NORTHWESTERN UNIVERSITY

Three Essays on the Dynamics of Creativity and Success in a Cultural Market

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Management and Organizations

By

Michael Gerard Mauskapf

EVANSTON, ILLINOIS

March 2020

Abstract

In this dissertation I leverage new data from the global music recording industry to study the social foundations of creativity and the relationship between product novelty, gender, and commercial success. In Chapter 1, I investigate how different kinds of social connection influence the creation of novel cultural products. Using data on over 25,000 musical artists and 600,000 songs, I construct a feature-based measure of song novelty and then estimate how musicians' collaboration networks, category memberships, organizational affiliations, and geographic neighbors affect their propensity to produce novel work. Results suggest that the most significant predictor of future song novelty is membership in a category or genre populated by other creative artists. This and other findings suggest that creativity stems not only from contact with diverse ideas accessed through collaboration, but also from exposure to creative alters with common cultural and organizational ties.

In Chapter 2, I use a subsample of the same data to explore the relationship between gender and creative output. I find no mean difference between men and women in terms of the novelty of the songs they produce, but after controlling for the size and gender composition of an artist's collaboration network, as well as the gender composition of her primary genre, I find that female artists create significantly *more* distinctive songs than their male counterparts. These results suggest that structural and cultural factors—rather than differences in raw ability—are responsible for gender disparities in creative production. To conclude, in Chapter 3 I unpack the relationship between atypicality and commercial success. Using additional data from *Billboard's* Hot 100 charts, I find that a song's perceived sonic proximity to its peers influences its position on the charts. Contrary to the claim that all popular music sounds the same, I find that songs

sounding too much alike—those that are highly typical—are less likely to succeed, while those exhibiting some degree of optimal differentiation are more likely to rise to the top of the charts.

Acknowledgements

Writing this dissertation was a labor of love—and at times something closer to hate—but we have finally arrived at the finish line. I say "we" because this document reflects the hard work of not one but many. What follows is an abbreviated list of the people who helped me along the way, with apologizes to anyone who I've left out. First, thank you to my dissertation chair and advisor, Klaus Weber, for showing so much kindness and patience while shepherding me through the process. Thanks also to my other committee members – Willie Ocasio, Brian Uzzi, and Damon Phillips – and the rest of the MORS/Kellogg faculty and student body for taking such a sincere interest in my development as a scholar and a person. Any mistakes found in the pages that follow are my own, but all merits should be attributed in part to Noah Askin, who is a close collaborator on all of this work, and a closer friend. Various chapters also benefited from the contributions of additional co-authors Eric Quintane, Joeri Mol, Sharon Koppman; research assistant Adam Kaplan; anonymous reviewers at the *American Sociological Review*; and countless conference participants and other scholars who have provided invaluable feedback and encouragement. The latter stages of the dissertation were completed while on faculty at Columbia Business School, and my new colleagues have been nothing but supportive as wrapped up the process. Never in my wildest dreams could I have anticipated the ups and downs that have brought me to this point, but any successes I've enjoyed along the way are a direct result of the unyielding support of my friends and family, and especially my wife, Rachel, without whom this dissertation would not have been written.

Table of Contents

List of Tables

List of Figures

Introduction

This dissertation has two primary objectives. First, I hope to enrich and expand current explanations of how cultural markets operate. These markets—which include fields as diverse as architecture, design, film, literature, music, television, and the visual art—have long stood at the center of social life, yet our understanding of their complex production and consumption dynamics is only beginning to come into focus. Cultural markets are unlike other markets in that they are largely taste-based and organized around complex production functions and subjective evaluation criteria, characteristics that make them difficult to study at scale. Thus, most extant research in this area takes one of two flavors: qualitative research that focuses on rich, case-based description and contextual understanding (at the expense of uncovering larger-scale patterns and dynamics), or "large-n" quantitative research that treats cultural markets as a case of some more generic market-based phenomenon (at the expense of contextual understanding). I leverage advances in the availability of large-scale, content-rich data on cultural products and markets to reconcile these competing approaches.

Second, I hope to reconsider how creativity is organized in these markets, and to what effect. Creativity is a central tenet of cultural production. The work of musicians, painters, writers, chefs, and even academics is concerned in part with producing something new, exciting, and different than what came before —not only as a means to an end, but as an end in and of itself. At the same time, art worlds are organized around conventions and routines—"people like what they know"—and the relationship between novelty and success is complicated, and can vary by audience type and other contextual factors. I theorize the dual role social context and cultural content play in organizing creative production, and (again) take advantage of advances in data availability to develop a new measure of creative output and empirically test the relationship between creativity, success, and a host of other related constructs.

I tackle these questions in the context of the global music recording industry, which is an ideal context to research the content-contingent nature of cultural production, creativity, and commercial success. As alluded to above, I employ data from a host of digital sources—including The Echo Nest, Spotify, Musicbrainz, Discogs, and Billboard—to study these dynamics at population scale, but with an eye toward aggregating fine-grained contextual information describing songs and the artists that produce them. Details about the data used for individual papers can be found in the "Data and Methods" section of each chapter, but it was culled from a self-built "parent" database that describes the production and consumption properties of all available commercially recorded songs released between 1955 and 2016. For the purposes of the dissertation, the most interesting data consists of ten algorithmically-derived sonic features assigned to each song, providing a simple but relatively intuitive and objective representation of how music sounds and is experienced by producers and consumers. I rely on this data in all three chapters to construct a novel measure of musical creativity (operationalized as "song novelty" in Chapters 1 and 2, and its inverse, "typicality," in Chapter 3). Each of the chapters is summarized below.

In Chapter 1 (**The Social Organization of Creativity: A Multi-Dimensional Perspective**), I investigate how different kinds of social connection influence the creation of novel cultural products. Leveraging original data on over 25,000 musical artists and 600,000 songs, I construct a feature-based measure of creative output (i.e., song novelty) and then estimate how musicians' collaboration networks, category memberships, organizational affiliations, and geographic neighbors affect their propensity to produce novel work. Results suggest that the most significant predictor of future song novelty is membership in a category or genre populated by other creative artists. This and other findings suggest that creativity stems not only from contact with diverse ideas accessed through collaboration, but also from exposure to creative alters with common cultural and organizational ties. Understanding how creative potential travels across these different "spheres of influence" generates new insights into the production of novelty in music and the social organization of creativity itself.

In Chapter 2 (**Do Women Produce More Novel Cultural Products than Men? The Gendered Effects of Networks and Genres on Musical Creativity**) I explore the relationship between gender and creative output. I find no mean difference between men and women in terms of the novelty of the songs they produce, but after controlling for the size and gender composition of an artist's collaboration network, as well as the gender composition of her primary genre, I find that female artists create significantly *more* distinctive songs than their male counterparts. These results suggest that structural and cultural factors—rather than differences in raw ability—are responsible for gender disparities in creative production. To conclude, in Chapter 3 (**What Makes Popular Culture Popular? Product Features and Optimal Differentiation in Music**) I unpack the relationship between musical creativity and commercial success. Using additional data on the week-to-week performance of 27,000 songs appearing on *Billboard's* Hot 100 charts between 1958 and 2016, I find that a song's perceived sonic proximity to its peers influences its position on the charts. Contrary to the claim that all popular music sounds the same, I find that songs sounding too much alike—those that are highly typical—are less likely to succeed, while those exhibiting some degree of optimal differentiation are more likely to rise to the top of the charts. These findings offer a new contingent perspective on popular culture by specifying how product features organize competition and consumption behavior in cultural markets. Considered as a whole, the dissertation contributes to the literatures on cultural markets and creative industries, economic and cultural sociology, organizational theory, and computational social science.

Chapter One

The Social Organization of Creativity: A Multi-Dimensional Perspective[1]

Creativity is a central tenet of cultural production. The work of musicians, painters, writers, chefs, and even academics is concerned in large part with producing something new, exciting, and different than what came before —not only as a means to an end, but as an end in and of itself (Jones, Lorenzen, and Sapsed 2015). This does not mean that all cultural producers are equally interested in or capable of great feats of creativity—most would probably agree that the music of David Bowie is more innovative than the music of Billy Joel, for example—but such variance highlights the role novelty plays in the everyday life of cultural production.

While considerable research has been conducted on the successful diffusion of creative ideas and the link between creativity and performance, both for individuals (e.g., Uzzi et al. 2013) and organizations (e.g., Ahuja 2000), less is known about the production of novelty itself, which serves as a critical first step in this process (Ruef 2002; Seidel and Greve 2016). Where do new ideas come from? Much of the early research on this topic hails from psychology and investigates when and why particular types of actors exhibit more creativity than others, highlighting individual differences in creative ability, skill, and motivation (for a review see Amabile and Pillemer 2012). However, an ever-growing stream of research in psychology, sociology, and related disciplines emphasizes the integral role of collective action and social structure in shaping the creative process (Becker 1982; Bourdieu 1993). Indeed, the propensity to create new things is determined not only by individual characteristics but also by connections

---

[1] This chapter is co-authored with: Eric Quintane (ESMT, Berlin), Noah Askin (INSEAD, Paris), and Joeri Mol (University of Melbourne).

between people and the web of relationships in which they are embedded (Fligstein and Dauter 2007; White 1981, 2002).

This literature has tended to focus on direct interactions within teams (Ruef 2010; Uzzi and Spiro 2005; Wuchty, Jones, and Uzzi 2007) or networks (Burt 2004; Obstfeld 2005; Uzzi et al. 2013), but there are also indirect sources of social influence that drive variation in creative output. Matty Karas, editor and curator for the musical arm of digital media platform REDEF, points out that “Even if you are a one-man band, writing and recording all the parts in your bedroom, you are collaborating with the ideas of your predecessors and your peers, both consciously and subconsciously, and you’re almost certainly nicking random bits of lyrics and melody out of the collective air” (Karas 2016). This interpretation of the creative process has led to a surge in music copyright lawsuits against those who have appropriated elements of other artists’ work without attribution, either intentionally or unintentionally (Runtagh 2016), but it also invites us to reconsider the nature of collaboration itself.

In this paper, we propose a new framework to empirically test how different types of social connection affect the creativity of cultural producers in music. As the quote above suggests, the act of influence need not occur through direct contact or even conscious borrowing. We identify four distinct “spheres of influence” to account for the multitude of ways in which cultural producers are embedded in their social environment. To summarize: the *collaborative sphere* is defined by face-to-face interaction between artists; the *cultural sphere* is a function of shared membership(s) to genre categories reflecting cultural similarity; the *organizational sphere* is defined by shared organizational affiliation(s), and; the *geographic sphere* is a function of physical proximity or co-location. Taken together, this framework provides a multidimensional perspective on social connection and influence, while offering a new perspective on the nature of "musical scenes" (Crossley 2009).

Note that, while the collaborative sphere assumes some kind of direct social contact, the other spheres reflect shared exposure to or affiliation with a particular category, organization, or geography. Nevertheless, we argue that each of these spheres defines an important class of potential influences, shaping their fellow actors' propensity to create new products through channels that are not adequately accounted for in current research. We are not the first to recognize the multiplexity of social life (see Zukin and DiMaggio 1990; Boschma 2005; Goldberg et al. 2016), but no one to our knowledge has generated a typology to understand how different domains of social connection independently shape creative output.

In addition to identifying these spheres of influence, we propose two mechanisms through which influence is exerted to shape producers' creativity. One explanation for how this process unfolds is "exposure to diverse ideas." This perspective describes a compositional process whereby producers encounter some set of material or symbolic resources and combine them in a new way. Exposure to and recombination of diverse ideas serves as an underlying assumption for research on structural holes (Burt 1992), boundary spanning (Fleming and Waguespack 2007), and entrepreneurial bricolage (Baker and Nelson 2005). We then postulate a second mechanism, which we argue is fundamentally distinct from the first: "exposure to creative alters."[2] Unlike the first mechanism, which underscores a resource-based perspective on

---

[2] We use the term "alter" to refer to all actors who are connected to a focal actor in a given sphere of influence, regardless of whether it refers to a direct network connection (i.e., the collaborative sphere) or not (i.e., the cultural, organizational, and geographic spheres).

creativity, exposure to creative alters highlights how even indirect connections to some broader community of creative producers can produce an unusually stimulating environment where the generation of new ideas or novel products becomes accepted, expected, valued, and thereby more likely.

To test how these mechanisms operate across spheres of influence to differentially shape the creation of novel products, we leverage data on the global music industry, a context in which creativity is particularly salient but poorly understood (Negus 1992; Ratliff 2016). The extensive dataset used for this analysis describes over 25,000 unique artists and 600,000 original songs recorded and released between 1955 and 2000. The scope and granularity of this data allow us to: (1) construct a measure of Song Novelty using algorithmically-derived features summarizing the sonic character of songs, and, and (2) calculate dynamic measures of Genre Diversity (i.e., exposure to diverse ideas) and Alter Creativity (i.e., exposure to creative others) across each artist-song's collaboration network (collaborative sphere), genre(s) (cultural sphere), record label (organizational sphere), and city or country (geographic sphere). We then estimate the effect of these measures on Song Novelty. Although accessing diverse ideas through collaboration may enhance the production of novelty, simply being proximate to other creative artists within one's home genre, label, or city can have a significant independent effect on future creative efforts. After reporting our results, we conclude by placing them in the context of other research and discussing their consequences for a more holistic understanding of the social foundations of creativity.

