evidence of the relatively high degree of similarity between songs across our dataset and in popular music more generally ($\mu = .81$; $\sigma = .06$; range = .26 to .92; see Figure 1).[13]

Finally, in addition to our yearly genre-weighted typicality measure, we constructed a second variable, *genre-weighted typicality (weekly)*, to investigate week-to-week competition between songs, which we test in our final set of models as a robustness check. Rather than calculating a single typicality score for each song based on its similarity to songs that charted during the 52 weeks prior to its chart debut, we calculated a unique typicality score for each week a song appears on the charts. To do this, we first measured the cosine similarity between each song's vector of features and those of other songs with which it shared a chart. For each week, we created a matrix $\mathbf{A}_t$ that has dimensions matching the number of songs on each week's charts (100 x 100), with cell A_{ijt} representing the similarity between song i and song j for that week. Because every song is perfectly similar to itself, we removed **A**'s diagonal from all calculations. As with our yearly typicality measure, we again weighted each cell in **A** by the similarity of each song-pair's genres from the year in which those songs were released. Once these weights were applied, we took the average of each row to give each song-week a *genre-weighted typicality (weekly)* value. This measure is designed to capture how similar a song is to other songs with which it is directly competing on the charts.

Additional predictor variables. We collected a handful of control variables to account for the multifaceted nature of musical production and ensure the robustness of our effects. First, we included a dummy variable coded to 1 if a song was released on a major record label, and 0 if it was from an independent label. Major labels typically have larger marketing budgets, higher production quality, closer ties with radio stations (see Rossman 2012), and bigger stars on their artist rosters. These factors suggest that songs released by major labels will not only appear more regularly on the charts (two-thirds of the songs in our dataset are major label releases), but that major label releases should have a comparatively easier time reaching the top of the charts. We include the major label dummy in all analyses to account for the benefits that such songs receive when striving to hit the top of the charts.

Second, we included a set of dummy variables in each model to account for the number of songs an artist had previously placed on the charts. Musicians receive different levels of institutional support (e.g., marketing or PR), which can affect their opportunities for success, but these differences are difficult to ascertain. These previous song count dummies capture artists' relative visibility or popularity at the moment of a song's release: (1) if a song is an artist's first on the charts, (2) if it is her second or third song on the charts, (3) if it is her 4th through 10th song on the charts, or (4) if she has had more than 10 songs in the Hot 100. These dummies also help capture "superstar" effects (Krueger 2005), which could account for the cumulative advantage popular artists experience as their songs become more likely to climb to the top of the charts.

We also constructed a variable called *multiple memberships* to account for artists who released songs under different names or band formations. Annie Lennox, for example, appears on the charts as a member of the Eurythmics and as a solo artist. The Eurythmics represent Lennox's first appearance on the charts, so we coded every subsequent appearance of hers as a solo artist as 1 for *multiple memberships*. We did this for every artist who appeared with multiple bands (or with a band and as a solo artist) on the Hot 100 (roughly 6 percent of our dataset). For these artists, we continued song counts from previous chart incarnations—meaning Lennox's first charting song as a solo artist was coded as her 15th song overall, because the Eurythmics charted 14 songs before she released her first solo hit. Whether a function of artists' skill in creating chart-friendly songs, labels' commitment to already established artists, or fans' loyalty to certain musicians, maintaining a comprehensive count of previous songs on the charts helps us account for any potential benefits chart veterans receive.

Fourth, we included a variable called *long song*, set to 1 if a song was unusually long, and 0 otherwise. Historical recording formats, along with radio, encouraged artists to produce songs that are shorter in length, typically between three and four minutes long (Katz 2010). Although the average length of a song on the charts has increased over time, longer songs were likely to get cut short or have trouble finding radio airtime during much of the timeframe covered by our data. We include this dummy to account for the possibility that these difficulties affect chart performance. For our analysis, any song that was two standard deviations longer than the average song for the year in which it was released was denoted a *long song*.

Fifth, we account for "crossover" songs—that is, songs affiliated with multiple genres and thus (potentially) appealing to multiple audiences.[14] In addition to the Hot 100, *Billboard* has several other, predominantly genre-based, charts to capture songs' popularity: mainstream rock, R&B, country, and others. Songs that cross genres and fandom boundaries may be more likely to succeed on the generalist chart (Brackett 1994; Lena 2012), although one could also argue that difficult-to-classify songs may suffer as the result of audience confusion (see Pontikes 2012; Zuckerman 1999). To capture the potential effects of genre-spanning, we created a variable *crossover*: This dummy is coded 1 for any song with more than one song-level genre designation (e.g., blues and country), *unless the two genre designations are pop and rock*, which for much of the chart's history were considered interchangeable and too mainstream to classify across multiple distinct fan bases. *Crossover* is coded as 0 for songs with only one genre classification. Using this method, roughly 24 percent of the songs in our data are considered crossovers, and on average they perform slightly better on the Hot 100 charts (average peak chart position of 43 versus 45 for crossovers and non-crossovers, respectively; t-test: –3.636, p = .0001).

Sixth, we constructed a dummy variable *reissue* for any song that was re-released and appeared on the charts for a second time. As an example, Prince's track "1999" originally charted in 1982, reaching #12 and staying on the charts for 27 weeks. It was reissued for New Years in 1999, when it charted again for a week. Such songs, already familiar to audiences and likely reissued due to their initial popularity, should have an easier time performing well on the charts when they re-enter them. To account for this potential advantage, we coded any song that was re-released in this manner as a *reissue*, and we included the dummy in all analyses.

Finally, we included nonparametric time dummies to account for historical variation in our results, partitioning 59 years of data into five-year blocks. We did this for two reasons. First, we wanted to capture the fact that producer and consumer tastes, as well as the sounds and boundaries of certain genres, change over time. Second, we needed a way to account for changes in the underlying calculation and meaning of chart rankings, particularly before and after the move to use SoundScan data (for further details and analyses, see note 8 and Part B of the online supplement). Using half-decade dummies allows us to estimate and control for the effects of these changes, which had an immediate impact on chart dynamics but took time to be fully understood and absorbed by industry stakeholders (Anand and Peterson 2000). Table 2 summarizes descriptive statistics and correlations for all the key variables in our analysis.

Estimation Strategy

To examine the relationship between songs' sonic features and their performance on the Hot 100 charts, we first ran pooled, cross-sectional OLS regressions for each of our two static outcome variables, *peak chart position (inverted)* and *weeks on charts* (Models 1 and 2). These models, run on the 25,102 songs for which we have complete data, are intended to provide correlational face validity of a relationship between our sonic features and chart outcomes.

To conduct a more formal test of the relationship between song typicality and chart performance, and to account for the fact that our *peak chart position* outcome variable is composed of discrete whole numbers derived from ranks, we ran a second set of models using an ordered logit specification (Models 3 and 4). We included various artist-level control variables (e.g., previous song and multiple band membership dummies) instead of artist fixed effects, because ordered logit models with fixed effects can produce inconsistent estimators (see Baetschmann, Staub, and Winkelmann 2015). Models estimating *weeks on charts* contain the same control variables but are estimated using zero-truncated negative binomial regression (Hilbe 2011), as the outcome is a count variable with a minimum value of one (Models 5 and 6).

The models described above are cross-sectional analyses using a song's typicality when it first appeared on the charts to predict its overall success. We know, however, that the Hot 100 charts are dynamic: they are released weekly and change just as frequently, with potentially dozens of songs entering, exiting, and shifting positions. Songs move an average of seven ranks from one week to the next, and they tend to have a relatively short shelf-life in the spotlight, with an average chart lifespan of only 11.5 weeks. Following the logic of our earlier prediction, we believe that songs' weekly chart movements

Table 2. Correlations and Descriptive Statistics for Select Variables in the Analyses

	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]	[12]	[13]	[14]	[15]	[16]	[17]	[18]	[19]	[20]	[21]
[1] Peak chart position (inverted)	1																				
[2] Weeks on charts	.72	1																			
[3] Genre-weighted typicality (yearly)	−.04	−.09	1																		
[4] Genre-weighted typicality (weekly)	−.03	−.08	.99	1																	
[5] Major label dummy	.07	.10	−.06	−.06	1																
[6] Long song	.03	.00	−.06	−.06	.02	1															
[7] Crossover track	.02	.03	−.04	−.04	.04	.00	1														
[8] Multiple memberships	.05	.01	−.01	−.01	.03	.02	.00	1													
[9] 2 to 3 previously charting songs	−.08	−.02	.00	.00	−.04	−.03	−.01	−.11	1												
[10] 4 to 10 previously charting songs	.01	−.01	.01	.01	.05	.00	−.01	−.01	−.32	1											
[11] > 10 previously charting songs	.05	−.06	.01	.01	.06	.03	.02	.26	−.32	−.40	1										
[12] Song tempo	−.01	−.02	.07	.07	−.01	−.02	−.01	.00	.00	−.01	.01	1									
[13] Song energy	−.01	.05	.16	.16	.01	−.03	.00	.03	.00	.00	−.02	.16	1								
[14] Song speechiness	−.01	.05	−.10	−.10	.00	.02	.00	−.02	.02	−.02	−.04	.01	.09	1							
[15] Song acousticness	−.04	−.19	.00	−.01	−.08	.00	−.02	−.05	−.01	.01	.01	−.08	−.56	−.11	1						
[16] Minor/major mode (0 or 1)	−.01	−.06	.66	.65	−.01	−.02	−.05	.00	.00	.00	.01	.03	−.07	−.10	.13	1					
[17] Song danceability	.03	.13	.12	.13	−.02	−.06	.07	.02	.03	−.03	−.05	−.14	.16	.19	−.28	−.14	1				
[18] Song valence	.00	−.05	.30	.30	−.11	−.10	.00	−.01	.04	−.02	−.06	.09	.30	.03	−.11	−.04	.47	1			
[19] Song instrumentalness	.00	−.03	−.30	−.30	−.05	.02	.01	.00	.02	−.03	−.05	.01	−.08	−.07	.12	−.02	−.03	.03	1		
[20] Song liveness	.05	.00	−.11	−.11	.01	.02	−.04	.01	−.01	.03	.00	.01	.15	.08	.02	.02	−.22	−.07	−.02	1	
[21] Song time signature = 4/4	.05	.08	.09	.10	.03	−.03	.00	.03	.01	.00	.00	−.03	.27	.01	−.27	−.05	.26	.18	−.07	.01	1
Mean	56.27	11.57	.81	.81	.67	.04	.24	.08	.21	.29	.28	119.09	.59	.07	.34	.74	.58	.62	.08	.24	.90
Standard Deviation	3.45	7.79	.06	.06	.47	.19	.42	.27	.40	.45	.45	27.70	.22	.08	.31	.44	.16	.24	.21	.22	.29
Minimum	1.00	1.00	.26	.26	0	0	0	0	0	0	0	0	0	.02	0	0	0	0	0	.01	0
Maximum	10.00	87.00	.92	.92	1.00	1.00	1.00	1.00	1.00	1.00	1.00	242.51	1.00	.96	1.00	1.00	.99	1.00	1.00	1.00	1.00

will also be influenced by their sonic differentiation from the competition on the Hot 100. Thus, our final set of models examines the dynamic effect of typicality on inter-song competition.