The Social Structure of Creativity

The production of culture perspective in sociology explicitly recognizes the role the social environment plays in shaping symbolic and material culture (Hirsch 1972; Peterson 1990;

Peterson and Anand 2004), and creativity is an important part of this process (Becker 1982). While creativity is defined in some literatures as the generation of something that is both novel and useful or appropriate in a given context (e.g., Amabile 1996), much of the extant research has focused on audiences' *evaluations* of creative goods. In this paper, we focus instead on the production of novel outcomes, which we operationalize as the relative distinctiveness of a product based on its constitutive features. Product novelty or distinctiveness is increasingly recognized as a critical first step in the innovation process, and many works of art, scientific discoveries, and technological inventions derive a large portion of their value from their novelty (Padgett and Powell 2012; Sgourev and Althuizena, 2014; Criscuolo et al., 2016; Johnson and Powell, 2017; Cattani et al., 2017). Novelty also serves as an important signifier of competence and achievement in many fields of cultural production, where "being creative" is highly valued by peers and critics alike (Cho and Mauskapf 2019). Consumers may be less inclined to value novelty for novelty's sake, but moderate levels of atypicality can also help producers differentiate themselves from their competition and provide significant ecological advantages (Askin and Mauskapf 2017; Uzzi et al. 2013).

The production of novelty, and creativity more generally, has antecedents in individual-level cognition, skills, dispositions, and motivations, but psychologists and sociologists alike recognize that these characteristics are moderated by the social environment (Amabile and Mueller 2008; Becker 1982). As early as 1944, Polanyi argued that economic action was embedded in social structure. Granovetter (1985) developed this idea further, showing how behavior was shaped not only by the quantity and quality of one's relationships but also by the configuration of those relations and how they positioned actors for market exchange. Much of the subsequent work in this area has studied the effect of embeddedness on firm performance (e.g., Uzzi 1996, 1997), peer evaluation (e.g., Cattani, Ferriani, and Allison 2014), and knowledge transfer (e.g., Reagans and McEvily 2003; Wang 2015).

This proliferation of research on embeddedness and network science has brought attention to the relational nature of all kinds of behavior (Emirbayer 1997), including creativity and innovation (e.g., Simonton 1984; Burt 2004). These and other studies, however, focus almost exclusively on how individual actors or firms are connected through explicit social or professional networks, where ties are defined by direct interaction. Yet we know that the explanatory power of social connectedness reaches well beyond traditional conceptualizations of networks (cf. Simmel, 1950). In fact, the entire subfield of economic sociology is largely predicated on the idea that interpersonal relationships generate macro-level market structures that influence the dynamics of production and consumption (Bourdieu 1993; Fligstein and Dauter 2007; White 1981, 2002). Moreover, recent work on the multiplexity of networks suggests that people are bound together by many different types of relationships (Boccaletti et al. 2014; Breiger and Puetz 2015; White, Powell, and Owen-Smith 2003). Given the inherent complexity of social life, it is not unreasonable to assume that actors simultaneously hold distinct positions in different networks, which in turn are likely to have independent effects on outcomes such as performance, reputation, and innovation (Heaney 2014). At the level of organizations and markets, Padgett and Powell (2012) argue that the emergence of novelty is, in essence, a function of the coevolution of multiple networks. Nevertheless, our understanding of how different dimensions of social connectedness inform creative production is incomplete.

To address this gap, we identify four distinct “spheres of influence” that we believe structure the production of novelty in cultural markets like music. While direct interpersonal networks continue to play an important role in this process, they are not the only means through which cultural producers are influenced by one another. Table 1.1 summarizes the four spheres: collaborative, cultural, organizational, and geographic. Broadly speaking, these can be conceived of as loosely bounded networks or communities of actors that share some form of connection or affiliation.[3] Different spheres are associated with different norms, values, and practices, and they define the relevant boundaries of social comparison in distinct ways. Artists who simultaneously inhabit the same sphere will experience an increased likelihood of influence via social proximity (Boschma 2005), which both increases the odds that they will cross paths and interact with one another, and provides a latent channel through which shared experiences, norms, resources, and creative energy flow (cf., Podolny 2001).

--Insert Table 1.1 here—

While we argue that these spheres are analytically distinct and at least partially independent from one another, we do not claim that they are exhaustive, nor that they represent the most significant forms of connection across all domains of social life. Our focus on these four spheres in particular is jointly informed by our knowledge of the music industry, as well as previous research on the multifaceted nature of social proximity and embeddedness. For example, Boschma (2005) integrates insights from economic geography to argue that social, cognitive, institutional, organizational, and geographic proximity all play a role in the emergence and diffusion of innovation. Similarly, Zukin and DiMaggio (1990) discuss structural, cognitive,

[3] One could conceptualize each of these forms of connection—collaboration, cultural similarity, organizational affiliation, or geographic proximity—as a binary or weighted variable, whereby an actor's connection could be considered strong or weak. For the purposes of this paper, we treat these connections as binary variables.

cultural, and political embeddedness in their book on structures of capital, although they focus on how these dimensions shape logics of exchange, rather than innovation or creativity. They do note, however, that “the time is ripe…to begin to compare, classify, and develop analytic theories about varieties of informal social structures” (Zukin and DiMaggio 1990: 18). Few scholars have answered their call (for exceptions, see de Vaan et al. 2015; Goldberg et al. 2016; Ruef 2002). Below, we draw on this and other related research to define each sphere of influence and situate them amongst each other.

Collaborative Sphere

The collaborative sphere of influence is defined by interpersonal interaction or direct contact, whereby two or more actors are connected when they explicitly work together to co-produce some creative output. Collaboration occurs, for example, when artists share a song-writing credit or record as members of the same band. This is the kind of connection most social scientists are talking about when they study how an individual’s network impacts their behavior. It constitutes the most well-theorized and empirically documented form of social influence, and it includes work on how brokerage (Burt 2004), closure (Soda, Stea, and Pedersen 2019), small world-iness (Uzzi and Spiro 2005), and other network characteristics shape an actor’s propensity to create something new.

Collaboration is generally considered an essential, or even necessary, condition for some creative tasks. Through collaboration with others, individuals and organizations can access the potentially diverse perspectives, experiences, and resources necessary to generate innovative solutions to complex problems. This is the primary reason teams and small groups outperform individuals across many empirical contexts, including entrepreneurial startups (Ruef 2010), the production of scientific knowledge (Uzzi et al. 2013), and Broadway musicals (Uzzi and Spiro

2005). Yet the effectiveness of collaboration can depend on the nature of the task and the overall structure of one's network. For example, if you collaborate with homogenous others on a simple task, it might lead to groupthink and conformity pressure. And in practice, while occupying a position of brokerage might help with the generation of new ideas, it can also hamper the diffusion of innovation and conferral of legitimacy in certain contexts (Burt 2005; Fleming, Mingo, and Chen 2007). Interpersonal collaboration is an important source of creative influence in music (Crossley 2009), although it can present difficulties in regards to the coordination and integration of diverse ideas.

Cultural Sphere

The cultural sphere is defined by shared category membership(s), whereby two or more actors are connected when they affiliate with the same category. In the context of music, this occurs when audiences tag artists—or gatekeepers tag artists, or artists tag themselves—with the same genre label(s). Categories signal an important form of cultural similarity or kinship, whereby category members share a common identity and cognitive schemas, or ways of thinking about the world (DiMaggio 1997).[4] Producers use categories as a guidepost for acceptable behavior consistent with shared norms and values, while audiences use them to classify producers and their products (Hsu and Hannan 2005) and build community (Lena 2012).

Unlike the collaborative sphere, the cultural sphere is comprised of connections that reflect common category membership, rather than direct contact, and thus does not require

[4] One alternative way to define cultural similarity is through shared language (e.g., Goldberg et al. 2016), but we instead follow Zukin and DiMaggio (1990) and DiMaggio's (1997) emphasis on shared cognitive schema and categories.

interpersonal interaction. Artists connected through this sphere may never work with or meet each other, but we argue that they are likely to share many of the same norms, values, expectations, and practices, which in turn provides a latent channel through which influence can travel. In music, we can interpret shared category membership through the lens of genre (Lena and Peterson 2008; Lena 2012). Genres act as important signifiers of meaning and community in music (Holt 2007). Scholars interested in genre have long debated how boundaries and affiliations should be defined – by common musical features, by communities of listeners with shared tastes, by industry gatekeepers, or by producers themselves—but the significance of genre has never been in question. Even today, "no ordering principle is as fundamental to culture as genre" (Lena 2015: 149).

Organizational Sphere

The organizational sphere of influence is defined by shared organizational affiliation, whereby actors are connected when they are employed or represented by the same parent organization. In music, artists occupy the same organizational sphere when they release a song or album under the same record label. Research suggests certain organizations work hard to design structures and employee managerial practices to breed creativity, while others do the opposite (Amabile 2006). Furthermore, members of the same organization develop strategies for creating and sharing knowledge (Walsh 1995), producing an "organizational advantage" that begets future performance and innovation gains (Nahpiet and Ghoshal 1998). Finally, organizational members develop shared culture (Schein 1990), producing stories and rituals that connect them through symbolic values and material practice (Lounsbury and Glynn 2001).

Like the cultural sphere, co-membership in the organizational sphere does not require direct contact or interaction. Being affiliated with the same organization may bring actors into contact with one another, but this is not necessary for a connection to exist. Simply being a part of the organization makes two artists more likely than average to be exposed to similar expectations, organizational resources, incentives, and constraints, and thus more likely to be influenced by one another. Most musical artists are represented by record labels that produce, record, market, and distribute their recordings. Although this relationship is one of representation rather than employment and is thus relatively weak, labels nevertheless serve as an important source of diversity in music (Dowd 2004). Certain record labels have developed effective strategies to address changes in the technical and resource environment, leading to systematic variation in the creativity (Peterson and Berger 1971) and success (Benner and Waldfogel 2016) of their artists.

Geographic Sphere

The fourth and final sphere of influence is the geographic sphere, in which a connection is defined by the physical proximity or co-location of two or more recording artists. A growing stream of research on the geography of innovation and entrepreneurship has produced considerable evidence that certain cities and regions are hubs for creativity (Florida, Adler, and Mellander 2017), innovation (Funk 2014), and entrepreneurship (Guzman and Stern 2015), due in part to the presence of like-minded actors, community support systems, and the accumulated resources necessary for experimentation (Cooper and Folta 2000; Lippmann and Aldrich 2016; Thornton and Flynn 2003). Moreover, geographic proximity can significantly affect the flow of information and resources that shape innovation, producing systematic variations in the capacity to innovate at local, regional, and national levels (Jian and Yongsheng 2009).

As in the cultural and organizational spheres, connections formed in the geographic sphere constitute latent associations based on shared local resources, audiences, and experiences, regardless of whether the actors in question have ever met or worked together. In the music industry, geographically-defined artist communities generate "scenes" that cultivate identities, audiences, and aesthetic or stylistic tendencies (Crossley 2009). Consider for example the importance of geography to the development of music scenes like the Motown sound in Detroit, country music in Nashville, or reggae in Jamaica.

Mechanisms of Creative Influence

Now that we have described the "where"—by identifying some of the sources of social influence that may be responsible for creative production—we turn our attention to the "how." We propose two distinct mechanisms through which these influences might be expressed and later integrated into the future work of creative producers: exposure to diverse ideas and exposure to creative alters.

Exposure to Diverse Ideas

One common explanation for how new ideas emerge and evolve is exposure to diverse ideas. As previously mentioned, this mechanism lies at the heart of most network-based explanations of innovation, across a wide variety of empirical contexts, levels of analysis, and disciplinary perspectives (Welch 1946; Galunic and Rodan 1998; Sternberg and Sternberg 1999; Carnabuci and Operti 2013; Jones, Lorenzen, and Sapsed 2015). The origins of this argument go back to Schumpeter's (1934) model of entrepreneurship. His model explains both technological innovation and scientific progress as byproducts of "elemental bricolage," or the recombination of existing resources and knowledge to create a new and unexpected whole.

Perhaps the clearest lineage of this perspective is Burt's theory of network brokerage and structural holes, defined as the gap between two or more individuals in a network who have complementary sources of information (Burt 1992). The presence of these gaps is associated with low levels of structural constraint, subsequently generating opportunities that entrepreneurs and other actors can exploit through brokerage. Robust empirical findings support the argument that the bridging of structural holes grants brokers access to diverse ways of thinking and doing, subsequently improving idea generation and reception (Burt 2004), the creation of game-changing products (de Vaan et al. 2015), and firm-level innovation (Zaheer and Bell 2005). This research has inspired several off-shoots and challengers, such as the "structural folds" perspective (Vedres and Stark 2010), but the underlying mechanism is the same. Because of their network position, brokers are more likely to encounter diverse perspectives and non-redundant information, which should aid them in their quest to (re)create through recombination.

This resource-based view of creativity and innovation also serves as one of the foundational assumptions of contemporary research on category spanning. Much of the early work in this area finds that spanning category boundaries confuses audiences, subsequently hurting economic performance (Zuckerman 1999). More recent findings, however, suggest that spanning behavior can also increase access to diverse viewpoints and sources of information, much like brokerage (Hsu 2006). This and other related work (e.g., Goldberg, Hannan, and Kovacs 2016, Kovacs and Hannan 2015) highlight the consequences of category (mis)fit for audience evaluation and consumption, but we can also apply these findings to producers, who engage in consumption throughout the cultural production process. As artists begin to create new work, they search the art worlds around them for inspiration, drawing on existing frameworks and features, and then recombining them to create something new and different (Becker 1982). The assumption underlying research on both network brokerage and category spanning is the same: an actor's position in social structure (or concept space) is a critical predictor of their access to diverse resources, which subsequently can be recombined to produce something new.

Exposure to Creative Alters

In contrast to this resource-based approach, we propose an alternative explanation for how a focal actor's social environment influences her creative output. Rather than recombining diverse ideas accessed through collaboration or some other source, actors might be incited to create something new and different simply by dint of their proximity to other creative producers. We argue that this explanation is distinct from the "exposure to diverse ideas" argument in several important ways. First, it does not assume that influence occurs only through direct contact or interaction with collaborators. Other dimensions of social connectedness that are less explicit may still produce an effective channel through which creative energy can flow. Second, while the recombination of diverse ideas suggests an intentional or even strategic process, succumbing to the influence of creative neighbors may occur unintentionally or even unconsciously. Finally, exposure to creative alters does not require that novel recombination occur for the production of new ideas to follow. Finding oneself surrounded by peers producing novel products is likely to increase one's propensity to do the same.

Put another way, producers might be "made" more creative through exposure to norms and practices promoting creativity, rather than exposure to diverse ideas per se. Social environments that are especially ripe for creativity are more likely to produce a heightened affective state that amplifies creative energy and expectations all around (see Amabile and Mueller 2008; also Collins 2005, Durkheim 1912). In such contexts, conformity pressure begins to fall away, and the experimentation and risk-taking associated with novelty production may become accepted or even encouraged practice, stimulating neighboring producers to "up their game" and try something new. We believe that this explanation represents a fundamentally different model of how social connections shape creativity, whereby novel ideas and processes can become an essential part of the shared experiences, norms, values, and expectations associated with a certain community of creative producers.

Given the novelty (no pun intended) of the theoretical framework developed in the first half of this paper, we have no formal a priori hypotheses about how these two mechanisms will operate across different domains of social connection. Extant research has largely focused on the positive effect of idea diversity or heterogeneity in the collaborative sphere, but in general we expect both idea diversity and exposure to creative alters to have a positive effect on the relative novelty of a focal artist's future production output. Moreover, we predict that the cultural sphere will have an outsized influence on this process, given its unique role in organizing the field of music production.

Data and Methods

We believe that the global industry for recorded music represents an ideal setting to study the social dynamics of novelty production for several reasons. First, creativity has a central role in musical composition and production. Most musicians and critics consider novelty a preferred outcome of the creative process (Cho and Mauskapf 2019), and consumers value moderate levels of atypicality as a means to differentiate songs from one another (Askin and Mauskapf 2017). Second, musical production is an inherently social process involving a range of different actors and interests (Becker 1982; Hirsch 1972; Negus 1992). Third, the music industry as an empirical setting is inherently relevant and interesting. With domestic wholesale and retail revenue surpassing $10 billion in 2014, and more than 50 million subscribers utilizing digital streaming services each month, music represents one of the largest fields of cultural production and consumption today (Friedlander 2014).

To conduct this study, we constructed a unique database describing more than 600,000 songs recorded and released between 1955 and 2000. We chose to focus on this timeframe because of its relevance to the history of modern popular music (Peterson 1990) and, more practically, because of the amount of missing data before 1955. We stop our analysis in 2000 because we believe that year represents a permanent shift in the way music was recorded and distributed. Although Napster and other illegal file-sharing services started to appear during the late 1990s, iTunes was released at the beginning of 2001, drastically reducing the barrier to self-produce and release music. Given this shift, we believe the platform's introduction serves as a good dividing line for our analysis, ensuring the internal consistency and validity of our dependent variable. We believe the resulting sample provides the most complete picture available of commercially recorded music during this period.

Data was collected from two digital sources. We collected data from The Echo Nest, a data science and Music Information Retrieval (MIR) company owned by Spotify. Before advancements in MIR, musicologists or other highly-trained individuals would have to listen to and hand code the sonic attributes of songs if they wanted to compare and analyze them, a labor- and time-intensive activity that necessarily limited the scope of what could be studied (e.g., Cerulo 1988, Dowd 2002). In much the same way that Natural Language Processing (NLP) techniques have enhanced our ability to analyze the content of text at scale (Blei, Ng and Jordan 2003, Bail 2014), MIR combines musicology with machine learning and computer science to extract unique acoustic "fingerprints" from audio files, providing both high- and low-level feature data describing the sonic attributes of tens of millions of songs. This data provides a simplistic but relatively consistent and comparable representation of how music sounds and is experienced by producers and consumers, and captures familiar musicological information as well as a variety of perceptual features (more details provided below).

We also used The Echo Nest to capture genre affiliations, which we use both to define the culture sphere of influence and measure exposure to diverse ideas. First, we collected all of the unique genre "terms" attributed to every artist in The Echo Nest's digital catalog. These attributions were generated using a combination of web-crawling and text parsing algorithms designed to identify the shared labels people use to describe and classify artists.[5] The resulting dataset contains over 900 unique terms, covering common (e.g., "blues", "opera", "jazz") and less common (e.g., "raggacore", "darkstep", "psychobilly") genres. Terms are attributed at the artist level, and artists can be labeled with more than one term. Each artist-term pair includes a weight that represents the relative "strength" of the affiliation. For the measures that follow, we include all artist-term attributions with weights of > 0.5.[6]

---

[5] See http://blog.echonest.com/post/73516217273/the-future-of-music-genres-is-here for a more detailed explanation.

[6] We chose a threshold of 0.5 for two reasons. First, we wanted to account for the fact that many artists and their music are affiliated with more than one genre. So, rather than select an artist's primary genre (weight = 1), we collected all available artist-term attributions. Second, after reviewing the data, we realized that, while all artists are tagged with at least one genre, and most are tagged with no more than three or four, there is a long right tail of (mostly highly visible) artists affiliated with many genres. We did not want over-attribution among popular artists to bias our results, nor did we want to include weak attributions that were made by only one or a few audience members. Thus, we only included terms that appeared often enough alongside a given artist that

Musicbrainz, a curated and crowd-sourced online music database, provided us with detailed artist credit information, including band members and other individuals involved in the writing, production, and/or performance of all original recordings in the dataset. We use this data to construct the collaborative sphere of influence, and collected additional metadata to control for other artist-level characteristics, such as gender and home city and country (which we used to construct the geographic sphere of influence). Finally, we collected the day-level release dates and record labels associated with every song to model changes in our dependent and independent variables over time and construct the organizational sphere of influence, respectively.

Dependent Variable

Song Novelty. We measure Song Novelty using a feature-based measure of musical distinctiveness for each song in our dataset. In music, preferences are linked to a series of high-level features that structure musical space, such as speed, repetition, sadness, and loudness (Greenberg et al. 2016). While previous research has recognized that creative production occurs through the recombination of both category labels and material features, most empirical studies have used only labels to describe and compare products (for exceptions, see Cerulo 1988, Askin and Mauskapf 2017).

All songs in our dataset were algorithmically-assigned a value for ten sonic features, each of which was designed by The Echo Nest to describe a song's most important characteristics. While these features necessarily distill the complexity of music into a handful of discrete summary statistics and thus fail to capture what makes music "art", recent scientific research

---

they warranted a weight above 0.5. Full descriptive statistics for this variable are included in Table 1.3.

suggests these features represent a universal musical grammar that is shared across human history and cultures (Mehr et al. 2019). They are also well suited for comparing songs to each other (Friberg et al. 2014; Mauch et al. 2015) and are used explicitly for that purpose by Spotify and other streaming recommendation services. Features include several standard musical attributes (e.g., "tempo," "mode," "key," "time signature"), as well as a series of algorithmically-derived measures that represent particular aural or emotive dimensions of music ("valence," "danceability," "acousticness," "energy," "liveness," "speechiness"; see Askin and Mauskapf 2017 for more details regarding these features).

After normalizing all ten features and collapsing them into a single vector for each song, we calculate the cosine similarity between all relevant song-pairs. To define the relevant comparison set for a focal song, we include all other songs released within a rolling 10-year window before the focal song's release date.[7] Thus, a song released on January 1st, 1990 is compared to all songs released between January 1st, 1980 and December 31st, 1989. Each comparison generates a pairwise similarity measure across our 10-dimensional feature space. For each song, we then average all pairwise similarities, invert the result (to represent distance), and multiply this value by 100, producing a single novelty score between 0 (maximum conformity)

---

[7] Note that this comparison set includes both between- and within-genre comparisons for each song, rather than simply within-genre comparisons. The latter specification may be appropriate for a consumer-centric measure of novelty, but we believe that a comprehensive comparison set is more appropriate for a producer-centric measure, as many artists span multiple genres. We also conducted analyses using 5-year, 2-year, and no time windows. Results are robust across each of these specifications; see Table A1.3.

and 100 (maximum distinctiveness). The mean novelty score for songs in our dataset is 20.5, highlighting the fact that most songs sound relatively similar to each other (See Table 1.3 for descriptive statistics).

---Insert Figure 1.1 here---

Independent Variables

To try and adequately account for the independent effects of exposure to diverse ideas and creative alters, we took great care in constructing measures that we believe are as close as possible to the phenomena we are trying to capture. This required us to (1) define the alter set for each sphere of influence at the time of every song release, (2) construct dynamic measures of Genre Diversity and (3) Alter Creativity (again across each sphere at the time of every song release), and (4) investigate and address some issues regarding zero values in the resulting measures to ensure that we did not bias the results of our analysis. We cover each of these steps in detail below, and Table 1.2 summarizes how we operationalize each of these variables.

---Insert Table 1.2 here---

# of Alters. A fundamental aspect of our theory is defining (at least partially) distinct sets of alters for each sphere of influence. Note that each of these values is recalculated at the time of each focal song's release, resulting in a rolling count of the number of alters in a given sphere at the artist-song release level. In addition to defining the relevant referent set for our Genre Diversity and Alter Creativity measures, we also include each of these counts as controls in our models. In the collaborative sphere, other artists are included in a focal artist's alter set if they are still active (i.e., have not released their last song) and share at least one track credit in the

decade before the focal song's release.[8] This includes songs previously released by an artist's bandmates, as well as more transient collaborators involved in the creative production process.[9] *(Mean # of Collaborative Alters: 4)* In the cultural sphere, other artists are included in the alter set if they remain active and have at least one genre affiliation in common with the focal artist (with a weight > 0.5). *(Mean # of Cultural Alters: 1771).* In the organizational sphere, other artists are included in the alter set if they remain active and have released at least one song with the record label associated with the focal song release in the last ten years *(Mean # of Organizational Alters: 59).* Finally, in the geographic sphere, other artists are included in the alter set if they remain active and are co-located with the focal artist at the time of a given song's release *(Mean # of Geographic Alters: 123).* We use city-level locations whenever possible; country otherwise.

Genre Diversity. Rather than use a structural proxy such as constraint to presume access to diverse ideas (cf., Burt 1992, 2004), we leverage the granularity of our dataset to construct a measure of idea heterogeneity that is substantively grounded in the context being studied.

---

[8] We included a rolling 10-year window to maintain consistency across our dependent and independent variables. Results are robust across different rolling window specifications, and are available upon request.

[9] Where possible, we try to isolate the original "cultural producer(s)" for each recording, identifying the individual(s) or group credited with the actual creation of a focal song. Many times, this coincides with the performer(s) credited on a recording, but not always. When the data does not allow us to adjudicate between these actors and their role in the cultural production process, we impute the performer(s) as cultural producer(s).

Specifically, we measure idea diversity by creating a scaled count of the number of unique genres a focal artist is exposed to across each of the alter sets defined above. Genres represent a relatively coherent (if contested) set of stylistic resources, and thus serve as a reasonable proxy for "ideas." An artist who is exposed to 5 unique genres through their collaboration network presumably has a greater opportunity to combine diverse ideas than an artist exposed to 2 genres in the same sphere of influence. For each focal artist, we count the number of unique genres she is exposed to via her set of collaborative, organizational, and geographic alters, respectively. For the cultural sphere, we simply count the number of unique genres with which a focal artist is herself affiliated; this is a measure of genre spanning.[10]

After constructing each of these counts, we then scale them by their relative (dis)similarity. The intuition here is simple: an artist exposed to two disparate genres, such as "heavy metal" and "country," potentially has more diverse material to recombine than an artist exposed to two similar genres, such as "heavy metal" and "thrash metal." To implement this scaling procedure, we first calculate the Jaccard distance between each unique genre pair in our dataset, where the distance is equal to 1 minus the likelihood of two genres co-occurring for a single artist in our dataset. Second, we calculate the average Jaccard distance across all relevant genre pairs associated with a focal artist's raw Genre Diversity score in a given sphere. Finally, we multiply this (0-1) value by the total number of unique pairwise combinations being

[10] We also constructed an alternative measure of cultural recombination that accounted for the number of unique genres a focal artist is exposed to herself, plus the other second-order genre affiliations of the alters in question. Results are robust using this specification, and are available upon request.

averaged. This allows us to simultaneously consider both the number of unique genres an artist is exposed to, as well as the relative similarity (or difference) between each of those genres.[11]

Alter Creativity. To measure a focal artist's exposure to creative alters, we calculate the average novelty score of the songs released by artists who appear in the appropriate alter set(s) defined above. Put another way, we re-appropriate our dependent variable to calculate the mean level of novelty surrounding a focal artist. We thought that this would serve as a defensible proxy for the relative presence (or absence) of creative energy, norms, and values associated with given sphere of influence at a certain point in time. The appropriate referent group is once again determined using a rolling window, whereby all songs released by active alters in the ten years before a focal song's release are included.[12]

Imputation Strategy

---

[11] Given their construction, it is unsurprising that this measure is highly correlated with the # of Alters measure summarized above. We include both of these measures in our models because we believe the size of one's referent group may have an effect on song novelty that is partially independent from "access to diverse ideas."

[12] Note that this is an "average of averages." We create a creativity score for each alter, calculated as the average of the novelty of all the songs released in the last 10 years for that artist. We then average the creativity scores of all relevant alters for a given focal artist. We constructed alternative versions of this measure by (1) averaging each alter's maximum Song Novelty score (rather than averaging each alter's mean Song Novelty score), and (2) changing the length of the rolling window. Results are robust across both specifications, and are displayed in Table A1.1 (final model).