To conduct this analysis, we model the *weekly* change in songs' chart position as a function of their *genre-weighted typicality (weekly)*. Note that this measure changes as new songs cycle in and out of the charts week-to-week. These models (7 and 8) include linear and quadratic control variables for the number of weeks a song has already been on the charts, as well as song-level fixed effects, which allow us to control for the time-invariant factors of each song, including the artist, the record label, the marketing budget, the song's individual sonic features, and all artist- and song-level controls included in Models 3 through 6. All independent and control variables are lagged one week—to (1) match the "natural" one week window used by *Billboard*, and (2) account for the constant short-term churn within the charts, which would render longer lags substantively meaningless. These and all other models presented here use robust standard errors.

RESULTS

Before presenting our main results, we first explore the historical relationship between song typicality and chart position. Figures 2a and 2b present a descriptive analysis of this relationship; they indicate substantial evolution in the typicality of charting songs over the life-course of our data. To construct these graphs, we took the average typicality of songs during their first week on the charts, and then compared (2a) songs that reached the top 40 with those that did not, and (2b) songs that reached number one with those that did not. Figure 2a indicates that songs peaking in the top 40 are comparable in their typicality to songs failing to reach this status. In fact, in the early years of the charts, top-40 songs are slightly *more* typical than the songs that peaked in positions 41 through 100. Conversely, Figure 2b indicates that, aside from a few years in which the average number-one hit was more typical than the average song on the charts, the most successful songs tended to be *less* typical than other songs, though that gap has narrowed in recent years. Although the average typicality of number-one songs is significantly different from that of their peers, they remain close enough to provide *prima facie* support for our optimal differentiation hypothesis.

It is worth noting the general trends in song typicality across our dataset: the chart's early history was marked by more homogenous, "typical" songs, with more atypical songs appearing in the 1970s, 1980s, and 1990s. This trend toward greater atypicality has reversed in recent years, as songs appearing on the charts after 2000 seem to be growing more typical. These trends tell us something interesting about *absolute* levels of feature-based typicality over time, but the models that follow allow us to assess how a song's typicality *relative to its contemporaneous competition* affects its performance on the charts.

Results from our first pair of multivariate models are depicted in Figure 3. This plot shows standardized estimates of the relationship between songs' sonic features, artists' previous success, and chart performance.[15] These results provide preliminary evidence that some of the sonic features in our dataset are significantly correlated with songs' chart performance, above and beyond the effects of genre, artist, and label affiliation. In Model 1 (represented with white circles), we find that a song's danceability, liveness, and the presence of a 4/4 time signature (as opposed to all other time signatures) lead to a more favorable peak chart position. Alternatively, energy (intensity/noise), speechiness, and acousticness lead to a lower peak position on the charts. Although theorizing the interpretation of these feature-specific effects is beyond the scope of this article, they provide some face validity that product features matter for songs' chart performance.

In addition to adjusting for social- and status-related effects on songs' chart position, the dummies for artists' song count reveal evidence of a "sophomore slump." This term refers to the common perception that musicians

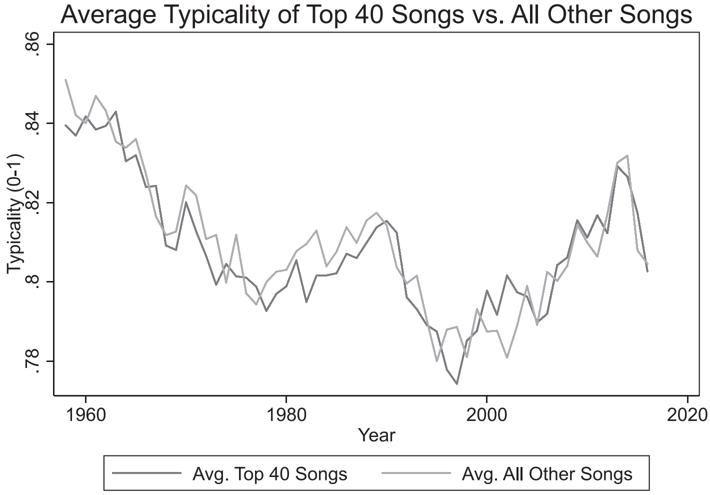

Figure 2a. Comparison of Average Typicality for Top-40 Songs and All Other Songs, 1958 to 2016

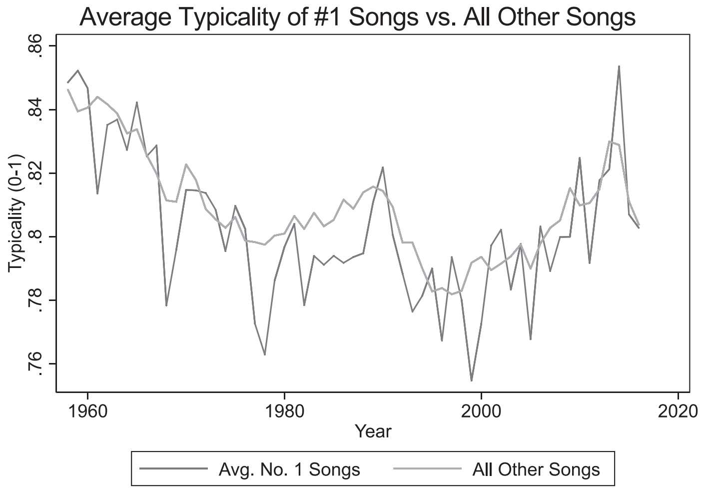

Figure 2b. Comparison of Average Typicality for #1 Songs and All Other Songs, 1958 to 2016

often fail to produce a second song or album as popular as their first. Our results provide evidence for this, as artists' second and third hit songs do not perform as well as their first. However, the positive coefficient for songs released by artists with more than 10 previous hits provides support for "superstar" effects, suggesting that artists experience cumulative advantage after achieving repeated success.

In Model 2 (represented with black circles), we estimate the effect of these same variables on songs' longevity on the charts (in weeks). These results reveal a similar pattern of effects. One difference is worth noting, however: although we find evidence of a sophomore slump, this effect does not reverse as an artist's number of previous hits increases. In other words, if an artist has already charted four or more songs, then subsequent hits will be more likely to experience shorter chart lives. This suggests that audiences may be more likely to grow tired of music released by artists they already know.

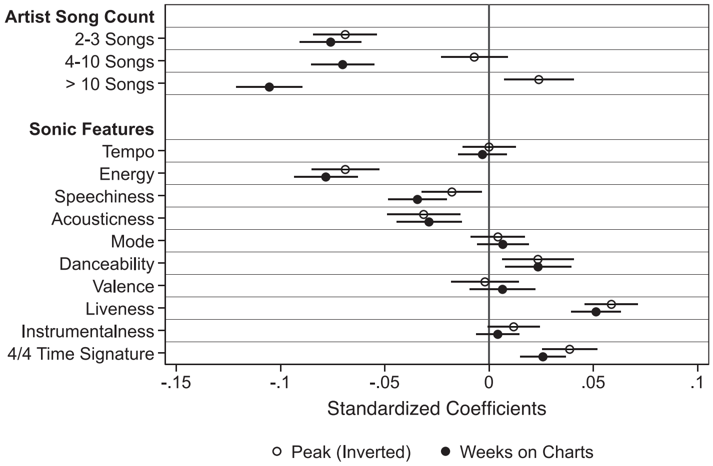

Figure 3. Select Standardized Coefficients from Pooled, Cross-Sectional OLS Models Predicting *Billboard* Hot 100 Peak Chart Position and Longevity (Models 1 and 2)
Note: Horizontal bars represent 95% CI. See Appendix Table A1 for full (unstandardized) results.

Results from Models 1 and 2 are instructive and provide *prima facie* evidence that sonic features are meaningfully correlated with chart performance. Nevertheless, because they appear independently in these models, the results reveal little about how bundles of features—that is, songs—are similar to or different from each other as whole, or how such differentiation affects chart performance. To address these questions, we move to our next set of models, which use our typicality measure to test how songs' differentiation across feature space affects their performance on the charts. We first discuss the results for our typicality variable of interest across models before examining key control variables.

Table 3 summarizes the coefficients for our key independent and control variables from Models 3 through 6 (see Appendix Table A2 for the full set of coefficient estimates). Recall that Model 3 is an ordered logit regression predicting a song's peak position using the song's typicality relative to other songs that appeared on the charts in the previous 52 weeks. The results show that song typicality leads to lower peak chart position. After controlling for genre affiliation, artist popularity, and a host of other song- and artist-level variables, songs that are more similar to their peers are less likely to reach the top of the charts.

In Model 4, we add a squared typicality term to test for the hypothesized inverted U-shaped relationship between typicality and chart performance. Results support the prediction, revealing the benefits of optimal differentiation. The most atypical songs in our dataset would benefit from being more similar to their peers, but as songs become more and more similar, this relationship is reversed—exhibiting too much typicality is associated with poorer chart performance.

Quadratic terms in ordered logit models are difficult to interpret (Karaca-Mandic, Norton, and Dowd 2012). To aid interpretation, Figure 4 plots the marginal effects of songs' typicality on their peak chart position. For purposes of clarity and interpretability, we partitioned peak chart position into meaningful intervals represented by the different lines in the figure. We then used the coefficients from Model 4 to calculate the marginal probability of songs with different typicality levels reaching certain peak positions.