After constructing each of these measures, we explored them descriptively and found a surprising number of zero values. Figures 1.2a-d display simple distributions of the "Number of Alters" measure for each sphere of influence. We were concerned that these zeroes may reflect a missing data issue, especially in the collaborative sphere. We realized that, while a few of the zero values might reflect incomplete data, most of them were legitimate, or reflected a logical consequence of our measure construction strategy. For example, if a focal artist early in her career had yet to collaborate with anyone she would have no collaborative alters, which would in turn leave her with Genre Diversity and Alter Creativity scores of 0. Similarly, if an artist was only exposed to one unique genre in a given sphere, her Genre Diversity score would appear as 0, as the simple count (1) is scaled by the average pairwise genre distance (which in this case would be 0, as you cannot calculate the distance between an object and itself!).

---Insert Figures 1.2a-d here---

To resolve our concern that this large number of zeroes might bias our findings, we developed the following imputation strategy. For Genre Diversity values of 0 where the number of unique genres = 1, we replaced them with the lowest non-zero value for Genre Diversity in our dataset. Note that this is different than being exposed to no genres at all (Genre Diversity = 0, unscaled unique genre count = 0). In these situations, we left the Genre Diversity scores as is. For Alter Creativity, we were concerned that all zero values would artificially drag down the effect of the variable. Thus, all instances where Alter Creativity = 0 were replaced with the mean value for this variable in a given year. Figures 1.3a-d and 1.4a-d display the impact of these imputations on the distributions of these measures. In the analyses that follow, all models employ the imputed version of these measures unless otherwise noted. To make sure this was not producing artificial results, we also ran versions of our full model with the raw data and without any observations that had at least one zero or missing value. The overall pattern of results is largely consistent across these specifications (see Table A1.1, Figure A1.1).

---Insert Figures 1.3a-d here---

---Insert Figures 1.4a-d here---

Control Variables

In addition to our eight explanatory variables of interest (i.e., Genre Diversity and Alter Creativity across each of the four spheres of influence), we also include a suite of song- and artist-level controls to account for potential confounds that may influence the production of novelty in music.

Country. During the timeframe covered in our analysis, 43% of the music industry's commercial recording activity took place in the United States. To control for the outsized influence that the US had on the means of production, we include a dummy variable set to 1 for any song released by a US-based focal artist.

Female or Male. We also wanted to account for artist gender. Approximately 38% of the songs in our dataset were created by musical groups or bands (these serve as the reference group), while 62% of the songs were created by solo focal artists. Of that number, men outnumber women by more than a 4-to-1 ratio.

Classical or Jazz. The vast majority (84%) of songs in our dataset are part of the popular music tradition, but some fall into the category of classical music and/or jazz. Due to the complex form, unique instrumentation, and other compositional characteristics associated with works in these genres, we worried that observations labeled with any variation of the term "jazz" or "classical" were likely to be systematically more novel. In addition to running models that exclude all songs affiliated with these two genres, our main models include a dummy variable for each.

Popular Success (Charting Artist). Popular success and widespread appeal are likely to shape an artist's propensity to produce novel songs as well. To measure a focal artist's popular success at the time of a given song's release, we include a dummy variable set to 1 for the 21% of artists who had previously placed one or more songs on the *Billboard* Hot 100 Charts, 0 otherwise.

Major Label. In addition to the organizational sphere of influence, we wanted to distinguish between major versus independent record label representation to account for differences in the quality of organizational support received by artists. Thus, we include a dummy variable set to 1 for the 20% of songs in our dataset released on a major label, 0 for independent label releases.

Artist Tenure. An artist's experience may also affect the relative novelty of their musical output. To operationalize experience, we construct a measure of artist tenure. This is calculated as the length of time (in days) since a focal artist's first release in our dataset. We also include a squared term for artist tenure in our main models, following previous research that has found a curvilinear effect of tenure on performance (Cole 1979, Jones and Weinberg 2011).

Artist Productivity. Part of the experimentation process is generating lots of outputs, knowing that there is always a possibility of producing something new and exceptional. Thus, the number of songs produced by a focal artist may covary with Song Novelty. We control for the total number of songs that an artist has released since the beginning of her career, standardized by the number of years she has been active in our dataset.[13]

Past Creativity. We know that there are large individual differences in creativity due to differences in motivation, skill, and cognitive style (Amabile and Pillemer 2012). To account for these individual differences, our main models control for a focal artist's past creativity, which is simply calculated by taking the average novelty score for all songs released previously by that artist.[14] The relative effect size for this variable will serve as a helpful benchmark for the rest of the results in our analysis.

Decade. To account for large-scale historical shifts, changes in production and distribution technology, and the emergence and evolution of exogenous fads and fashions, we include dummy variables for the decade in which a song was released.[15] The period from 1955–1960 serves as the reference group. Figure 1.5 plots the general historical trend of musical

---

[13] We also ran an alternative specification where productivity was measured as the number of songs released by an artist in the previous two years only. Results were robust across both specifications.

[14] We also ran specifications (1) where past creativity was measured as the average of an artist's songs released in the previous two years only, and (2) using artist-level fixed effects rather than Past Creativity as a control variable (see Table A1.2). Results were robust across all specifications.

[15] We also ran an alternative specification using year-level fixed effects rather than decade dummies (see Table A1.2). Results were robust using this method.

creativity over time, suggesting that the 1960s were particularly promising for the production of novelty, while the 1980s were not.

---Insert Figure 1.5 here---

We do not include artist- or year-level fixed effects in our main models, but results using these specifications are largely consistent with our primary models and are reproduced in Table A1.2. Table 1.3 includes descriptive statistics and simple correlations for each of these variables.

---Insert Table 1.3 here---

Modeling Strategy

The analysis that follows is divided into two sections. We begin by running pooled, cross-sectional OLS regressions to estimate the effect of each of our primary explanatory variables (Table 1.4, Figure 1.6). To assess whether there might be any tradeoffs between these measures and their relative effect on novelty production, we then ran a series of OLS models (by sphere) with interaction terms between Genre Diversity and Alter Creativity (Table 1.5). We also generated a series of interaction plots to aid our interpretation of these results, (Figure 1.7a-d). In an informal appendix, we present several robustness checks, re-estimating our main model: with non-imputed measures, without zeros, for artists' first songs only, and with an alternative measure of Alter Creativity (Table A1.1, Figure A1.1); with year- and artist-level fixed effects (Table A1.2); and with alternative specifications for our song novelty measure (Table A1.3). These tables appear at the end of the dissertation.

Results

Our primary interest concerns the multi-dimensional effects of social influence on creative production, but we begin by summarizing some of the other predictors of novelty included in our main models (see Table 1.4). The results presented in Model 1 shed light on some of the individual-level and historical factors that shape creative production. Not surprisingly, we find that an artist's propensity to create novel songs in the past is a consistently strong predictor of future creativity. This result signals the importance of individual differences among cultural producers and supports the notion that some artists are simply more skilled at or motivated to create novel products. We also find evidence for a u-shaped relationship between artist tenure and the production of novelty. Young artists may be less constrained by the norms and expectations surrounding certain fields of cultural production, while those in their prime are more likely to be burdened by audience expectations and conformity pressure. Artists who continue to work into the twilight of their careers, however, accrue sufficient social, political, and artistic capital to experiment without incurring the penalty assigned to producers in the middle of their careers. Other signals of success and status—such as previous appearances on the *Billboard* Hot 100 Charts, or representation by a major record label—may lead to other positive career outcomes, but they tend to promote conformity rather than creativity.

---Insert Table 1.4 here---

In Models 3–5, we estimate the independent effects of Genre Diversity and Alter Creativity across each sphere of influence, combining them with all of our controls in Model 6. As expected, all eight measures seem to be positively associated with Song Novelty, although Genre Diversity in the organizational sphere and Alter Creativity in the geographic sphere lose statistical significance in the fully-specified model. In Figure 1.6, we standardize the effect sizes from this model to compare the substantive significance of these effects. Here it becomes clear that Alter Creativity in the cultural sphere—and to a lesser extent in the organizational sphere—has by far the largest impact on a focal producer's future creativity. This result is robust across a variety of specifications and measurement strategies (see robustness checks in Table A1.1 and Figure A1.1).

---Insert Figure 1.6 here---

The substantive interpretation of this pattern of findings is profound. First, it appears that exposure to creative alters is an important driver of novelty production, independent of exposure to genre diversity. Simply being surrounded by other creative producers increases the distinctiveness of your own creative output. Second, the cultural sphere, defined here via shared genre affiliation, clearly comes out on top as the dominant organizing principle of creative production in music. This is reflected by the fact that the variance explained by Model 3 is much greater than the variance explained by the models including the other spheres of influence (Models 2, 4, and 5). Although we find some support for the positive effects of Genre Diversity accessed through direct collaboration and other spheres of influence, it appears that occupying one or more genres populated by other artists producing novel work—think progressive rock in the late 1906s, or grunge in the early 1990s—is the best predictor of Song Novelty in our dataset.

Interactions

To develop a better understanding of how Genre Diversity and Alter Creativity operate together, we ran an additional suite of models that interacted these variables within each sphere of influence.[16] The results summarized in Table 1.5 suggest that there is in fact a tradeoff between these two variables. While the coefficients for each of these interactions is negative and statistically significant, however, their substantive significance is less clear. Based on the plots

---

[16] We also ran interactions between spheres but they were not significant.

represented by Figures 1.7a-d, it appears as though the collaborative sphere may be the only context in which these two mechanisms are truly at odds with each another.

---Insert Table 1.5 here---

---Insert Figures 1.7a-d here---

Discussion

The findings presented in this paper provide compelling evidence that creativity—and the generation of distinctive cultural products in particular—is shaped by multiple dimensions of social connection or proximity, but not in the way previous theory suggests. Existing research has focused primarily on how direct ties defined by interpersonal interaction shape new ideas and products. In this paper, we posit that there are other channels through which actors influence one another—including shared category membership, organizational affiliation, and geographic co-location. These other spheres of influence do not require direct contact or interaction to exert an effect on creative production. For example, the cultural sphere operates by invoking a shared commitment to norms, values, and assumptions that transcend direct social interactions or collaborations. Artists who are members of the same genre category may be more likely to cross paths and collaborate; but even if they do not, they still share common ground and compete for audience attention. The same can be said for the organizational or geographic sphere, which describes how actors immersed in the same local or regional community are likely to share resources, experiences, and audiences, whether or not they come into direct contact with each other. Through our analysis of these distinct spheres of influence, we find that exposure to creative alters in one's home genre(s), rather than the recombination of diverse ideas accessed through collaboration, is a primary driver of novelty production in music.

These results provide a more nuanced and holistic perspective on how individual characteristics, networks, and other forms of social connection organize the creative process. Nevertheless, our work has some important limitations that merit consideration. First, we assume that our dependent variable is a valid proxy for the "actual" novelty or distinctiveness of a song. Rather than aggregate subjective evaluations made by consumers, we chose to use the sonic features of songs to construct a materially-informed and hopefully more objective measure of product distinctiveness. This measure has its limitations: for example, we do not account for song lyrics as a potential source of differentiation (see Berger and Packard 2018), although such a comparison would reduce the breadth of our analysis and further complicate calculations of song novelty. We try to take full advantage of the quantity and quality of music-related data available today to construct a scalable measure of product novelty that is both internally and externally valid.

Our measures of exposure to diverse ideas (Genre Diversity) and creative alters (Alter Creativity) are also limited in their ability to explain exactly how these mechanisms operate. For example, does being surrounded by creative alters motivate artists to be more creative themselves? And if so, does this unfold through a process of social learning (Bandura and Walters 1977), or by following a mimetic norm for atypicality (Amabile 1996, DiMaggio and Powell 1983), or through heightened levels of affective energy that make novelty production more likely (Amabile and Mueller 2008, Collins 2005, Durkheim 1912)? The data used in this paper does not allow us to disentangle these or other possible explanations. We hope future research will continue to develop and refine these constructs theoretically and deploy them empirically to better understand the effect of social influence on creativity. Finally, the four spheres of influence and two mechanisms described in this paper do not represent an exhaustive review of the many ways in which social context affects creative output. It may not be surprising that genre plays such an important role in the production of new music, but exactly which dimensions of social connectivity matter most, and in what ways, is likely domain dependent. Future research should continue to tease apart the relationships between these dimensions and test how they operate across different empirical contexts.

Conclusions

We believe our paper makes several important contributions. First, by conceiving of social structure as a multi-dimensional construct, we recognize the fact that information and influence flow not only through collaboration but also other types of connection that are defined by social proximity rather than direct contact or interaction. Second, we unpack two mechanisms that highlight two very different ways in which these spheres of influence shape creative production. Finally, the data and methods used in this study present a host of exciting opportunities for empirical research on product novelty and cultural production more generally. Creativity research has historically relied on subjective ratings to capture the distinctiveness of a particular product or idea. We instead leverage advances in the field of music information retrieval (MIR) to generate a more objective measure of product novelty using songs' sonic features. We hope that our study, along with others capitalizing on algorithmically-generated, large-scale feature data (Bail 2014, Lazer and Radford 2017), will encourage scholars to continue exploring the nature of creativity in new and different ways.

Chapter Two

Do Women Produce More Novel Cultural Products than Men?

The Gendered Effects of Networks and Genres on Musical Creativity[1]

Women have exhibited higher rates of participation and consumption across fields of cultural production for decades (Lizardo 2006, Christin 2012, Schmutz, Stearns and Glennie 2016). In the United States, for example, women earn 62 percent of degrees in the fine and performing arts (U.S. Department of Education 2017), a trend that seems to reinforce the widely-held belief that artistic and creative activities are intrinsically "feminine" (Miller 2016, Lagaert, Van Houtte and Roose 2017). Yet women remain disadvantaged in creative careers. When compared to their male peers, women are less likely to reach the top of their profession (Miller 2016), are paid less (Lindemann, Rush and Tepper 2016, Brown 2019), and enjoy muted recognition from media, critics, and peers (Schmutz and Faupel 2010, Lincoln et al. 2012, Stokes 2015). Why do these inequalities persist?

Scholars interested in understanding this puzzle have historically focused on two explanations. First, grounded in the notion that men are naturally more creative than women, scholarship examined the extent to which there were gender differences in creative ability. However, fifty years of experimental research suggests that there is little, if any, evidence to support such a claim (for reviews see Kogan 1974, Baer and Kaufman 2008). Nevertheless, when asked why the 2018 Grammy Award recipients were almost exclusively male, the award show's president said that "It has to begin with…women who have the creativity in their hearts

---

[1] This chapter is co-authored with: Noah Askin (INSEAD, Paris), Sharon Koppman (University of California, Irvine), and Brian Uzzi (Northwestern University).

and souls, who want to be musicians, who want to be engineers, producers, and want to be part of the industry on the executive level. [They] need to step up," suggesting that women's poor showing was somehow a function of their own incompetence or lack of motivation (Yahr 2018).[2] More recently, scholars have focused on the presence of gender bias in evaluations of creative products such as songs, novels, and designs (e.g., Goldin and Rouse 2000, Schmutz and Faupel 2010, Proudfoot, Kay and Koval 2015, Stokes 2015). Peers, critics, and mass audiences make unequal attributions regarding the quality of creative products produced by male versus female creators.

In this paper, we focus our attention on a third and largely unexplored way in which gender may shape these outcomes: structural and cultural differences in the work contexts of creative producers. Creating something new and different is not simply a function of an individual's ability and motivation; it is shaped by social context (Childress 2017, Becker [1982] 2008). Research on career advancement suggests that men and women tend to work under different structural conditions: for example, women benefit differently from informal social connections (Brass 1985, Ibarra 1992) and cluster into different roles and categories of work than men (Gorman 2006, Roth 2006, Cech 2013). These differences presumably shape the content of creative products generated by men and women, yet we do not have robust empirical evidence to explain when this occurs and to what effect. Unpacking these dynamics is important for understanding gender inequality in creative careers. If men are producing more creative work

[2] In 2019, 50% of the major Grammys were awarded to women—a dramatic departure from past years—but female representation in the industry at large remains substantially below population averages (Smith et al. 2019).

than women, this could explain greater levels of achievement and recognition; conversely, if women are producing more creative work than men, this would suggest that gender bias in the reception and evaluation of creative products is even more pronounced than previously assumed. Put another way, if the social contexts that shape creative production are different for men and women, or impact men and women in systematically different ways, this could in turn be associated with their creative output and explain some of the gender inequality we observe.

We set out to rigorously examine the extent to which there are gendered differences in the novelty of creative products in the music industry. Specifically, we ask: Are there are material differences in the creative output of female versus male artists? And if so, what structural and cultural factors are linked to these differences? To answer these questions, we analyze an extensive dataset comprising over 250,000 songs produced and released by male and female solo artists between 1955 and 2000. Using an algorithmically-derived, feature-based measure of musical novelty, we initially find no mean difference in the creative production of women versus men. After controlling for contextual factors associated with artists' creativity, however, we find that female artists produce songs that are significantly *more* novel than those produced by their male peers. We also find that these contextual factors can operate differently for men and women. Female artists benefit creatively when they have larger collaboration networks, but male artists do not. Moreover, while previous work suggests that the tendency to cluster in female-majority networks and genres can hinder women's career prospects, its association with women's creative output appears relatively weak. In contrast, men experience a larger penalty for affiliations with "female" genres but benefit creatively from collaborating with women. After summarizing and interpreting these results, we conclude by discussing implications for our understanding of how gender and creativity cooperate in the musical marketplace and beyond.

Gender and the Novelty of Creative Products

Gender inequalities exist well beyond the context of creative production—they pervade nearly every corner of work life. Disparities between men and women persist across a number of phenomena, including wages (Blau and Kahn 2000, England 2010), hiring and promotion into upper management (Cohen, Broschak and Haveman 1998), patenting (Whittington and Smith-Doerr 2008), entrepreneurial activity (Thébaud 2015), and access to social capital (Ibarra 1992, Ibarra 1997, Yang, Chawla and Uzzi 2019), awards recognition (Ma et al. 2019), and financing for new ventures (Kanze et al. 2017, Oliveira et al. 2019). Many of these disparities are pronounced and enduring, buoyed by a complex machinery of structural and cultural forces that have proven difficult to change.

The scope of gender inequality in fields of cultural production is particularly puzzling, given the rates of women participating in activities such as art, music, and literature. As highlighted in the introduction, extant research has predominately highlighted two explanations of gender inequality in creative production. The first attributes the paucity of eminent women in the creative professions to man's "naturally greater" creative ability. Darwin himself suggested that "man is more courageous, pugnacious and energetic than woman, and has a more inventive genius" (1872: 557). Decades of psychological research, however, has produced considerable evidence that there are no consistent gender differences in creative potential (for reviews, see Kogan 1974, Baer and Kaufman 2008). The vast majority of this research relies on tests of divergent thinking—e.g., asking subjects to list different uses for common objects like paper clips or bricks—and finds, if anything, that women are slightly more creative than men (Baer and

Kaufman 2008). Studies that employ Amabile's (1982) consensual assessment tool, in which subjects create an object like a poem or story that is later evaluated by experts, are similarly devoid of significant gender differences (Amabile 1982, Kaufman, Baer and Gentile 2004). Nevertheless, among both the general public and many creative producers, the belief that men are naturally more creative than women persists (Proudfoot, Kay and Koval 2015).

Other scholars, primarily but not exclusively sociologists, have tended to focus their efforts on explaining how creative producers and their products are categorized and valued by audiences. Stereotypes about a producer's gender extend to their products and the product markets within which they operate, shaping consumer evaluations (Tak, Correll and Soule 2019). Creative products made by women tend to be assigned less value, assessed less positively, and garner less attention and recognition (Bielby and Bielby 1992, Lindemann, Rush and Tepper 2016). Attributions of creativity and quality by both experts (Goldin and Rouse 2000) and lay audiences (Proudfoot, Kay and Koval 2015) are biased toward men, while female designers and musicians win fewer awards (Schmutz and Faupel 2010, Miller 2014, Stokes 2015). Even when women are recognized for their creative achievements, they accrue less prestige and respect than their male counterparts (Ma et al. 2019). In summary, these studies focus on the evaluation of creativity: female artists and their creative products are devalued by audiences, critics, and other gatekeepers, who remain biased by gendered stereotypes about what it means to be creative.[3]

[3] A third but related stream of research has found that women are more likely to be underrepresented in fields in which genius is believed to be necessary for success, because genius is stereotypically associated with men (Leslie et al. 2015, Meyer, Cimpian and Leslie 2015). This stereotype may shape women's underrepresentation among creative producers by

Taken together, research on gender inequality in creative careers can more or less be divided along disciplinary boundaries. Studies in psychology tend to focus on the antecedents of creativity, specifically novelty, but ignore the extent to which creative production is embedded in a complex web of activity among artists, critics, and audiences (Becker [1982] 2008). Studies in sociology tend to focus on the external valuation of creative work, but neglect the material content of creative products themselves—a choice usually justified by their inherently subjective and difficult-to-measure character (Childress 2017). This omission is regrettable, particularly given the potential for a sociological approach to study the production of novel work. Although some scholars have examined how creative producers evaluate novelty (e.g., Rosenblum (1978) on photographs; Guetzkow, Lamont and Mallard (2004) on fellowship proposals; and Koppman (2014) on advertising materials), these studies analyze perceptions of novelty, i.e., what novelty means to different groups, rather than material differences in the novelty of creative products themselves. This leaves questions about material differences in the novelty of creative products made by men and women—and the contextual factors underlying these differences—unexplored.

Most creative professionals—including musicians, painters, scientists, chefs, writers, and architects—are engaged in the production of cultural goods (for a review see Peterson and Anand 2004). Like all forms of production, creating cultural goods involves transforming inputs into outputs (Perrow 1967). Here, inputs are material and symbolic elements such as sonic attributes, aesthetic conventions, and substantive knowledge, while outputs are particular arrangements of these elements in the form of songs, paintings, or research articles. These outputs—"creative

---

coloring the decisions of artists to enter and remain in the profession, as well as gatekeepers' decisions to allow or prohibit market entry (e.g., Gorman 2005, Cech et al. 2011).

products"—may be more or less similar to (or different from) previous arrangements. We refer to this as "product novelty." Creativity is often defined as the generation of products, ideas, or other outcomes that are both novel and valuable (Amabile 1996), but many works of art, scientific discoveries, and technological inventions derive at least part of their value from their novelty. Although we recognize the importance of audiences' evaluations of creative products, in this paper we focus on a product's relative novelty or distinctiveness, which we feel is at the very heart of creative production.[4] By conceptualizing products and their constitutive properties as creative outputs shaped by contextual factors, we can advance knowledge on the role of gender in creative production.

The Social Context of Creative Work

The literatures described above attempt to understand gender inequality in creative careers based on competing arguments about the extent and direction of gender differences in the novelty of creative products. Extant work has approached these questions with either *i*) a (largely unsupported) assumption that men's creative products are more novel than women's, or *ii*) a

---

[4] The term "creativity" refers to a complex and dynamic process that unfolds both in peoples' minds and through social interaction (i.e., it is socially constructed), but for the purposes of this study we focus our attention on differences in creative output—specifically operationalized as product novelty or distinctiveness. This means that we do not capture attributions of social recognition (Barron and Harrington 1981), appropriateness (Hennessey and Amabile 2010), or recognizability (Seidel and Greve 2017)—each necessary for its respective definition of creativity or innovation—which are subject to the types of audience-side social influences that we are attempting to circumvent in this paper.

(largely untested) assumption that there are no material differences in the creative products of men and women. This leaves a third (unexplored) argument: female artists' creative products may actually be *more* novel when compared to men in similar structural and cultural positions. To advance this argument, we integrate research on creativity with the literature on how gender shapes women's experiences at work.

Although creativity is sometimes portrayed in the popular imagination as an individual stroke of insight or genius, we know that it is embedded in a complex web of social experiences and connections that include gatekeepers, peers, audiences, and collaborators (Becker 1984, Childress 2017). And Kanter (1977) showed us that women who are underrepresented in their occupations—what we refer to as female minorities—are treated differently than men by their peers, clients, and managers. We suspect that female minorities' creative products may be more novel than those created by structurally equivalent men, in part because women tend to be held to a higher standard—by both external evaluators and themselves. Kanter was the first to document that female minority saleswomen had to "work twice as hard to prove their competence" (Kanter 1977:973). Although her study focused on workgroup dynamics, the tendency for women's performance to be discounted reflects a much broader phenomenon inside and outside organizations (Castilla 2011). For the same levels of performance, women tend to receive more negative evaluations than men (for a review see Heilman 2012), and they have to outperform men to receive comparable evaluations (Foschi 2000, Botelho and Abraham 2017). To overcome this "double standard," female minorities work harder and seek to demonstrate their performance through more "objective" metrics, such as qualifications, knowledge, and track record (Davies-Netzley 1998, Ragins, Townsend and Mattis 1998). For instance, female consultants advance by establishing technical competence over time through over-preparation

(Ibarra and Petriglieri 2017), congresswomen outperform congressmen in delivering federal funding to their districts and sponsoring legislation (Anzia and Berry 2011) and, holding quality constant, papers by female economists tend to be better written than those by men (Hengel 2017).

We argue that, in much the same way women's performance is discounted by managers, female creators are discounted by critics, audiences, and gatekeepers, which in turn shapes the material content of their creative products. We already know that implicit biases and stereotypes seep into evaluations of creative potential (Koppman 2016) and products (Childress and Nault 2019); thus, women may need to produce more novel creative products than men to secure equivalent industry and audience support. And this is exacerbated by women's stark minority status in some creative professions. At the 2019 Grammy Awards, singer H.E.R. declared, "I had to work twice as hard. I had to earn my respect as a musician growing up as a little girl because you don't expect a little black girl to pick up the electric guitar" (Fekadu 2018).

The Gendered Effects of Networks and Genre Affiliation

We propose that, all else equal, female minorities may produce more novel creative products than men. Yet all else is typically not equal: men and women often face different structural and cultural conditions at work. To better understand gender inequalities in creative production, we need to not only study variation in the novelty of creative products but also examine *when* gender disparities emerge. In particular, we draw from research on key differences in the work contexts of men and women in relation to career advancement (Brass

1985, Ibarra 1992, Bielby and Bielby 1996, Burt 1998, Gorman 2006, Cech 2013) to identify three key variables that may differentially affect the creative work of male versus female artists.[5]

Collaboration Network Size

Women can become socially isolated and disadvantaged in certain professional contexts relative to men because of their limited ability to capture value from large networks (Ibarra 1992, Burt 1998, Yang, Chawla and Uzzi 2019). For example, while networks rich in structural holes facilitate entrepreneurial behavior—strategic exploitation of networks by knowing, shaping, and controlling rewarding opportunities—women are less likely to reap career rewards from these networks (Burt 1998: 10-11). Although Burt attributes this difference to women's discounted legitimacy, it is also true that there are different norms about appropriate networking strategies for men versus women (Carli 1989, Ibarra 1997). Thus, the entrepreneurial behavior that allows men to extract returns from more network connections may reflect a strategy that is more socially acceptable for men than women. Similarly, to facilitate advancement in creative careers, artists must exploit their networks in an entrepreneurial way—e.g., musicians must seek out gigs, promoters, and audiences—but these self-promotion activities are often more socially acceptable for men (Miller 2016).