Moving from the top of the figure to the bottom (i.e., from the worst position on the charts to the best), we find that the most atypical and most typical songs are likely to fall *outside* of the top 40 (the white and black circles). These two curves do not reflect the inverted-U shape that we find in Model 4

Table 3. Select Variables from Pooled, Cross-Sectional Ordered Logit and Negative Binomial Models Predicting *Billboard* Hot 100 Peak Chart Position and Longevity, 1958 to 2016

	3. Ordered Logit	4. Ordered Logit	5. Negative Binomial	6. Negative Binomial
Outcome Variable:	Peak Position (Inverted)	Peak Position (Inverted)	Weeks on Charts	Weeks on Charts
Genre-weighted typicality (yearly)	−2.419**	7.672*	−.538**	1.791
	(.429)	(2.987)	(.150)	(1.051)
Genre-weighted typicality (yearly)^2		−6.805**		−1.570*
		(2.004)		(.698)
Major label dummy	.145**	.145**	.0246**	.0245**
	(.0255)	(.0255)	(.00883)	(.00882)
Long song	.262**	.265**	.0291	.0290
	(.0609)	(.0608)	(.0193)	(.0193)
2 to 3 previously charting songs	−.306**	−.306**	−.138**	−.138**
	(.0353)	(.0353)	(.0119)	(.0119)
4 to 10 previously charting songs	−.0305	−.0298	−.118**	−.118**
	(.0331)	(.0331)	(.0108)	(.0108)
10+ previously charting songs	.0874*	.0878*	−.168**	−.168**
	(.0347)	(.0347)	(.0115)	(.0115)
Crossover track	.151**	.149**	−.00556	−.00590
	(.0303)	(.0303)	(.0107)	(.0107)
Multiple memberships	.146**	.147**	.0554**	.0559**
	(.0417)	(.0417)	(.0133)	(.0133)
Reissued track	−.204*	−.204*	−.0812*	−.0814*
	(.0923)	(.0921)	(.0409)	(.0409)
Half-Decade Dummies				
1987 to 1991	.265**	.232**	.440**	.432**
	(.0697)	(.0702)	(.0217)	(.0218)
1992 to 1996	−.282**	−.328**	.567**	.557**
	(.0701)	(.0714)	(.0239)	(.0241)
Observations	25,077	25,077	25,077	25,077

Note: Robust standard errors are in parentheses. Reference categories for dummy variables: pop (genre), independent label, first charting song (previously charting songs), key of E-flat, and all non-4/4 time signatures.
*p < .05; **p < .01 (two-tailed tests).

across the entirety of our dataset. This makes sense: songs that sound too much (or not enough) like their peers have a higher likelihood of staying outside the top of the charts. The remaining curves—which predict likelihoods of reaching the top 40, top 20, top 10, top 5, and #1, respectively—all show the expected inverted-U shape relationship, albeit with decreasing likelihoods as each echelon becomes more difficult (and unlikely) for songs to reach. The songs that climb to the top of the charts have a higher marginal probability of doing so when they are in the middle of the typicality distribution—that is, when they are optimally differentiated.

As an example of what constitutes an optimally differentiated song, The Beatles' 1969 hit "Come Together" reached the top of the charts on November 29, 1969, and featured a typicality score of .66 the week it debuted—over two standard deviations less typical than the average song released that year. Digging into the song's individual features, we find that much of its novelty can be attributed to its low energy (1.2σ below the mean) and low valence (1.9σ below the mean). Although this

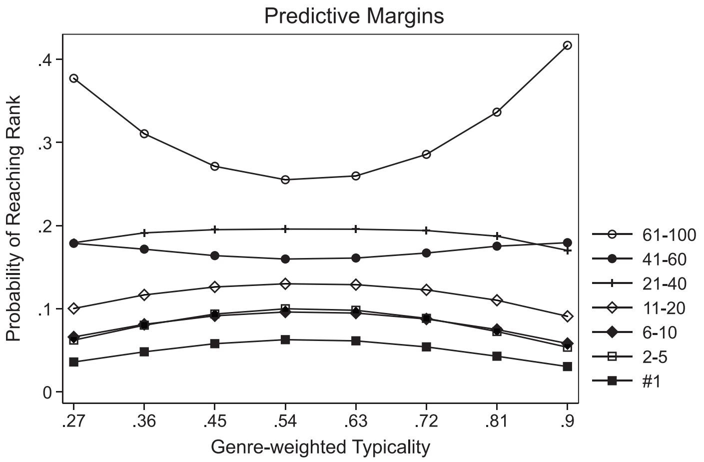

Figure 4. Predicted Marginal Probability of Songs Achieving Selected Peak Position (by Typicality) from Ordered Logit Model (Model 4)
Note: Although we inverted chart position in our models to assist readers with a more straightforward interpretation (e.g., positive coefficients reflect better performance), we revert to the originally coded chart positions for our marginal-effects graphical analysis. In the figure, the predicted positions are coded as they would be on the charts (i.e., #100 is the lowest, #1 the highest).

example does not statistically represent our entire dataset, it does speak to some of the factors that drive our typicality measure and song differentiation in general.

Models 5 and 6 use zero-truncated negative binomial regression to estimate the effect of song typicality on chart longevity. When entered as a linear term, typicality is again negatively associated with length of stay on the charts. However, when we include the squared term (Model 6), we find the expected inverted U-shaped relationship. For the most novel songs in our dataset, higher levels of typicality would increase their odds of remaining on the charts, whereas the most typical songs would remain in the spotlight longer if they were more differentiated. *Ceteris paribus*, a song that is a single standard deviation below mean typicality (.75 versus .81) is likely to remain on the charts for roughly a half week longer than a song at the mean (11.5 weeks and 11 weeks, respectively).

Across Models 3 through 6, we find that songs are more likely to attract and maintain consumers' attention if they are differentiated from other songs on the charts, but not so different that they fail to meet prevailing expectations. We also find theoretically intuitive results for several of our key control variables. For example, songs released by major labels tend to reach higher chart positions and to last longer on the charts, as we anticipated. Somewhat surprisingly, however, we find that song length is positively related to chart performance (Models 3 and 4). This could be attributed to a few outliers (e.g., Don McLean's "American Pie" is 8:36 minutes and spent a month at number 1; The Beatles' "Hey Jude" clocks in at 7:11 and spent nine weeks at number 1), or it could be evidence of yet another mechanism through which songs achieve some degree of differentiation (although this would not be picked up by our typicality variable). This result seems to indicate that long songs are more salient to listeners than their average-length peers.

As in models predicting peak chart position (1 and 2), we again find support for an artist's sophomore slump and for superstar effects. When looking at the dummies for artists' previous success (the reference category here is an artist's first song on the chart), we find that artists' second and third songs do not exhibit the same longevity as their chart debuts. In addition, songs released by artists with more than 10 previously charting songs reach higher chart positions than do artists' first songs, but they do not stay on the charts as long.

The positive coefficient of *multiple memberships* provides further evidence of superstar effects. This result suggests that veteran musicians, having already amassed a following as a solo artist or member of a band, are likely to see their songs perform better when they hit the charts under a different moniker. Having a pre-established fan base surely benefits artists who, having already proven themselves capable of producing hits, decide to go solo, form a new band, or join a different band altogether. Similarly, we find that crossovers benefit from broader audience appeal: songs that span multiple genres are more likely to climb to the top of the charts, although they do not appear to stay on the charts any longer than their single-genre peers.[16]

Finally, we turn to the half-decade time dummies, with the chart's first several years (1958 to 1961) comprising the omitted reference group. Recall that the five years before (1987 to 1991) and after (1992 to 1996) the introduction of SoundScan are of particular interest (1991 is included in the pre-SoundScan era, as the change took place in November of that year). We find that songs were more likely to perform better on the charts prior to SoundScan; in every period thereafter it has become more difficult to reach the top of the charts. Moreover, results from Models 5 and 6 reveal that songs released in the late 1980s and 1990s remained on the charts longer than they did during the earliest years of the Hot 100. This is especially true for songs released directly after the introduction of SoundScan. These results support Anand and Peterson's (2000) claim about how the shift in chart-ranking calculation slowed chart churn.

The results presented thus far support our hypothesis that optimally differentiated songs perform better on the charts in general. They do not, however, allow us to speak to the relationship between song typicality and weekly changes in chart position. To examine these week-to-week dynamics, we turn to the results of our fixed-effects models, presented in Table 4.

In Model 7, the coefficient for *genre-weighted typicality (weekly)* is negative and significant, indicating that songs sounding more similar to their peers are likely to see their performance suffer in subsequent weeks (recall that all covariates and controls are lagged one week in these models). Controlling for the natural decay that songs typically experience on the charts (i.e., the negative coefficient for *weeks on charts*), a single standard deviation increase in typicality results in a song descending more than an *additional* .6 positions each week—this is substantial given the relatively low debut position of most songs on the charts (i.e., #82). The squared *weeks on charts* coefficient is small but positive, reflecting the ever-diminishing distance that songs can drop as they remain on the charts week after week.

In Model 8 we add a quadratic term for song typicality. We find that, all else equal, more typical songs tend to fare worse on subsequent weeks' charts than do songs that are optimally differentiated. Indeed, only the most novel songs in our dataset benefit from being more similar to the songs around them, suggesting that some degree of typicality is beneficial for success. This means songs from heavily underrepresented genres—or songs from mainstream genres that are particularly unique—benefit from the entrance of similar sounding songs, or "sonic neighbors," on the charts. These songs may serve as a kind of bridge for listeners to compare and reconsider songs that are otherwise distinctive. Conversely, songs that would otherwise be deemed too atypical by audiences may perform better when there are even more unusual songs already on the charts. For the majority of observations in our dataset, however, increased levels of typicality suggest a subsequent drop in chart position.

DISCUSSION

Our results provide evidence that the features of cultural products affect consumption behavior, both in terms of their intrinsic qualities and in the way they structure how audiences compare and evaluate products in relational space (cf. de Vaan et al. 2015; Lena

Table 4. Results of Fixed-Effects Models Predicting *Billboard* Hot 100 Songs' Weekly Change in Position, 1958 to 2016

Outcome Variable:	7. Change in (Inverted) Chart Position	8. Change in (Inverted) Chart Position
Genre-weighted typicality (weekly)	−10.84**	37.98*
	(2.323)	(15.43)
Genre-weighted typicality (weekly)^2		−32.44**
		(10.33)
Weeks (on charts)	−1.941**	−1.941**
	(.0129)	(.0129)
Weeks (on charts)^2	.0334**	.0334**
	(.000478)	(.000477)
Constant	21.95**	3.805
	(1.873)	(5.848)
Observations	263,715	263,715
R-squared	.432	.432

Note: Robust standard errors are in parentheses.
*$p < .05$; **$p < .01$ (two-tailed test).

and Pachucki 2013). Controlling for many of the social and industry-specific factors that contribute to a song's success, we find that listeners' assessments of popular music are shaped, in part, by the qualities of songs themselves, perhaps suggesting that consumers are more discerning than we sometimes give them credit for (cf. Salganik et al. 2006).

Revisiting our initial question, "What makes popular culture popular?" we can add to the list of explanations: (1) the underlying features of products, and (2) the relative position of those products within feature space. Our empirical proxy for this second explanation—typicality, a concept that can easily be adapted to other domains of cultural analysis—significantly predicts how songs perform on the *Billboard* Hot 100 charts. Specifically, we find that most popular songs suffer a penalty for being too similar to their peers, although this effect is attenuated and even reversed for the most novel songs. These effects extend to songs' overall performance, which we measured using peak chart position and longevity, and week-to-week changes in chart position. Our findings support the prediction that songs that manage the similarity–differentiation (or familiarity–novelty) tradeoff are more likely to achieve success.

These findings provide important insights into the consumption dynamics of a multi-billion-dollar cultural market. However, we also recognize several important limitations. Although the data we use to measure sonic features are relatively comprehensive and sophisticated, they represent a substantial distillation of a song's musical complexity. Reducing such a high dimensional object into 11 fixed features inevitably simplifies its cultural fingerprint and alters its relationships to other like-objects. As MIR tools improve, so too will our ability to map connections between songs. Our data also do not allow us to account for listeners' idiosyncratic experience of features or lyric similarity between songs. Moreover, the bounded nature of the Hot 100—it includes only songs that achieve enough success to appear on the charts in the first place—raises the issue of selection bias and the generalizability of our conclusions. Based on the discussion in Part A of the online supplement, we believe that while the most popular cultural products are slightly more typical *on average*, the difference is not so vast as to circumscribe our findings. The patterns we encountered likely extend beyond our sample in the music industry and to other creative industries as well.

Our analyses show that, conditional on entering the charts, certain songs outperform others, suggesting that not all popular culture is created equal. In future research, we hope to conduct more fine-grained, dynamic analyses to understand the nature and implications of specific idiosyncrasies that appear in our dataset. Carving the chart into distinct segments, estimating effects for different time periods, and identifying scope conditions for the arguments presented here will undoubtedly provide additional insight into the dynamic and historically contingent nature of our findings.

Additionally, although we provide robust evidence for how musical features affect songs' chart performance, our explanation of evaluation outcomes is limited to characteristics of the production environment. Thus, the analyses presented here do not account for external consumption dynamics, making it difficult to identify the cognitive mechanisms that drive listeners' selection decisions. Although we suspect that the patterns of optimal differentiation we find are relevant across empirical domains, it remains unclear whether and how these findings could be extended, or whether the concepts herein can prove fruitful for researchers interested in the ecological dynamics of products that are firmly outside the creative industries.

We expect, however, that the curvilinear relationship between typicality and popularity will carry over to other realms of cultural production, such as art, television, and movies. Even the biggest budget productions will likely be viewed less favorably than their competition if audiences perceive them to be derivative or too similar to existing productions. We believe that the continued use of the concepts and measures developed in this article can be generative in a variety of empirical contexts, serving as a useful tool for social scientists interested in how product features shape consumption behavior. More generally, we hope scholars continue to try and integrate production- and consumption-side narratives to highlight the interdependencies between these processes and their associated outcomes.

Conclusions

Without denying the role of social dynamics, it is important to acknowledge and theorize the effect of product features on success in cultural markets. Both independently and in concert, the content and qualities of cultural products need to be considered more seriously when investigating success in cultural markets. In this article, we showed how a song's feature-derived position amongst its competition—whether considered over the span of a year or a week—contributes to its success. We hope this article, including its data and methodological approach, can serve as a model for more content-driven explorations of large-scale empirical puzzles in cultural sociology and beyond.

This article makes several important contributions. First, we import methods traditionally associated with computer science and big data analytics to enhance our understanding of large-scale consumption dynamics and performance outcomes. These tools necessarily simplify the intrinsic high-dimensionality of cultural objects, but they also empower us to generate new insights in historically opaque contexts. Although many new measurement tools originate from advances in computer science and other disciplines, social scientists must critically develop and apply them appropriately and thoughtfully (Bail 2014). Other scholars have mapped meaning structures (Mohr 1994, 1998), charted diffusion patterns (Rossman 2012), and introduced the link between cultural content and consumption behavior (Jones et al. 2012; Lena 2006). However, there has been no systematic attempt to theorize and measure how product features influence the emergence and diffusion of consumption patterns. In this article, we introduced and exploited a rich dataset capable of exploring these dynamics, generating new insights into the world of popular music and cultural markets more broadly.

Second, we measured and tested the effects of product features and the associations they generate among audiences. Our conception of feature-similarity space can serve as a tool to map ecosystems of cultural products, and as a means to understand selection dynamics in markets that require subjective evaluation. The system of associations between products is theoretically and analytically distinct from—although integrally connected to and mediated through—networks of producers and consumers. This raises the possibility that cultural content and qualities assert an autonomous influence over evaluation outcomes through product crowding and differentiation.

This conceptualization of culture is dynamic and will ideally push scholars to continue developing new ways to talk about culture and its consequences. One path forward involves importing the tools of network science to study perceived similarities and associations between cultural products. Although existing research on networks focuses largely on interpersonal or interorganizational ties, substantive relationships exist between all sorts of actors, objects, and ideas (Breiger and Puetz 2015). These relationships serve as conduits for information or signals of quality (Podolny 2001), but also as a spatial metaphor for how markets are structured (Emirbayer 1997). Continuing to redefine what constitutes "nodes" and "edges" might help scholars rethink how cultural objects of all types—including products, practices, and ideas—assert influence or agency, thereby addressing a critical issue in social theory more broadly (Berger and Luckmann 1966). Such a reconception may also change how scholars think about taste formation, which will no longer reside in a theoretical black box.

Taking these ideas about culture and agency a step further, the dynamics of optimal differentiation also provide a mechanism to support and explain endogenous cultural change (cf. Kaufman 2004; Lieberson 2000). If optimally differentiated products perform better at time t, producers seeking success are likely to try and replicate those products in the future. However, given a growing population of producers trying to match the attributes of successful products, and the inevitable lag between production and consumption, the most popular products released at time $t + 1$ will likely come not from producers who earlier chose a replication strategy, but from those who release products that are now optimally differentiated from the competition at $t + 1$. As this pattern continues, popular culture will shift and evolve, with products becoming more (and less) typical over time, just as we saw in Figures 2a and 2b. The most successful producers, to paraphrase a well-known saying, will aim to produce something for where the cultural context is headed, rather than where it currently resides.

Third, our conceptualization of products and feature space contributes to the literature on categories and market structure (e.g., Kovács and Hannan 2015; Pontikes 2012). Research in this area has explored the origins and consequences of categorical classification on firms and products. However, our results suggest a more fine-grained approach may be necessary to understand how markets are structured by categories. Combinations of features likely play an integral role in the way products, organizations, and even individuals are perceived and evaluated. In our analysis, we included both product features (sonic attributes) and category labels (genres) to ensure that a computer-driven reduction in complexity did not cause inappropriate interpretation.

Future work should dive even deeper into the relationship between features and labels. For example, how do product features help create the categorical structure of musicological genres? To draw an analogy, researchers have looked at networks of recipe ingredients on the one hand (Teng, Lin, and Adamic 2012), and the categorization of food and its consequences for market outcomes on the other (Kovács and Johnson 2014; Rao, Monin, and Durand 2003). Integrating these perspectives to explore the relationship between ingredients and the way food is categorized and evaluated appears to be an obvious next step. We hope our findings encourage category scholars to work toward this integration in the study of music, food, and beyond.

Finally, our findings speak to the inherent difficulty—and folly—in practicing "hit song science" (Dhanaraj and Logan 2005; Pachet and Roy 2008). It is certainly true that a small cabal of writers and producers are responsible for many of the most popular songs in recent years (Seabrook 2015). It is also true that artists have more tools and data at their disposal than ever before. These tools provide them with incredibly detailed information about the elements of popular songs, which might in turn help them craft their own hits (Thompson 2014). Nevertheless, although writing recognizable tunes may become easier with the emergence of these tools, our results suggest that artists trying to reverse engineer a hit song may be neglecting two important points. First, songs that sound too similar to the competition are going to have a more difficult time attracting and holding audience attention. Second, and most important, the characteristics of contemporaneous songs produced by other artists will have a significant impact on that song's success. Content *and* context matter. Because a song's reception is partially contingent on how differentiated it is from its peers, and artists cannot precisely forecast or control which songs are released concurrently with their own, the crafting of a hit song should be more art than science.