In the context of creative production, however, extracting value from large networks comes at a cost—bigger may not always be better. Although large networks tend to be associated with greater creativity (Burt 2004, Uzzi and Spiro 2005), collaboration can hinder the production

[5] While we focus our discussion below on why and how these factors might shape women versus men's creative products, we do not develop formal hypotheses. We believe that extant theory and empirical evidence is either too sparse or too inconclusive to generate strong *a priori* predictions.

of novelty in certain contexts, as the coordination required triggers certain widely-shared conventions and rituals (Becker 1984). Extracting new and different ideas from networks and integrating them into a coherent whole requires a specific set of collective behaviors, such as actively seeking help from and providing help to others (Hargadon and Bechky 2006); idea taking and idea giving (Elsbach and Flynn 2013); as well as asking questions, trying out ideas, and collectively negotiating solutions (Harrison and Rouse 2014). Without these practices in place, more collaborators may increase coordination costs and decrease creative production (Harrison, Hagvedt and Askin 2019).

The tradeoff between access to diverse ideas and the coordination costs associated with large collaborator networks may vary by gender as well. Unlike entrepreneurial behavior, which is often associated with masculine norms, some behaviors associated with collective creativity are more socially acceptable for women than men (e.g., seeking help, asking questions, idea taking). Indeed, the dominant form of masculinity in Western societies conceptualizes these behaviors as contradictory to the self-reliance and independence associated with "manhood" (Connell [1995] 2005). Thus, in the context of creative production, female artists may actually benefit *more* from large collaboration networks than male artists, because the latter are constrained by expectations that "real men" do not engage in behaviors like seeking help.

Network Composition

In addition to the *size* of their networks, women may also experience differential returns to the *composition* of their networks. One of the foundations of social networks research is that connections are often formed along the lines of demographic similarity (McPherson, Smith-Lovin and Cook 2001). As a result, women tend to have other women in their networks. This in turn can be negatively associated with power and legitimacy in the workplace, as women tend to be statistically underrepresented in positions of power in organizations (Brass 1985, Ibarra 1992). Among the research that has examined differential effects on men's and women's careers, results are mixed. One study of film actors' careers shows that cohesive project teams tend to be more damaging to women's than men's careers, which is partially attributed to women's tendency toward having more female-dominated networks (Lutter 2015). On the other hand, a recent study of MBA students found that women's but not men's careers benefited from having other women in their "inner circle," to help them maximize the benefits associated with a large number of diverse connections (Yang, Chawla and Uzzi 2019).

It follows that the relationship between the relative gender composition of one's network and novelty in creative production is likely be moderated by gender. Collaborating with dissimilar others usually has a positive effect on creativity: diverse contacts provide creators with access to new perspectives and information, serving as fuel for creative ideas (Burt 1992). In professions like music where women are numerically underrepresented (Smith et al. 2019), most collaboration networks will be populated with more men than women. Thus, having more women in one's network may introduce an artist to more novelty-rich diversity, especially if the artist in question is a man.

Genre Composition

We also posit that the gender composition of an artist's primary genre or category affiliation will shape creative production. Research suggests that women tend to cluster in specific specialties, roles, and categories within occupations. For example, female lawyers cluster in family and estate law (Gorman 2006), female engineers in non-core social work activities (Cech 2013, Cardador 2017), and female investment bankers in public finance (Roth 2006). When the composition of an occupation or category is skewed toward a single gender, that category becomes synonymous with that gender (Ridgeway 2011). Moreover, due to enduring beliefs that women are less status-worthy and authoritative than men, certain categories tend to lose status when large numbers of women enter them (England 1992, Reskin and Roos [1990] 2009). Work-related roles and categories that are dominated by or associated with women are thus likely to acquire lower status and authority through their mere presence, a process known as "occupational gendering."

The analog to occupational categories in spheres of creative production is genre. Genres represent a critical cultural resource for creative producers (Zukin and DiMaggio 1990, Bourdieu 1996), binding them together with like-minded colleagues and consumers while generating and reinforcing prescriptions that influence artistic identity and sound (Frith 1996), as well as audience selection and consumption behavior (Askin and Mol 2018). Much like other categories, genres with a larger proportion of women tend to be lower status. For example, male-dominated musical genres have historically been considered more "serious" and open to experimentation (Schmutz 2009). As a result, artists affiliated with female-dominated genres may experience status insecurity, becoming more likely to conform to conventional production practices than those in male-dominated genres (Phillips and Zuckerman 2001).

We once again predict that the relationship between the gender composition of an artist's primary genre affiliation and the novelty of her products will be moderated by gender. As a result of men's higher status (on average), we expect that working in a gendered role has stronger effects for men than women (Doering and Thébaud 2017). Affiliation with a female-dominated genre is thus likely a greater status threat for men because their higher status means that they have farther to fall. As the higher-status gender, men may have more to lose from their association with a lower-status category, thereby increasing their status insecurity and likelihood to conform, rather than produce novelty.

Data and Methods

The recorded music industry represents an ideal setting to investigate gender differences in creative production. Indeed, creativity plays a central role in musical composition: much like painters, writers, stage directors, chefs, and scientists, the work of musicians is concerned with the production of novelty—not only as a means to an end, but also as an end unto itself. The most successful hit songs exhibit an optimal degree of sonic differentiation (Askin and Mauskapf 2017), and novelty is typically rewarded by critics (via awards like the Grammys) and peers (Choo and Mauskapf 2019).

Moreover, recent advances in the quality, granularity, and accessibility of digital data make it possible to study musical creativity at scale in a way that was not possible ten or even five years ago. It is not an exaggeration to say that the availability of data on production and consumption patterns in the music industry is unrivaled by any other field of cultural production. These advances in computing power and sophistication allowed us to construct a unique dataset describing more than 250,000 songs written and released by 15,000 unique focal artists between 1955 and 2000. We focus on this timeframe because of its relevance to the history of modern popular music (Peterson 1990) and, more practically, the amount of missing data prior to this period. We stop our analysis in 2000 because we believe that year represents a permanent shift in the way music was recorded and distributed. Although Napster and other illegal file-sharing services started to appear during the late 1990s, iTunes was released at the beginning of 2001, drastically reducing the barrier to self-produce and release music. Given this shift, we believe the platform's introduction serves as a good dividing line for our analysis, ensuring the internal consistency and validity of our dependent variable. We believe that this dataset is the best available approximation of the world of commercially recorded music released by individual focal artists for whom we were able to collect complete gender and sonic feature data.[6]

A few additional notes about the two primary sources we used to construct the dataset. MusicBrainz, a curated and crowd-sourced online music database, provided us with detailed artist credits, describing who composed, wrote, and/or performed on each of millions of tracks or albums. MusicBrainz also provided record label and day-level release date information, as well as gender designations for all individuals involved in the writing, recording, and production of each song. Unlike many other large-scale studies of gender differences that rely on computational methods to assign gender, we have access to gender coding that was the result of crowd-sourcing by community members and was double-checked by staff moderators. Based on repeated spot-checks, we have a high degree of confidence that the gender codes assigned to the artists in our dataset are accurate.

To construct our song-level measures, we collected data from The Echo Nest, a data science and Music Information Retrieval (MIR) company owned by Spotify. Before advancements in MIR, musicologists or other highly-trained individuals would have to listen to and hand code the sonic attributes of songs if they wanted to compare and analyze them, a labor- and time-intensive activity that necessarily limited the scope of what could be studied (e.g.,

[6] We recognize that any dataset of this type suffers from some selection bias. We don't capture songs that were never commercially recorded or released, and our data skews slightly toward western popular music. However, our sample is far broader than, for example, the approximately 20,000 songs that appeared on the *Billboard* Hot 100 charts during this time period.

Cerulo 1988, Dowd 2002). In much the same way that Natural Language Processing (NLP) techniques have enhanced our ability to analyze the content of text at scale (Blei, Ng and Jordan 2003, Bail 2014), MIR combines musicology with machine learning and computer science to extract unique acoustic "finger prints" from audio files, providing both high- and low-level feature data describing the sonic attributes of tens of millions of songs.[7] This data provides a simplistic but relatively accurate representation of how music sounds and is experienced by producers and consumers, and captures familiar musicological information as well as a variety of perceptual features.

Dependent Variable

Song Novelty. We measure song novelty using a feature-based measure of musical distinctiveness for each song in our dataset. Work in psychology suggests that low-level descriptive attributes play an important role in shaping how actors interpret and make sense of products (Tversky 1977). In music, preferences are linked to a series of features that structure musical space, such as speed, repetition, sadness, and loudness (Greenberg et al. 2016). While previous research has recognized that creative production occurs through the recombination of both category labels and material features, most empirical studies have used only labels to describe and compare products (for exceptions, see Cerulo 1988, Askin and Mauskapf 2017). All

---

[7] For our analyses, we use the summary features, rather than the beat-level representations for each song. This is because (1) beat-level data is available for fewer songs, and we were interested in including as many songs as possible in our analyses, and (2) it is significantly more complex to aggregate, analyze, and compare beat-level data across songs. We do not account for the lyrical content of songs in our measure of novelty.

songs in our dataset were algorithmically assigned a value for ten features, each of which was developed by The Echo Nest to describe a song's most important sonic characteristics. Features include several standard musical attributes (e.g., "tempo," "mode," "key," "time signature"), as well as a series of algorithmically-derived measures that represent particular aural or emotive dimensions of music ("valence," "danceability," "acousticness," "energy," "liveness," and "speechiness"; citation removed to preserve author anonymity, more background information available upon request). While these sonic features distill the considerable complexity of music into a handful of discrete statistics, and thus fail to capture what makes music "art", they are well suited for comparing songs to each other (Friberg et al. 2014).

After normalizing all ten features and collapsing them into a single vector for each song, we calculate the cosine distance between all relevant song-pairs. To define the relevant comparison set for a focal song, we include all other songs released within a rolling 10-year window prior to the focal song's release date.[8] Thus, a song released in January 1990 is compared to all songs released between January 1980 and December 1989. Each comparison generates a pairwise distance measure across our 10-dimensional feature space, which we then sum, average, and multiply by 100, producing a single novelty score between 0 (maximum

[8] Note that this comparison set includes both between- and within-genre comparisons for each song, rather than simply within-genre comparisons. The latter specification may be appropriate for a consumer-centric measure of novelty, but we believe that a comprehensive comparison set is more appropriate for a producer-centric measure, as many artists span multiple genres. We also conducted analyses using 5-year, 2-year, and no time windows. Results are robust across each of these specifications and are available upon request.

conformity) and 100 (maximum novelty) for each song.[9] The mean novelty score for songs in our dataset is just over 21, highlighting the fact that most songs sound relatively similar to each other. Given the right-tailed skew of this variable, we take its natural log for the analyses that follow (See Table 2.1 for raw descriptive statistics).

Primary Moderating Variable

Female Artist. Our primary moderating variable is the gender of individual artists in our data set. We recognize that gender is a social construct (Butler 1988) but follow convention and use it as a proxy to study differences between men and women. Although we have extensive data on bands and the gender composition of their members, we opted to limit the analyses presented in this paper to individual focal artists (rather than bands) for the sake of simplicity and interpretation.[10] We use an indicator variable equal to 1 for "Female" artists and 0 otherwise

---

[9] Although this is a continuous measure of novelty, Table A2.4 presents results from a series of logistic regressions that predict the release of an "optimally" versus "radically" novel song (defined as falling in the 70^{th}-to-90^{th} percentile vs. the 90^{th}-to-100^{th} percentile of our measure). Results from these models suggest that there is no systematic difference in the effect of being female on producing songs exhibiting moderate vs. radical novelty.

[10] We chose not to include musical groups with fixed membership, because (1) identifying and modeling a band's gender "assignment" becomes much more complex, and (2) we believe that the gender dynamics of production within teams with fixed membership is substantively different than solo artists with more flexible collaboration networks. Note that we do account for collaborators who worked with individual focal artists, which include writers, composers, and often times performers (if they are credited as a creator by MusicBrainz).

("Male" or "Other").[11] In our data, we see an approximate 3.5:1 ratio between male and female artists, a ratio that very nearly matches contemporary studies of female selection into in the professional music industry (Smith et al. 2019). Interestingly, this under-representation is more-or-less constant across the entire time horizon of our analysis (1955 to 2000).

Independent Variables

Network Size. Data from MusicBrainz contains extensive artist credits for individual songs, allowing us to construct a rolling count of active unique collaborators for each focal artist in our dataset. Common types of artist credits include "composer," "writer," "producer," "instrumentalist," and "vocalist," among others. Regardless of the type of connection, we believe each person in an artist collaboration network is a potential creative influence, and thus we treat them all equally. "Active" means that a collaborating artist has (1) already appeared on a release with a given focal artist and (2) has not retired, and "unique" means that we only count each focal artist-collaborating artist pair once. We choose to use a simple in-degree measure of connectedness, rather than a composite measure of network position, because each ego network is non-hierarchical with a density of 1.

Network Composition. We include the gender composition of each artist's collaboration network. Using the above count of a focal artist's active and unique collaborators at the time of a given song's release, this variable is measured by taking the simple proportion of female (relative to male) collaborators in a given network.

---

[11] We have run our analyses both with and without the very small number of "other" artists in our data set; results remain nearly identical due to the very small percentage of the data set falling into this category.

Genre Composition. Genres are one of the most salient means of categorization in the music industry. Labels like "rock," "country," and "hip hop" were originally created for marketing purposes, making radio advertising more effective (Peterson 1990). Genres quickly became the key organizing force in the music industry, shaping the behavior of artists, record labels, critics, and other consumers (see Bourdieu 1993, Lena 2012). Although they continue to evolve and can be highly contested within and across communities (Lena and Pachucki 2013), genres continue to play a central role in creative production.

The emergence and attribution of genre labels was historically determined by a complex interplay between fans, institutional gatekeepers, and artists (Lena 2012), but advances in technology have produced a more emergent labeling process. To measure the genre gender composition for each song release, we first identified the focal artist's primary genre. We collected all of the unique genre "terms" attributed to every artist in The Echo Nest's digital catalogue. These attributions are generated using a combination of web-crawling and text parsing algorithms designed to identify the shared labels people use to describe and classify artists.[12] Our dataset contains over 900 unique terms, covering common (e.g., "blues", "opera", "jazz") and less common (e.g., "raggacore", "darkstep", "psychobilly") genres. Terms are attributed at the artist level, and artists can be labeled with more than one term. Each artist-term pair includes a weight that represents the relative "strength" of the attribution; for this measure, we only include each focal artist's primary genre (i.e., the artist-term pair with the highest weight).