APPENDIX: FULL MODEL RESULTS

Table A1. Results from Pooled, Cross-Sectional OLS Models Predicting *Billboard* Hot 100 Peak Chart Position and Longevity (Models 1 and 2 in Figure 3)

Outcome Variable:	1. Peak Position (Inverted)	2. Weeks on Charts
Major Label Dummy	2.544**	.298**
	(.431)	(.101)
Long Song	4.220**	.165
	(1.007)	(.228)
2 to 3 Previously Charting Songs	−5.201**	−1.465**
	(.589)	(.145)
4 to 10 Previously Charting Songs	−.477	−1.210**
	(.550)	(.133)
10+ Previously Charting Songs	1.616**	−1.825**
	(.576)	(.140)
Crossover Track	2.543**	−.0812
	(.510)	(.124)
Multiple Memberships	2.279**	.550**
	(.706)	(.158)
Reissued Track	−3.297*	−.882*
	(1.642)	(.415)
Genre Dummies		
Blues	−12.17**	−.734
	(1.954)	(.383)
Brass and Military	−8.505	−3.422
	(12.84)	(1.760)
Children's	5.264	−3.011
	(7.229)	(1.600)
Classical	−5.536	.436
	(8.032)	(1.923)
Electronic	1.259	1.238**
	(.862)	(.235)
Folk, World, and Country	−8.663**	.102
	(.822)	(.208)

(continued)

Table A1. *(continued)*

Outcome Variable:	1. Peak Position (Inverted)	2. Weeks on Charts
Funk/Soul	−4.299**	.316*
	(.721)	(.155)
Hip Hop	1.472	.508
	(.962)	(.272)
Jazz	−11.98**	−.678**
	(1.333)	(.246)
Latin	−18.15**	−3.137**
	(3.430)	(.832)
Non-music	.149	−.0113
	(4.154)	(.793)
Reggae	−3.166	.873
	(3.974)	(1.010)
Rock	1.004	.672**
	(.637)	(.151)
Stage and Screen	−4.968	−.685
	(4.937)	(1.324)
Sonic Features		
Tempo	−1.33e–05	−.000887
	(.00713)	(.00168)
Energy	−9.710**	−2.822**
	(1.171)	(.282)
Speechiness	−6.409*	−3.160**
	(2.645)	(.659)
Acousticness	−3.098**	−.728**
	(.886)	(.202)
Mode (1 = Major Key)	.297	.118
	(.464)	(.113)
Danceability	4.584**	1.176**
	(1.719)	(.405)
Valence	−.253	.213
	(1.073)	(.266)
Instrumentalness	1.686	.154
	(.914)	(.194)
Liveness	8.100**	1.812**
	(.898)	(.214)
Key = C	−.685	−.331
	(.953)	(.243)
Key = C-Sharp	.705	.152
	(1.020)	(.270)
Key = D	−1.232	−.519*
	(.985)	(.249)
Key = E-Flat	1.440	.170
	(1.246)	(.312)
Key = E	−1.085	−.495
	(1.039)	(.264)
Key = F	−.215	−.255
	(1.014)	(.255)
Key = F-Sharp	−1.305	−.396
	(.961)	(.246)

(continued)

Table A1. *(continued)*

Outcome Variable:	1. Peak Position (Inverted)	2. Weeks on Charts
Key = G	.408	−.213
	(1.096)	(.278)
Key = G-Sharp	.252	−.216
	(.978)	(.248)
Key = A	−.132	−.229
	(1.053)	(.267)
Key = B-Flat	.194	−.119
	(1.065)	(.282)
4/4 Time Signature Dummy	4.010**	.686**
	(.700)	(.149)
Half-Decade Dummies		
1962 to 1966	−.624	−.644**
	(.885)	(.140)
1967 to 1971	−1.549	−.329*
	(.924)	(.154)
1972 to 1976	1.810	1.592**
	(.977)	(.177)
1977 to 1981	3.809**	3.304**
	(.996)	(.206)
1982 to 1986	2.865**	3.753**
	(1.044)	(.215)
1987 to 1991	5.547**	4.856**
	(1.113)	(.231)
1992 to 1996	−2.379*	6.881**
	(1.143)	(.280)
1997 to 2001	.638	7.401**
	(1.141)	(.297)
2002 to 2006	.500	7.583**
	(1.155)	(.302)
2007 to 2011	−4.455**	4.294**
	(1.101)	(.305)
2012 to 2016	−5.970**	5.527**
	(1.224)	(.380)
Constant	54.82**	9.537**
	(2.015)	(.460)
Observations	25,102	25,102
R-squared	.047	.171

Note: Robust standard errors are in parentheses. Reference categories are pop (genre), independent label, first charting song (previously charting songs), key of E-flat (key), and all time signatures other than 4/4.
*$p < .05$; **$p < .01$ (two-tailed tests).

Table A2. Results from Pooled, Cross-Sectional Ordered Logit Models Predicting *Billboard* Hot 100 Peak Chart Position and Longevity, 1958 to 2013 (see Table 3)

	3. Ordered Logit	4. Ordered Logit	5. Negative Binomial	6. Negative Binomial
Outcome Variable:	Peak Position (Inverted)	Peak Position (Inverted)	Weeks on Charts	Weeks on Charts
Genre-Weighted Typicality (yearly)	−2.419**	7.672*	−.538**	1.791
	(.429)	(2.987)	(.150)	(1.051)
Genre-Weighted Typicality (yearly)^2		−6.805**		−1.570*
		(2.004)		(.698)
Major Label Dummy	.145**	.145**	.0246**	.0245**
	(.0255)	(.0255)	(.00883)	(.00882)
Long Song	.262**	.265**	.0291	.0290
	(.0609)	(.0608)	(.0193)	(.0193)
2 to 3 Previously Charting Songs	−.306**	−.306**	−.138**	−.138**
	(.0353)	(.0353)	(.0119)	(.0119)
4 to 10 Previously Charting Songs	−.0305	−.0298	−.118**	−.118**
	(.0331)	(.0331)	(.0108)	(.0108)
10+ Previously Charting Songs	.0874*	.0878*	−.168**	−.168**
	(.0347)	(.0347)	(.0115)	(.0115)
Crossover Track	.151**	.149**	−.00556	−.00590
	(.0303)	(.0303)	(.0107)	(.0107)
Multiple Memberships	.146**	.147**	.0554**	.0559**
	(.0417)	(.0417)	(.0133)	(.0133)
Reissued Track	−.204*	−.204*	−.0812*	−.0814*
	(.0923)	(.0921)	(.0409)	(.0409)
Genre Dummies				
Blues	−.727**	−.732**	−.123**	−.123**
	(.111)	(.111)	(.0446)	(.0447)
Brass and Military	−1.173	−.911	−.481*	−.413*
	(1.003)	(.961)	(.223)	(.210)
Children's	−.157	−.0278	−.347	−.334
	(.461)	(.431)	(.210)	(.204)
Classical	−.723	−.642	−.0555	−.0241
	(.484)	(.507)	(.164)	(.167)
Electronic	.122*	.127*	.100**	.101**
	(.0520)	(.0521)	(.0190)	(.0190)
Folk, World, and Country	−.477**	−.482**	.0134	.0123
	(.0463)	(.0464)	(.0189)	(.0190)
Funk/Soul	−.250**	−.250**	.0230	.0229
	(.0421)	(.0422)	(.0148)	(.0148)
Hip Hop	.127*	.130*	.0458*	.0467*
	(.0576)	(.0576)	(.0212)	(.0212)
Jazz	−.749**	−.750**	−.0947**	−.0947**
	(.0816)	(.0815)	(.0291)	(.0291)
Latin	−1.240**	−1.224**	−.353**	−.343**
	(.213)	(.217)	(.0919)	(.0926)
Non-music	−.298	−.201	−.117	−.0991
	(.215)	(.213)	(.0851)	(.0836)
Reggae	−.389	−.323	.0442	.0548
	(.263)	(.256)	(.0764)	(.0760)
Rock	.0569	.0581	.0638**	.0640**
	(.0375)	(.0375)	(.0137)	(.0137)

(continued)

Table A2. *(continued)*

Outcome Variable:	3. Ordered Logit Peak Position (Inverted)	4. Ordered Logit Peak Position (Inverted)	5. Negative Binomial Weeks on Charts	6. Negative Binomial Weeks on Charts
Stage and Screen	−.588**	−.502	−.113	−.0806
	(.278)	(.280)	(.139)	(.141)
Sonic Features				
Tempo	1.94e–05	5.75e–05	−3.24e–05	−2.80e–05
	(.000419)	(.000419)	(.000146)	(.000146)
Energy	−.473**	−.442**	−.216**	−.209**
	(.0717)	(.0722)	(.0242)	(.0244)
Speechiness	−.524**	−.539**	−.252**	−.258**
	(.157)	(.156)	(.0531)	(.0531)
Acousticness	−.187**	−.170**	−.0763**	−.0722**
	(.0527)	(.0528)	(.0178)	(.0180)
Mode (1 = Major Key)	.246**	.304**	.0597**	.0727**
	(.0489)	(.0511)	(.0173)	(.0179)
Danceability	.397**	.425**	.138**	.143**
	(.104)	(.104)	(.0358)	(.0358)
Valence	.0844	.118	.0324	.0400
	(.0666)	(.0670)	(.0233)	(.0236)
Instrumentalness	−.104	−.143*	−.0148	−.0231
	(.0644)	(.0650)	(.0214)	(.0216)
Liveness	.399**	.378**	.139**	.134**
	(.0549)	(.0552)	(.0182)	(.0184)
Key = C	−.109*	−.129*	−.0383*	−.0428*
	(.0544)	(.0547)	(.0189)	(.0190)
Key = C-Sharp	.00510	−.00605	.0143	.0117
	(.0584)	(.0585)	(.0203)	(.0203)
Key = D	−.0928	−.0934	−.0371	−.0374
	(.0552)	(.0553)	(.0193)	(.0193)
Key = E-Flat	.0897	.100	.0270	.0291
	(.0719)	(.0720)	(.0254)	(.0254)
Key = E	−.0459	−.0332	−.0205	−.0179
	(.0586)	(.0588)	(.0208)	(.0208)
Key = F	.0226	.0371	−.00155	.00175
	(.0582)	(.0584)	(.0200)	(.0200)
Key = F-Sharp	.0396	.0557	.0171	.0206
	(.0630)	(.0632)	(.0223)	(.0223)
Key = G	−.0510	−.0322	−.0198	−.0157
	(.0548)	(.0551)	(.0191)	(.0191)
Key = G-Sharp	.0459	.0609	−.00407	−.000956
	(.0630)	(.0632)	(.0218)	(.0218)
Key = A	.0227	.0327	−.00348	−.00139
	(.0548)	(.0549)	(.0189)	(.0189)
Key = B-Flat	.000810	.00501	−.00452	−.00398
	(.0603)	(.0603)	(.0207)	(.0207)
4/4 Time Signature Dummy	.262**	.266**	.0742**	.0747**
	(.0414)	(.0413)	(.0147)	(.0147)

(continued)

Table A2. *(continued)*

	3. Ordered Logit	4. Ordered Logit	5. Negative Binomial	6. Negative Binomial
Outcome Variable:	Peak Position (Inverted)	Peak Position (Inverted)	Weeks on Charts	Weeks on Charts
Half-Decade Dummies				
1962 to 1966	−.0647	−.0707	−.0964**	−.0980**
	(.0530)	(.0529)	(.0177)	(.0177)
1967 to 1971	−.156**	−.179**	−.0586**	−.0637**
	(.0556)	(.0559)	(.0188)	(.0189)
1972 to 1976	.0234	−.0107	.157**	.150**
	(.0606)	(.0614)	(.0201)	(.0202)
1977 to 1981	.0811	.0418	.315**	.307**
	(.0606)	(.0616)	(.0213)	(.0215)
1982 to 1986	.0425	.00613	.353**	.345**
	(.0638)	(.0647)	(.0215)	(.0216)
1987 to 1991	.265**	.232**	.440**	.432**
	(.0697)	(.0702)	(.0217)	(.0218)
1992 to 1996	−.282**	−.328**	.567**	.557**
	(.0701)	(.0714)	(.0239)	(.0241)
1997 to 2001	−.108	−.156*	.603**	.593**
	(.0709)	(.0726)	(.0245)	(.0247)
2002 to 2006	−.0931	−.136	.623**	.614**
	(.0704)	(.0716)	(.0249)	(.0251)
2007 to 2011	−.350**	−.379**	.392**	.385**
	(.0665)	(.0670)	(.0277)	(.0277)
2012 to 2016	−.414**	−.433**	.495**	.492**
	(.0731)	(.0734)	(.0305)	(.0305)
Observations	25,077	25,077	25,077	25,077

Note: Robust standard errors are in parentheses. Reference categories for dummy variables: pop (genre), independent label, first charting song (previously charting songs), key of E-flat, and all non-4/4 time signatures.
*$p < .05$; **$p < .01$ (two-tailed tests).

Acknowledgments

For their valuable ideas, comments, and support, we are grateful to Matt Bothner, Frédéric Godart, Alejandro Mantecón Guillén, Brayden King, Jenn Lena, John Levi Martin, John Mohr, Klaus Weber, Erich Ludwig, and the rest of The Echo Nest team, as well as members of the University of Chicago's Knowledge Lab and Northwestern University's Institute on Complex Systems (NICO). We also thank Omar Lizardo, the rest of the *ASR* editorial team, and three anonymous reviewers for their thoughtful and thought-provoking feedback, which made the paper better.

Supplemental Material

Please find the online supplement for this article on *ASR*'s website. It includes two parts: Part A provides a comparative analysis of songs that appear in the Hot 100 with a sample of those that do not; Part B investigates the effects and implications of the introduction of SoundScan and other historical variations on chart outcomes.

Notes

1. To avoid repetition, we use the terms "features," "attributes," and "characteristics" interchangeably to refer to the fixed, material elements that constitute cultural products. For example, in the context of a painting, relevant features might include the different colors used, and whether the painting is a portrait or a landscape.
2. We recognize that category labels might themselves be considered just another product attribute or feature, but we treat them here as distinct entities. This distinction is analytical as well as phenomenological, as category labels convey a qualitatively different kind of information than do the underlying features of products.

3. In this article, we use the term "success" to connote mass or popular appeal, rather than critical acclaim or other legitimate measures of performance.
4. It is worth noting here that channels of influence between networks and taste run in both directions (Lizardo 2006). Just as social networks can alter cultural outcomes, so too can those networks be altered by prevailing tastes and practices, recasting culture and social structure as mutually constitutive (Pachucki and Breiger 2010; Vaisey and Lizardo 2010). This view—one that highlights culture's causal role—is supported by the "strong program" in cultural sociology (e.g., Alexander and Smith 2002) and related work on the materiality of culture (Rubio 2012). Rather than passive symbolic structures, culture is endowed with real properties that can influence actors' preferences, behaviors, and affiliation patterns.
5. We do not explicitly invoke network terminology to describe our theory—in part because we do not use network measures to test it—but the notion of a product "association network" can serve as a salient image to help visualize this space. Although networks have historically been used to study information transfer between people, groups, or organizations, they are increasingly used in a variety of contexts, including the study of co-occurrences of or associations between narrative elements (Smith 2007), cultural objects (Breiger and Puetz 2015), multimedia content (Meng and Shyu 2012), and even food flavors (Ahn et al. 2011) and human genes (Schafer and Strimmer 2005). In the context of music, the nodes in the network would be songs, and the edges between them might represent varying degrees of feature overlap or similarity.
6. Another familiar metaphor that approximates this idea is the cultural "milieu" or "fabric." This concept encompasses the population of cultural products that producers or consumers have access to in a given context. In the market for popular music, this might include all current and previously released songs, which can be connected to one another, however distantly, based on their shared feature sets. The theoretical and empirical implications of this idea extend beyond the scope of this article, but the imagery of a cultural fabric may help motivate our rationale for extending the concept of networks to cultural products and their constitutive features.
7. Because we focus on songs that appear in the Hot 100, our analysis may suffer from considerable selection bias, an issue we address in Part A of the online supplement. However, we believe that any bias in our data does not present a major limitation, as charting songs constitute an appropriate sample for answering our initial research question: what makes popular culture popular? Furthermore, this sample is consistent with studies that explore the differentiated outcomes for cultural products that get shortlisted for prizes versus those that win (e.g., Kovács and Sharkey 2014; Sorensen 2007). Other factors, such as artist popularity and marketing support, play an important role in driving certain songs into the Hot 100, but we are primarily interested in understanding why, conditional on entering the charts, certain songs outperform others.
8. The initial algorithm for determining the charts included a combination of radio airplay and a survey of selected record stores across the country. This methodology had several flaws, as it relied on human reporting for a large portion of the input and was therefore subject to personal biases and external influence. In November 1991, *Billboard* replaced self-reported sales data with SoundScan's point-of-sale data from most of the record stores in the United States (for more on the history of the algorithm and the consequences of the shift, see Anand and Peterson [2000]). We ran a series of supplementary analyses to test how this development influences our results, and we found that the effect of SoundScan falls within the range of expected historical variation across our dataset (see Part B of the online supplement).
9. In case consumer selection decisions occur at the artist rather than song level, we also ran our models using artist-level genre attributions. Results are consistent.
10. One of the weaknesses in our data is that these genre codes were applied in the early twenty-first century, rather than the year in which each song was originally released. Although genre attributions are admittedly dynamic (Lena and Peterson 2008), we believe it is reasonable to assume that historical attributions are for the most part consistent with our data. Furthermore, although genres appear and disappear over the course of our data, and those that persist have evolved, such changes have their provenance predominantly at the subgenre or style level (e.g., "hard rock" and "roots rock" versus "rock"). Using primary, song-level genre assignments means misattributions are unlikely or should be relegated to fringe cases.
11. Each feature is weighted equally to calculate our pairwise cosine similarity measure. We considered prioritizing certain features over others (e.g., weighting tempo more heavily than mode), but conversations with musicologists and computer scientists specializing in MIR provided no consistent rationale for using weights. Moreover, the 11 features included in our analysis were designed to encompass the most important dimensions of songs in a relatively evenhanded and comparable way, with the possible exception of mode, which has a slightly outsized influence due to its binary (0,1) rather than continuous scale.
12. When two songs were the only representatives of their respective genres over the previous year (a rare occurrence, largely confined to the early years of the chart), we used the minimum similarity between any pair of genres from the year prior to the focal song's debut week to construct our weighted measure. For example, if a focal song has

a primary genre of *vocal*, and is the only such track to appear for an entire year on the charts, then we used the minimum weight (i.e., the largest distance between two genres' vectors of average features) as the weighting multiplier for that song's cosine similarity with every other song on the charts during the previous year.

13. In another set of models (available by request), we checked the robustness of our results vis-à-vis different levels of reliance on genre classification. To do this, we calculated two additional typicality variables—*all pair typicality (yearly)* and *within-genre typicality (yearly)*—which seek to provide further evidence that our results hold across multiple specifications. *All pair typicality* is again a cosine similarity measure, but it is the simple, unweighted average of each song with all other songs that appeared on the charts in the previous 52 weeks (see also Part B of the online supplement). It is our main independent variable without any genre-based weighting. *Within-genre typicality* is, as its name implies, the average cosine similarity between each song and the average feature vector for all other songs affiliated with the same genre in a given year. This version of the measure captures how typical a song is for its given genre. We found substantively similar results using both of these variables across Models 3 through 6.
14. We are grateful to an anonymous reviewer for bringing this issue to our attention.
15. All control variables—including genre affiliation, dummies for each musical key (C through B), and dummies for major label, long song, multiple memberships, crossover, reissue, and half-decade time—are included in these models but are not shown in the figure (see Appendix Table A1 for full results).
16. In addition to including the crossover dummy in our models, we also ran separate versions of Models 4 and 6 for crossover songs and non-crossover songs. Our main findings hold for *non*-crossover songs—they benefit from being optimally differentiated—but not for crossovers. However, crossovers do comprise a higher proportion of #1 songs (30 percent) than their overall chart presence would suggest (24 percent of all songs are crossovers by our measure).

References

Ahn, Yong-Yeol, Sebastian E. Ahnert, James P. Bagrow, and Albert-László Barabási. 2011. "Flavor Network and the Principles of Food Pairing." *Scientific Reports* 1:196. doi:10.1038/srep00196.

Alexander, Jeffrey C., and Philip Smith. 2002. "The Strong Program in Cultural Theory: Elements of a Structural Theory." Pp. 135–50 in *Handbook of Sociological Theory*, edited by J. Turner. New York: Kluwer Academic/Plenum Publishers.

Alexander, Peter J. 1996. "Entropy and Popular Culture: Product Diversity in the Popular Music Recording Industry." *American Sociological Review* 61(1):171–74.

Anand, N., and Richard A. Peterson. 2000. "When Market Information Constitutes Fields: Sensemaking of Markets in the Commercial Music Industry." *Organization Science* 11(3):270–84.

Anderson, John R. 1991. "The Adaptive Nature of Human Categorization." *Psychological Review* 98(3):409–429.

Baetschmann, Gregori, Kevin E. Staub, and Rainer Winkelmann. 2015. "Consistent Estimation of the Fixed Effects Ordered Logit Model." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 178(3):685–703.

Bail, Christopher A. 2014. "The Cultural Environment: Measuring Culture with Big Data." *Theory and Society* 43(3–4):465–82.

Barroso, Alicia, Marco S. Giarratana, Samira Reis, and Olav Sorenson. 2014. "Crowding, Satiation, and Saturation: The Days of Television Series' Lives." *Strategic Management Journal* 37(3):565–85.

Becker, Howard. 1982. *Arts Worlds*. Berkeley: University of California Press.

Berger, Jonah, and Gael Le Mens. 2009. "How Adoption Speed Affects the Abandonment of Cultural Tastes." *Proceedings of the National Academy of Sciences* 106(20):8146–50.

Berger, Peter, and Thomas Luckmann. 1966. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. Garden City, NY: Doubleday & Company.