---

[12]See http://blog.echonest.com/post/73516217273/the-future-of-music-genres-is-here for a more detailed explanation.

To measure genre composition, we calculate the simple proportion of female (relative to male) artists active in a focal artist's primary genre at the time of a given song release.

Control Variables

Past Creativity. There can be large individual differences in creativity due to differences in motivation, skill, and cognitive style (Amabile and Pillemer 2012). To account for these individual differences, our main models control for a focal artist's past creativity, which is simply calculated by taking the average novelty score for all songs released previously by that artist.[13] The relative effect size of this variable can serve as a helpful benchmark for the rest of the results in our analysis. Note that we do not include artist-level fixed effects, as they would eliminate our ability to interpret a main effect on artist gender, which is the primary goal of this paper.

Artist Productivity. Linus Pauling, winner of two Nobel prizes (Chemistry and Peace), once quipped that, "If you want to have good ideas, you must have many ideas."[14] Part of the experimentation process is generating lots of outputs, knowing that there is always a possibility of producing something new and exceptional. Thus, the number of songs produced by a focal artist may covary with song novelty. We control for the total number of songs that an artist has

---

[13] We also ran a specification where past creativity was measured as the average of an artist's songs released in the previous two years only. Results were robust across both specifications.

[14] http://scarc.library.oregonstate.edu/events/1995paulingconference/video-s1-2-crick.html

released since the beginning of her career, standardized by the number of years she has been active in our dataset.[15]

Artist Tenure. An artist's past experience may also affect the relative novelty of their musical output. Experience represents the knowledge, skills, and competencies that an individual accumulates over time, sometimes referred to as "human capital" (Becker 1994). Experience can confer legitimacy, trust, and the resources necessary to explore and experiment. The more time someone spends embedded in a given field of production, however, the more likely she is to internalize its norms and conventions. To operationalize experience, we construct a measure of artist tenure. This is calculated as the length of time (in days) since a focal artist's first release in our dataset. We also include a squared term for artist tenure in our main models, as much of the previous research finds a curvilinear effect of tenure on performance (Cole 1979, Jones and Weinberg 2011).

Popular Success (Charting Artist). Popular success and widespread appeal are likely to shape an artist's propensity to produce novel songs as well. The relative (a)typicality of a creative product plays an important role in its success (Uzzi et al. 2013, Askin and Mauskapf 2017), but this relationship is bi-directional: an artist's previous successes (or failures) can shape their future production efforts as well. For example, while achieving rock-star status might seem like a license to experiment freely, it might also generate market pressures to conform to audience expectations. By reproducing a recent hit, an artist might hope to "catch lightening in a bottle," thereby reliving their success and reaping the outsized rewards reserved for superstars

[15] We also ran a specification where productivity was measured as the number of songs released by an artist in the previous two years only. Results were robust across both specifications.

(Rosen 1981). While we cannot disentangle which of these forces is at work, we do know that popular artists are slightly less likely to produce novel songs overall (see Askin and Mauskapf 2017). To measure a focal artist's popular success at the time of a given song's release, we include a dummy variable set to 1 if she had previously placed one or more songs on the *Billboard* Hot 100 Charts, 0 otherwise.

Major Label. We distinguish between major versus independent record label representation to account for differences in the organizational support received by artists. Some organizations develop structures, practices, and competencies that support creative exploration and innovation, while others do not (Amabile 2012). In the music industry, commercial recording artists usually receive organizational support from one or more record labels. For much of the time period covered in our dataset, a handful of so-called "major" labels were responsible for releasing the vast majority of popular music that listeners could hear on the radio or purchase to consume privately (Peterson 1990). Some of these labels developed strategies to respond to or even preempt changes in technology and consumer taste, increasing the creativity (Peterson and Berger 1971) and success (Benner and Waldfogel 2016) of artists on their rosters. For example, Lopes (1992) found that, in the 1970s and 1980s, major labels were more likely than independents to support an open production environment, leading to more diversity and innovation in musical production. Today, however, "indie" labels may be better positioned to recruit and cultivate creative artists, as their smaller size and peripheral position in the market make it easier for them to drive change (Hesmondhalgh 1999). To account for this, we include a dummy variable set to 1 for all songs released on a major label, 0 for independent label releases.

Classical or Jazz. The vast majority (>80%) of songs in our dataset are part of the popular music tradition, but some fall into the category of classical music and/or jazz. Due to the complex form, unique instrumentation, and other compositional characteristics associated with works in these genres, we worried that observations labeled with any variation of the term "jazz" or "classical" were likely to be systematically more novel. In addition to running models that exclude all songs affiliated with these two genres, our main models include a dummy variable for each.

# of Genres Spanned. Many cultural producers, including musicians, affiliate or identify with multiple genres. Spanning categories is often penalized in labor and product markets, as it can signal failure (Zuckerman et al. 2003) or a lack of commitment (Leung 2014). More recent research, however, suggests that genre-spanning can increase access to diverse viewpoints and other cultural and symbolic resources (Hsu 2006). As artists begin to create new work, they may mine multiple art worlds for inspiration and raw materials, recombining them to create something new and different (Becker [1982] 2008). Most work assumes that genre spanning and product novelty are positively correlated, but in contexts such as music, where the integration of resources from different categories is complex and can lead to cognitive overload, this is not always the case (Mauskapf et al. 2019). To measure the extent to which a song has spanned genres, we generate a simple count of the unique "terms" attributed to a given focal artist in the Echo Nest's digital catalogue. For this variable, we include all artist-term attributions with weights of > 0.5.[16]

---

[16] We chose a threshold of 0.5 for two reasons. First, we wanted to account for the fact that many artists and their music are affiliated with more than one genre. So, rather than select an artist's primary genre (weight = 1), we collected all available artist-term attributions. Second, after reviewing the data, we realized that, while all artists are tagged with at least one genre, and most

Country. During the timeframe covered in our analysis, the majority of the music industry's commercial activity took place in the United States. Most of the major labels were based here, and advances in the market meant that by the second half of the 20th century, it had become the epicenter of musical production. This reality is underscored by the fact that our data skew toward English-language music: 45% of the artists in our data are US-based, and another 25% are based in the UK. To control for the outsized influence that the US had on the means of production, we include a dummy variable set to 1 for any song released by a US-based focal artist.

Decade. To account for large-scale historical shifts, changes in production and distribution technology, and the emergence and evolution of exogenous fads and fashions, we include dummy variables for the decade in which a song was released. The 1960's, one of music's most novel eras, serve as the reference group.

Analytical Strategy

The analysis that follows is divided into three sections. Because we did not have strong *a priori* predictions about whether (or when) we would find gender differences in the relative novelty of creative products, we begin by conducting some descriptive explorations, including a

---

are tagged with no more than three or four, there is a long right tail of (mostly highly visible) artists affiliated with many genres. We did not want over-attribution among popular artists to bias our results, nor did we want to include weak attributions that were made by only one or a few audience members. Thus, we only included terms that appeared often enough alongside a given artist that they warranted a weight above 0.5. Full descriptive statistics for this variable are included in Table 2.1.

simple comparison of the means and distributions of song novelty for male versus female artists (Table 2.1, Figure 2.1). Next, we ran pooled, cross-sectional OLS regressions to estimate the effects of being a female artist on song novelty, gradually incorporating our controls and the three contextual variables described above (Table 2.3, Figure 2.3). To better understand the conditions under which women are more (or less) creative than men, we then ran another series of OLS models with interaction terms between gender and the contextual variables (Table 2.4). To aid our interpretation of these results, we generated a series of figures that plot (a) the raw distributions of the contextual variable in question, and (b) its marginal effects on song novelty, by gender (Figures 2.4a-c). Finally, we also conducted several robustness checks, re-estimating our main model: by gender sub-sample (Table A2.1); with Coarsened Exact Matching (Iacus, King and Porro 2012) (Table A2.2); by artist tenure sub-sample (first song vs. tenure > median), to try and identify the role of selection on our results (Table A2.3); and with logistic regression, estimating entry into bins of "optimal" versus "radical" novelty (Table A2.4). These tables appear in an informal appendix without further comment. We are happy to provide additional specifications upon request.

Results

Descriptive Analysis

Before summarizing our regression results, we first explore simple differences between male and female artists and their creative output. Tables 2.1 and 2.2 summarize some descriptive statistics and correlations for our key variables of interest. Figure 2.1 consists of two histograms that compare the raw distribution of our dependent variable—song novelty—for male and female musicians. At first glance, the distributions look remarkably similar: there seems to be little if any gender difference in the overall levels of creative production in our dataset. We also conducted a series of t-tests to compare the mean song novelty of men and women, as well as Kolmogorov-Smirnov (KS)-tests to compare the two distributions. In light of the 3.5:1 ratio of male to female artists in our data, however, we wanted to compare apples to apples; that is, we wanted to compare an equal number of male and female artists. To accomplish this, we down-sampled musicians from each gender, comparing songs by 2000 male and 2000 female artists to each other. We then bootstrapped each of these statistical tests (t- and KS-tests), running each on 1000 randomly-generated samples to ensure that we were not picking up noise that could be biasing any individual subsample result. The mean t-statistic across 1000 samples is -0.09, indicating no statistically-significant difference in mean-levels of creative output between genders. The average p-value from 1000 bootstrapped two-sided KS-tests is .006, reflecting a small but statistically significant difference in the underlying distributions of song novelty. Not surprisingly, male artists have more songs in both tails of the distribution, which is likely caused by female artists' underrepresentation in the musical labor market.

---Insert Tables 2.1 and 2.2 here---

---Insert Figure 2.1 here---

In Figure 2.2, we provide a two-dimensional depiction of how 10,000 randomly-selected songs are distributed across sonic feature space (Maaten and Hinton 2008). Each dot in the figure represents a song, and the distance between dots represents the extent to which two songs "sound" different from one another based on their sonic features. This plot suggests that, while there are clearly many more songs produced by men (orange) than women (blue), they do not occupy systematically different regions of feature space. In addition to visualizing the sonic foundations of our dependent variable, this plot should alleviate any concerns that the underrepresentation of women in our data is somehow biasing female song novelty scores and, by proxy, our results.

---Insert Figure 2.2 here---

Regression Results

Table 2.3 includes results from our main suite of regression models predicting song novelty. Model 1 is a simple regression containing only our female indicator variable, and we again find that in the absence of contextual factors and other controls, there is no significant relationship between artist gender and song novelty. Once we begin to add in controls, however, the coefficient for female artists becomes positive and statistically significant, and the magnitude of the effect continues to grow as we account for additional contextual variables. For example, in Model 3 we find that our main effect for artist gender holds even when accounting for an artist's past creativity, which is unsurprisingly a strong positive predictor of subsequent creative output.

---Insert Table 2.3 here---

Most of our controls behave as expected. We find a U-shaped effect of artist tenure on song novelty. Early on in a musician's career, song novelty tends to decrease with each release, but that trend reverses itself as careers advance. The exact turning point appears to be after about 10 years in the field. Songs by more popular artists and major labels tend to be less novel, while those affiliated with only one genre (49% of female artists and 44% of male artists in our data) are likely to produce more novel songs than artists who span two or three genres. The real benefits of genre recombination for creative production are only felt for the small subset of artists in our dataset who sample from a wider range of genres.

Finally, we turn to our full model (Model 6). Figure 2.3 displays standardized beta coefficients from this model to highlight the substantive significance of our variables of interest.

While the two largest effects are associated with an artist's prior creativity (positive) and popular appeal (negative), being female is among the strongest predictors of song novelty. Indeed, after incorporating all of the structural and cultural conditions that we believe may aid or inhibit creative production, the coefficient for *Female* remains positive and significant. For example, we find that having a larger network is negatively associated with song novelty on average, highlighting the potential coordination costs associated with collaboration. Although this negative effect is relatively small and challenges some of the research linking network size with creative output, it supports more recent work that suggests smaller teams may be responsible for more disruptive and radical innovations (Wu, Wang and Evans 2019), while large networks tend to be homophilous, stunting creativity (Uzzi and Spiro 2005).[17] As expected, we also find that the effect of being affiliated with a female-majority genre is negative, while the effect of having more women in one's network does not meet the threshold of statistical significance in this model. Thus, artists with larger networks tend to release less novel songs, regardless of the network's gender composition.

---Insert Figure 2.3 here---

Interaction Effects

The results presented in Table 2.3 suggest that female artists with equivalent access to critical resources produce more creative songs than their male counterparts. To better understand when women (and men) put out more novel songs, we introduce a new set of models with

---

[17] We also tried including a squared term for network size, to test for the possibility of a non-linear relationship between the number of collaborators an artist has and her song novelty, but the effect was not significant.

interactions. Rather than walk through all of the results included in Table 2.4, we summarize the key findings using Figures 2.4a-c, which highlight marginal effects from the fully specified model (10) in Table 2.4. These plots are presented alongside raw histograms for each of three structural conditions discussed above, separated by gender.

---Insert Table 2.4 here---

---Insert Figure 2.4a-c here---

Remember that in Model 6, we found a small but negative relationship between collaboration network size and song novelty. Figure 2.4a provides evidence of a significant *boost* in novelty for female artists who have larger networks, while the opposite is true for men. Most artists have a relatively small number of collaborators (median = 3 for women, 4 for men), but for the nearly 10% of artists whose networks feature 10 or more unique collaborators, the benefits (costs) associated with women's (men's) creativity becomes significantly more substantial. Even at their median network sizes, women produce more novel songs than their male counterparts. Female artists may benefit from having more collaborators because key behaviors used to take advantage of collaborators' diverse knowledge (e.g., help-seeking, asking questions, and taking ideas) are more socially acceptable for women, while the coordination costs for men seem to outweigh the potential benefits associated with collaboration.