Bielby, William T., and Denise D. Bielby. 1994. "'All Hits Are Flukes': Institutionalized Decision Making and the Rhetoric of Network Prime-Time Program Development." *American Journal of Sociology* 99(5):1287–1313.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *The Journal of Machine Learning Research* 3:993–1022.

Bothner, Matthew S., Jeong-han Kang, and Toby E. Stuart. 2007. "Competitive Crowding and Risk Taking in a Tournament: Evidence from NASCAR Racing." *Administrative Science Quarterly* 52(2):208–247.

Bourdieu, Pierre. 1993. *The Field of Cultural Production: Essays on Art and Literature*. New York: Columbia University Press.

Brackett, David. 1994. "The Politics and Practice of 'Crossover' in American Popular Music, 1963 to 1965." *The Music Quarterly* 78(4):774–97.

Brashears, Matthew E., and Eric Quintane. 2015. "The Microstructures of Network Recall: How Social Networks Are Encoded and Represented in Human Memory." *Social Networks* 41:113–26.

Breiger, Ronald L., and Kyle Puetz. 2015. "Culture and Networks." Pp. 557–62 in *International Encyclopedia of Social and Behavioral Sciences*, Vol. 5, edited by J. D. Wright. Oxford, UK: Elsevier.

Brewer, Marilynn B. 1991. "The Social Self: On Being the Same and Different at the Same Time." *Personality and Social Psychology Bulletin* 17(5):475–82.

Carroll, Glenn R., Olga M. Khessina, and David G. McKendrick. 2010. "The Social Lives of Products: Analyzing Product Demography for Management Theory and Practice." *The Academy of Management Annals* 4(1):157–203.

Caves, Richard E. 2000. *Creative Industries: Contracts between Art and Commerce*. Cambridge, MA: Harvard University Press.

Cerulo, Karen A. 1988. "Analyzing Cultural Products: A New Method of Measurement." *Social Science Research* 17(4):317–52.

Chan, Cindy, Jonah Berger, and Leaf Van Boven. 2012. "Identifiable but Not Identical: Combining Social Identity and Uniqueness Motives in Choice." *Journal of Consumer Research* 39(3):561–73.

Deephouse, David L. 1999. "To Be Different, or to Be the Same? It's a Question (and Theory) of Strategic Balance." *Strategic Management Journal* 20(2):147–66.

DellaPosta, Daniel, Yongren Shi, and Michael W. Macy. 2015. "Why Do Liberals Drink Lattes?" *American Journal of Sociology* 120(5):1473–1511.

de Vaan, Mathijs, Balazs Vedres, and David Stark. 2015. "Game Changer: The Topology of Creativity." *American Journal of Sociology* 120(4):1144–94.

Dhanaraj, Ruth, and Beth Logan. 2005. "Automatic Prediction of Hit Pop Songs." *Proceedings of the International Society for Music Information Retrieval* pp. 488–91.

DiMaggio, Paul. 1987. "Classification in Art." *American Sociological Review* 52(4):440–55.

DiMaggio, Paul, Manish Nag, and David Blei. 2013. "Exploiting Affinities between Topic Modeling and the Sociological Perspective on Culture: Application to Newspaper Coverage of U.S. Government Arts Funding." *Poetics* 41:570–606.

Douglas, Mary, and Baron C. Isherwood. 1996. *The World of Goods: Towards an Anthropology of Consumption*. New York: Routledge.

Dowd, Timothy J. 2004. "Concentration and Diversity Revisited: Production Logics and the U.S. Mainstream Recording Market, 1940–1990." *Social Forces* 82(4):1411–55.

Eastman, Jason T., and Terry F. Pettijohn II. 2015. "Gone Country: An Investigation of Billboard Country Songs of the Year across Social and Economic Conditions in the United States." *Psychology of Popular Media Culture* 4(2):155–71.

Emirbayer, Mustafa. 1997. "Manifesto for Relational Sociology." *American Journal of Sociology* 103(2): 281–317.

Evans, James A. 2010. "Industry Induces Academic Science to Know Less about More." *American Journal of Sociology* 116(2):389–452.

Fleming, Lee, Santiago Mingo, and David Chen. 2007. "Collaborative Brokerage, Generative Creativity, and Creative Success." *Administrative Quarterly* 52(3):443–75.

Fligstein, Neil, and Luke Dauter. 2007. "The Sociology of Markets." *Annual Review of Sociology* 33(1):105–128.

Friberg, Anders, Erwin Schoonderwaldt, Anton Hedblad, Marco Fabiani, and Anders Elowsson. 2014. "Using Listener-Based Perceptual Features as Intermediate Representations in Music Information Retrieval." *Journal of the Acoustical Society of America* 136(4):1951–63.

Frith, Simon. 1996. "Music and Identity." Pp. 108–125 in *Questions of Cultural Identity*, edited by S. Hall and P. Du Gay. London, UK: Sage.

Frow, John. 1995. *Cultural Studies and Cultural Value*. Oxford, UK: Clarendon Press.

Frow, John. 2006. *Genre: The New Critical Idiom*. New York: Taylor and Francis.

Godart, Frédéric C., Andrew V. Shipilov, and Kim Claes. 2014. "Making the Most of the Revolving Door: The Impact of Outward Personnel Mobility Networks on Organizational Creativity." *Organization Science* 25(2):377–400.

Goldberg, Amir. 2011. "Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined." *American Journal of Sociology* 116(5):1397–1436.

Goldberg, Amir, Michael T. Hannan, and Balázs Kovács. 2016. "What Does It Mean to Span Cultural Boundaries: Variety and Atypicality in Cultural Consumption." *American Sociological Review* 81(2):215–41.

Goldberg, Amir, Sameer B. Srivastava, V. Govind Manian, William Monroe, and Christopher Potts. 2016. "Fitting In or Standing Out? The Tradeoffs of Structural and Cultural Embeddedness." *American Sociological Review* 81(6):1190–1222.

Hadida, Allegre. 2015. "Performance in the Creative Industries." Pp. 219–49 in *The Oxford Handbook of Creative Industries*, edited by C. Jones, M. Lorenzen, and J. Sapsed. New York: Oxford University Press.

Hamlen, William A. 1991. "Superstardom in Popular Music: Empirical Evidence." *The Review of Economics and Statistics* 73(4):729–33.

Hannan, Michael T., and John H. Freeman. 1977. "The Population Ecology of Organizations." *American Journal of Sociology* 82(5):929–64.

Hilbe, Joseph M. 2011. *Negative Binomial Regression*, 2nd ed. Cambridge, UK: Cambridge University Press.

Hirsch, Paul M. 1972. "Processing Fads and Fashions: An Organization-Set Analysis of Cultural Industry Systems." *American Journal of Sociology* 77(4):639–59.

Holt, Fabian. 2007. *Genre in Popular Music*. Chicago: University of Chicago Press.

Hsu, Greta. 2006. "Jacks of All Trades and Masters of None: Audiences' Reactions to Spanning Genres in Feature Film Production." *Administrative Science Quarterly* 51(3):420–50.

Hsu, Greta, and Michael T. Hannan. 2005. "Identities, Genres, and Organizational Forms." *Organization Science* 16(5):474–90.

Hsu, Greta, Giacomo Negro, and Fabrizio Perretti. 2012. "Hybrids in Hollywood: A Study of the Production and Performance of Genre Spanning Films." *Industrial and Corporate Change* 21:1427–50.

Huron, David. 2013. "A Psychological Approach to Musical Form: The Habituation–Fluency Theory of Repetition." *Current Musicology* 96:7–35.

Jones, Candace, Massimo Maoret, Felipe G. Massa, and Silviya Svejenova. 2012. "Rebels with a Cause: Formation, Contestation, and Expansion of the De Novo Category 'Modern Architecture,' 1870–1975." *Organization Science* 23(6):1523–45.

Kahl, Steven J. 2015. "Product Conceptual Systems: Toward a Cognitive Processing Model." Pp. 119–46 in *Cognition and Strategy*, Vol. 32, edited by G. Gavetti and W. Ocasio. Bingley, UK: Emerald Group Publishing.

Karaca-Mandic, Pinar, Edward C. Norton, and Bryan Dowd. 2012. "Interaction Terms in Nonlinear Models." *Health Services Research* 47(1):255–74.

Katz, Mark. 2010. *Capturing Sound: How Technology Has Changed Music*. Berkeley: University of California Press.

Kaufman, Jason. 2004. "Endogenous Explanation in the Sociology of Culture." *Annual Review of Sociology* 30:335–57.

Keuschnigg, Marc. 2015. "Product Success in Cultural Markets: The Mediating Role of Familiarity, Peers, and Experts." *Poetics* 51:17–36.

Kovács, Balázs, and Michael Hannan. 2015. "Conceptual Spaces and the Consequences of Category Spanning." *Sociological Science* 2:252–86.

Kovács, Balázs, and Rebeka Johnson. 2014. "Contrasting Alternative Explanations for the Consequences of Category Spanning: A Study of Restaurant Reviews and Menus in San Francisco." *Strategic Organization* 12(1):7–37.

Kovács, Balázs, and Amanda Sharkey. 2009. "The Paradox of Publicity: How Awards Can Negatively Affect the Evaluation of Quality." *Administrative Science Quarterly* 59(1):1–33.

Krueger, Alan B. 2005. "The Economics of Real Superstars: The Market for Rock Concerts in the Material World." *Journal of Labor Economics* 23(1):1–30.

Lancaster, Kelvin. 1975. "Socially Optimal Product Differentiation." *American Economic Review* 65(4):567–585.

La Rue, Jan. 2001. "Fundamental Considerations in Style Analysis." *Journal of Musicology* 18(2):295–312.

Lena, Jennifer C. 2006. "Social Context and Musical Content of Rap Music, 1979–1995." *Social Forces* 85(1):479–95.

Lena, Jennifer C. 2012. *Banding Together: How Communities Create Genres in Popular Music*. Princeton, NJ: Princeton University Press.

Lena, Jennifer C. 2015. "Genre: Relational Approaches to the Sociology of Music." Pp. 149–60 in *Routledge International Handbook of the Sociology of Art and Culture*, edited by L. Hanquinet and M. Savage. New York: Routledge.

Lena, Jennifer C., and Mark C. Pachucki. 2013. "The Sincerest Form of Flattery: Innovation, Repetition, and Status in an Art Movement." *Poetics* 41(3):236–64.

Lena, Jennifer C., and Richard A. Peterson 2008. "Classification as Culture: Types and Trajectories of Music Genres." *American Sociological Review* 73(5):697–718.

Lieberson, Stanley. 2000. *A Matter of Taste: How Names, Fashions, and Culture Change*. New Haven, CT: Yale University Press.

Lizardo, Omar. 2006. "How Cultural Tastes Shape Personal Networks." *American Sociological Review* 71(5):778–807.

Lizardo, Omar. 2014. "Omnivorousness as the Bridging of Cultural Holes: A Measurement Strategy." *Theory and Society* 43(3–4):395–419.

Lounsbury, Michael, and Mary Ann Glynn. 2001. "Cultural Entrepreneurship: Stories, Legitimacy, and the Acquisition of Resources." *Strategic Management Journal* 22(6–7):545–64.

Mark, Noah. 1998. "Birds of a Feather Sing Together." *Social Forces* 77(2):453–85.

Meng, Tao, and Mei-Ling Shyu. 2012. "Leveraging Concept Association Network for Multimedia Rare Concept Mining and Retrieval." *IEEE International Conference on Multimedia and Expo (ICME) Proceedings,* pp. 860–65.

Mohr, John W. 1994. "Soldiers, Mothers, Tramps, and Others: Discourse Roles in the 1907 New York City Charity Directory." *Poetics* 22:327–57.

Mohr, John W. 1998. "Measuring Meaning Structures." *Annual Review of Sociology* 24:345–70.

Montauti, Martina, and Filippo C. Wezel. 2016. "Charting the Territory: Recombination as a Source of Uncertainty for Potential Entrants." *Organization Science* 27(4):954–71.

Negus, Keith. 1992. *Producing Pop: Culture and Conflict in the Popular Music Industry*. London, UK: Edward Arnold.

Nunes, Joseph C., and Andrea Ordanini. 2014. "I Like the Way It Sounds: The Influence of Instrumentation on a Pop Song's Place in the Charts." *Musicae Scientiae* 18(4):392–409.

Pachet, Francois, and Pierre Roy 2008. "Hit Song Science Is Not Yet a Science." *Proceedings of the International Society for Music Information Retrieval*, pp. 355–60.

Pachucki, Mark A., and Ronald L. Breiger. 2010. "Cultural Holes: Beyond Relationality in Social Networks and Culture." *Annual Review of Sociology* 36:205–224.

Peterson, Richard A. 1990. "Why 1955? Explaining the Advent of Rock Music." *Popular Music* 9(1):97–116.

Peterson, Richard A. 1997. *Creating Country Music: Fabricating Authenticity*. Chicago: University of Chicago Press.

Peterson, Richard A., and David C. Berger. 1975. "Cycles in Symbol Production: The Case of Popular Music." *American Sociological Review* 40(2):158–73.

Phillips, Damon J. 2011. "Jazz and the Disconnected: City Structural Disconnectedness and the Emergence of a Jazz Canon, 1897–1933." *American Journal of Sociology* 117(2):420–83.

Phillips, Damon J. 2013. *Shaping Jazz: Cities, Labels, and the Global Emergence of an Art Form*. Princeton, NJ: Princeton University Press.

Phillips, Damon J., and Young-Kyu Kim. 2008. "Why Pseudonyms? Deception as Identity Preservation among Jazz Record Companies, 1920–1929." *Organization Science* 20(3):481–99.

Podolny, Joel M. 2001. "Networks as the Pipes and Prisms of the Market." *American Journal of Sociology* 107(1):33–60.

Podolny, Joel M., Toby E. Stuart, and Michael T. Hannan. 1996. "Networks, Knowledge, and Niches: Competition in the Worldwide Semiconductor Industry." *American Journal of Sociology* 102(3):659–89.

Pontikes, Elizabeth G. 2012. "Two Sides of the Same Coin: How Ambiguous Classification Affects Multiple Audiences' Evaluations." *Administrative Science Quarterly* 57(1):81–118.

Pontikes, Elizabeth G., and Michael T. Hannan. 2014. "An Ecology of Social Categories." *Sociological Science* 1:311–43.

Rao, Hayagreeva, Philippe Monin, and Rodolphe Durand. 2003. "Institutional Change in Toque Vile: Nouvelle Cuisine as an Identity Movement in French Gastronomy." *American Journal of Sociology* 108(4):795–843.

Rosen, Sherwin. 1981. "The Economics of Superstars." *American Economic Review* 71(5):845–58.

Rossman, Gabriel. 2012. *Climbing the Charts: What Radio Airplay Tells Us about the Diffusion of Innovation*. Princeton, NJ: Princeton University Press.

Rossman, Gabriel, and Oliver Schilke. 2014. "Close, But No Cigar: The Bimodal Rewards to Prize-Seeking." *American Sociological Review* 79(1):86–108.

Rubio, Fernando D. 2012. "The Material Production of the Spiral Jetty: A Study of Culture in the Making." *Cultural Sociology* 6(2):143–61.

Salganik, Matthew J., Peter S. Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science* 311(576):854–56.

Salganik, Matthew J., and Duncan J. Watts. 2008. "Leading the Herd Astray: An Experimental Study of Self-Fulfilling Prophecies in an Artificial Cultural Market." *Social Psychology Quarterly* 71(4):338–55.

Salganik, Matthew J., and Duncan J. Watts. 2009. "Web-Based Experiments for the Study of Collective Social Dynamics in Cultural Markets." *Topics in Cognitive Science* 1(3):439–68.

Schafer, Juliane, and Korbinian Strimmer. 2005. "A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics." *Statistical Applications in Genetics and Molecular Biology* 4(1):Article 32.

Seabrook, John. 2015. *The Song Machine: Inside the Hit Factory*, 1st ed. New York City: W. W. Norton & Company.

Serrà, Joan, Álvaro Corral, Marian Boguñá, Martín Haro, and Josep L. Arcos. 2012. "Measuring the Evolution of Contemporary Western Popular Music." *Scientific Reports* 2:521. doi:10.1038/srep00521.

Smith, Tammy. 2007. "Narrative Boundaries and the Dynamics of Ethnic Conflict and Conciliation." *Poetics* 35(1):22–46.

Sonnett, John. 2004. "Musical Boundaries: Intersections of Form and Content." *Poetics* 32(3–4):247–64.

Sorensen, Alan T. 2007. "Bestseller Lists and Product Variety." *The Journal of Industrial Economics* 55(4):715–38.

Storey, John. 2006. *Cultural Theory and Popular Culture: An Introduction*. Upper Saddle River, NJ: Pearson Prentice Hall.

Swaminathan, Anand, and Jacques Delacroix. 1991. "Differentiation within an Organizational Population: Additional Evidence from the Wine Industry." *Academy of Management Journal* 34(3):679–92.

Teng, Chun-Yuen, Yu-Ru Lin, and Lada A. Adamic. 2012. "Recipe Recommendation Using Ingredient Networks." Pp. 298–307 in *Proceedings of the 3rd Annual ACM Web Science Conference*, ACM.

Thompson, Derek. 2014. "The Shazam Effect." *The Atlantic Magazine*, December.

Tversky, Amos. 1977. "Features of Similarity." *Psychological Review* 84(4):327–52.

Uzzi, Brian, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. "Atypical Combinations and Scientific Impact." *Science* 342(6157):468–72.

Uzzi, Brian, and Jarrett Spiro. 2005. "Collaboration and Creativity: The Small World Problem." *American Journal of Sociology* 111(2):447–504.

Vaisey, Stephen, and Omar Lizardo. 2010. "Can Cultural Worldviews Influence Network Composition?" *Social Forces* 88(4):1595–1618.

Whitburn, Joel. 1986. *Joel Whitburn's Pop Memories, 1890–1954: The History of American Popular Music: Compiled from America's Popular Music Charts 1890–1954*. Menomonee Falls, WI: Record Research.

Whitburn, Joel. 1991. *Joel Whitburn's Top Pop Singles, 1955–1990: Compiled from Billboard's Pop Singles Charts, 1955–1990*. Menomonee Falls, WI: Record Research.

Wijnberg, Nachoem M., and Gerda Gemser. 2000. "Adding Value to Innovation: Impressionism and the Transformation of the Selection System in Visual Arts." *Organization Science* 11(3):323–29.

Yogev, Tamar. 2009. "The Social Construction of Quality: Status Dynamics in the Market for Contemporary Art." *Socio-Economic Review* 8(3):511–36.

Zajonc, Robert B. 1968. "Attitudinal Effects of Mere Exposure." *Journal of Personality and Social Psychology* 9(2p2):1–27.

Zerubavel, Noam, Peter S. Bearman, Jochen Weber, and Kevin N. Ochsner. 2015. "Neural Mechanisms Tracking Popularity in Real-World Social Networks." *Proceedings of the National Academy of Sciences (PNAS)* 112(49):15072–77.

Zhao, Eric Y., Greg Fisher, Michael Lounsbury, and Danny Miller. 2017. "Optimal Distinctiveness:

Broadening the Interface between Institutional Theory and Strategic Management." *Strategic Management Journal* 38(1): 93–113.

Zuckerman, Ezra W. 1999. "The Categorical Imperative: Securities Analysts and the Illegitimacy Discount." *American Journal of Sociology* 104(5):1398–1438.

Zuckerman, Ezra W. 2016. "Optimal Distinctiveness Revisited: An Integrative Framework for Understanding the Balance between Differentiation and Conformity in Individual and Organizational Identities." Pp. 183–99 in *Handbook of Organizational Identity*, edited by M. G. Pratt, M. Schultz, B. E. Ashforth, and D. Ravasi. Oxford, UK: Oxford University Press.

Noah Askin is an Assistant Professor of Organizational Behavior at INSEAD on their European campus. His research focuses on social and cultural networks, the antecedents and consequences of creativity and innovation (particularly in the creative industries), the production and consumption of culture, and the dynamics of organizational and individual status. He holds a joint PhD from the Booth School of Business and the Department of Sociology at the University of Chicago.

Michael Mauskapf is an Assistant Professor of Management at Columbia Business School. His research investigates the dynamics of innovation and success in cultural markets and creative industries. Other research interests include categories, computational social science, cultural entrepreneurship, and the formation and diffusion of taste. His dissertation was awarded the Art Stinchcombe Dissertation Prize in Organization Studies, and was recognized as a finalist in the 2016 INFORMS/Organization Science Dissertation Proposal Competition. He earned degrees from the University of Pennsylvania (BA), the University of Michigan (MA, PhD), and Northwestern University (PhD).