Interestingly, we also find that having a collaboration network populated by more women is more positively associated with creative output for male rather than female focal artists (Figure 2.4b). Given the underrepresentation of women in the field, both genders tend to have networks with considerably more men than women. In fact, most men have never collaborated with a female artist (or other type of collaborator). Yet when men do work with women, it enhances their creativity. The interaction between being female and having a greater-than-average proportion of female collaborators is significant and *negative*. Collaborating with other women does not necessarily hurt female artists' propensity to produce song novelty, however—it simply does not help them as much as it helps male artists. Presumably the difference here stems from the robust positive effects associated with greater access to diverse ideas, which women may not necessarily experience when collaborating with other women. Network diversity is key. Unfortunately, while men's propensity to create novel songs increases with more female collaborators, such collaborations are not particularly common occurrences.

On the other hand, we find a significant *positive* interaction between being female and having a greater-than-average proportion of women in one's primary genre. Figure 2.4c shows that both men and women tend to affiliate with genres that are majority male, which is again unsurprising given the relative scarcity of female artists in this context. Yet while the main effect of affiliating with a "female" genre like vocal pop on novelty is strongly negative, that negative association is stronger for men than it is for women. Men who operate in female-dominated genres are simply more likely than women to experience a significant decline in creativity. As previously described, we suspect that this is due to male artist's greater status insecurity and the subsequent conformity pressures they face when affiliating with lower-status genres.

Discussion

Our results reveal several important insights regarding the presence and nature of a gender gap in creativity in music. On one hand, there is no statistical difference in raw product novelty between male and female artists, whether simply comparing group means or regressing a gender indicator on creativity with no additional control variables. On the other hand, once we begin to account for relevant contextual variables such as the size of an artist's collaboration network and the gender composition of their network and primary genre, we find robust evidence that women create more novel songs than their male peers. Structural and cultural factors clearly play an important role in the creative production process, differentially shaping the creative output of men and women.

We believe that these results provide support for the claim that women are held to a higher standard in the market for recorded music. Consistent with past research on women's career advancement, we find that the men and women in our study work under different structural conditions: female musicians are more likely than men to work in female-dominated genres and have collaboration networks that are smaller and more female. We diverge from past research, however, by finding that women's tendency to cluster in certain genres and networks has mixed effects on creative production. Men experience a larger penalty for being affiliated with "female" genres, but benefit more from having women in their collaboration networks. We also find that female artists benefit creatively from having larger collaboration networks, whereas male artists do not.

One possible alternative explanation for our findings could be that women are *too* creative on average for conservative audiences and industry gatekeepers. However, if this were the case, we would expect to see significant differences in raw creative output, and we find none. This explanation is also contradicted by our finding that there are more men than women in the right tail of the novelty distribution We argue instead than women face a higher bar when it comes to breaking into the music industry, echoing previous research that identifies the double standards faced by women in the context of peer review (Hengel 2017) and electoral politics (Anzia and Berry 2011).

Furthermore, although our analysis focuses on differences among cultural producers, we see some evidence of these higher standards shaping not only the work experiences of men and women, but also women's selection into the music market and their survival over time. We do not have access to the upstream decisions about which artists are scouted, signed, and subsequently produced by record labels, nor do we know who gets selected out of the various stages of this process, but evidence from academia suggests that it is harder for women to enter fields like music composition and philosophy in which genius is considered important for professional success (Leslie et al. 2015, Meyer, Cimpian and Leslie 2015). The extent to which this contributes to female underrepresentation in music is beyond the scope of our study, but, much like stereotypes in other occupations (e.g., Gorman 2005, Cech et al. 2011), it may effect women's likelihood to enter—or be allowed entry into—the music profession. To try and assess whether the differences we find between men's and women's creative output exist upon entry, we estimated our main model for both artists' first song release and songs released after median tenure (see Table A3). These estimates provide at least some evidence that, all else equal, women produce more novelty upon entry into the market, and this difference persists over time. We hope that future research on this topic will identify the specific mechanisms responsible for entry and survival in creative fields like music, and the role creativity plays in these processes.

Limitations

While we believe that our findings provide compelling evidence for the gendered role of networks and genre in musical creativity, we rely on several important assumptions. First, we assume that our measure is a valid proxy for the "actual" novelty or distinctiveness of a song. Rather than aggregate subjective evaluations made by consumers, we chose to use the sonic features of songs to construct a materially-informed and hopefully more objective measure of product distinctiveness. This measure has its limitations: for example, we do not account for song lyrics as a potential source of differentiation (e.g., Berger and Packard 2018), although such a comparison would reduce the breadth of our analysis and further complicate calculations of song novelty. We try to take full advantage of the quantity and quality of music-related data available today to construct a scalable measure of product novelty that is both internally and externally valid.

Our analysis also presupposes that novelty, and creativity more generally, is valued in the domain of music. Empirically, we find musicians who have previously appeared on the *Billboard* Hot 100 Charts and those represented by major record labels tend to produce less novel songs. We recognize that audiences' appreciation of creativity and novelty is worthy of study, but we also believe that novelty is an important outcome in and of itself, especially in music and other creative industries. Moreover, while novelty and commercial success are clearly orthogonal constructs, they are not antithetical. Given that the median of our novelty measure is relatively low (21.6 on a 0-100 scale) and the interquartile range is narrow (17.4-24.4), songs displaying some small or "optimal" degree of differentiation may simultaneously achieve popular success and be heard by audiences as creative or different (cf., Askin and Mauskapf 2017). Further research is needed to unpack this relationship and understand when and how it is moderated by gender.

Another limitation is that, while our data allow us to account for many of the structural and cultural conditions that shape creativity, there are other factors that we cannot address—including the influence of individual record label representatives who influence the production process and. the pressure artists put on themselves to be more novel or conventional, to name two. Furthermore, our analysis stops in the year 2000. Although we have little reason to believe that things have changed in the intervening twenty years, we hope future research investigates how creativity and gender dynamics unfold in the twenty-first century. In this paper we focus exclusively on individual focal artists and their collaboration networks, choosing not to address the additional complexities that may arise in the context of musical groups or bands. Others have explored the effects of group gender composition on various performance-related outcomes (Apesteguia, Azmat and Iriberri 2012), and we intend to examine how changes in group composition may affect creative output in future work.

We also focus exclusively on between-artist differences in creativity. Although we control for individual artist's historical novelty production, we do not test for differences in the internal heterogeneity of creative output. The marginal effects of artist tenure for female versus male artists are nearly identical, but it would be worthwhile to compare career trajectories of creativity more rigorously. Future research should investigate how gender intersects with other minority categories such as race and ethnicity to affect creativity as well. Finally, while this paper emphasizes production dynamics, industry gatekeepers and consumers inevitably play an important role in shaping the production of novelty. Exploring the demographic distribution of consumers and linking differences in product novelty with consumer evaluations of creative achievement will help us to understand how both sides of the market shape one another.

Conclusions

We believe our findings make four primary contributions. First, we demonstrate that novelty in creative production is correlated with gender. Whereas prior research assumed that men's creative products were more novel than women's, or that there were no differences at all, we find that female minorities' products exhibit more novelty when compared to the products of men in similar structural and cultural positions. This suggests that gender bias surrounding the evaluation of creative products may be even *stronger* than previously assumed. Not only are women less likely than men to be recognized as creative—they are also more likely to produce novel creative products.

Second, we show that key structural and cultural conditions related to women's career advancement are also related to their creative output. When we compare men and women with similar collaboration networks and genre affiliations, we find that women are producing more novel work than men. This suggests that their creativity may be suppressed by different structural and cultural conditions of work. At the same time, we find that these conditions operate differently for men and women. Resources such as network size, network composition, and genre composition may enhance or inhibit artists' creative efforts. Recognizing when and how gender moderates these effects is a critical first step in addressing inequalities in the music labor market. And, because these resources and dynamics extend well beyond the domain of music, our findings can contribute to our understanding of gender differences in cultural markets more generally.

This study also contributes to the nascent stream of research on women's career advancement that examines whether men and women in the same occupations differ in the work they produce. Studies of gender differences in the *quantity* of work performed are rare and results are mixed: women in science tend to publish fewer articles (Leahey 2006) and patent fewer inventions (Whittington and Smith-Doerr 2008), but those in legal occupations do not significantly differ from men in their number of billable hours (Kay and Hagan 1998). In terms of the *quality* of work produced, research in science has found that women are less likely than men to produce a professional body of work that is internally homogenous or specialized (Leahey 2007). More relevant to our study, scholars have shown that scientists who produce more novel research tend to be more prominent and more likely to win prizes, but also less productive (Foster, Rzhetsky and Evans 2015, Leahey, Beckman and Stanko 2017). Research in this area has not yet considered the extent to which men and women differ in terms of the novelty of their work, which is likely to have important professional consequences in fields like science and music that value creativity.

Finally, we introduce a new means of studying and measuring creative output and cultural production more generally. Creativity research has historically relied on subjective ratings to capture the distinctiveness of a particular product or idea. We instead leverage advances in the field of music information retrieval (MIR) to generate a more objective measure of product novelty using song's sonic features. We hope that our study, along with others capitalizing on algorithmically-generated, large-scale feature data (Bail 2014, Lazer and Radford 2017), will encourage scholars to continue exploring the nature of creativity in new and different ways.

Just as importantly, however, we hope that our study will encourage individuals and institutions involved in creative production to think more deeply about how to address gender inequality. The underrepresentation of women in these roles and professions has recently been brought to the forefront of public conversations among industry gatekeepers (Werde 2018), and all-female supergroups like the "The Highwomen" have been celebrated for "knock[ing] a bigger dent in the … music industry's systemic misogyny" (Richards 2019). Yet as our results highlight, even after women enter these occupations, more needs to be done to level the playing field. Country music has made headlines for keeping women off of its radio stations and charts (Watson 2019), but it is not the only context where female artists face discrimination. Cultivating environments with a more equal distribution of resources and opportunities, coupled with the continued evolution of technology and societal norms, can help to erase the double standards facing women in music, and will ultimately empower all artists to tap into their creative potential.

Chapter Three

What Makes Popular Culture Popular?

Product Features and Optimal Differentiation in Music[1]

What makes popular culture popular? Scholars across the humanities and social sciences have spilled considerable ink trying to answer this question, but our understanding of why certain cultural products succeed over others remains incomplete. Although popular culture tends to reflect, or is intentionally aimed toward, the tastes of the general public, 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 utilize diverse information that might signal the quality and value of new products (Keuschnigg 2015), including the characteristics and networks of cultural producers (Peterson 1997; Uzzi and Spiro 2005; Yogev 2009), audience preferences and social influence dynamics (Lizardo 2006; Mark 1998; Salganik, Dodds, and Watts 2006), elements in the external environment (Peterson 1990), and various institutional forces (Hirsch 1972).

Each of these signals plays an important role in determining which products audiences select, evaluate, and recommend to others. Nevertheless, while these choices and the preferences they represent vary widely over time and across individuals, research suggests that the inherent quality of cultural products also affects how audiences classify and evaluate them (Goldberg, Hannan, and Kovács 2015; Jones et al. 2012; Lena 2006; Rubio 2012; Salganik et al. 2006). Certain product features may independently signal quality and attract audience attention (e.g., Hamlen 1991), but we believe that these features matter most *in toto*, both by creating a multi-

---

[1] This chapter is co-authored with: Noah Askin (INSEAD, Paris).

dimensional representation of products and by positioning those products across the plane of possible feature combinations.[2] Rather than existing in a vacuum, cultural products are perceived in relation to one another, 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, which highlights how social classification systems organize consumers' expectations and preferences (Hsu 2006; Zuckerman 1999) and help them draw connections between products. We agree that categories play a significant role in structuring taste and consumption behavior (Bourdieu 1993), but much of the work in this area makes the implicit assumption that category *labels* remain tightly coupled with a set of underlying *features*. Recent research notes, however, that these features may not cluster or align with prevailing classification schemes (Anderson 1991; Kovacs and Hannan 2015; Pontikes and Hannan 2014).[3] Category labels (e.g., "country" in the case of musical genres) work well when navigating stable product markets with clearly defined category boundaries, but they do not always reflect how

---

[2] 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, along with whether the painting is a portrait or a landscape.

[3] 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 the underlying features of products.

audiences actually make sense of the world in which they are embedded, especially in contexts where products are complex and tastes are idiosyncratic and dynamic (Lena 2015). In these domains, extant categories may not provide adequate or accurate information to consumers, who must instead rely on the underlying features of products to draw comparisons and make selection decisions.

We build on these insights to propose a new explanation for why certain cultural products outperform their competitors to achieve widespread success. In the context of popular music, we argue that audiences use musical attributes or features to draw latent associations between songs. These associations, which are conceived independently from traditional categories, help to organize the choice set from which audiences select and evaluate products, positioning certain songs more advantageously than others. We hypothesize that hit songs are able to successfully manage a similarity-differentiation tradeoff, simultaneously invoking conventional feature combinations associated with previous hits while inciting some degree of novelty that distinguishes 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 (Lounsbury and Glynn 2001; Uzzi et al. 2013; Zuckerman 2016).

To test this prediction and better understand the relationship between product features and success in music, we construct 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 attributes or features that describe a song's sonic quality. Sonic features range from relatively objective musical characteristics, such as "key," "mode," and "tempo," to more perceptual features that quantify a song's "acousticness," "energy," and "danceability," among others. After demonstrating the baseline significance of individual features in predicting a song's peak position and longevity on the charts, we use these features to construct a measure of sonic similarity or typicality and test its effect on chart performance. While 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), 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, cultural content matters, particularly in the way it structures songs' relationships to each other. These findings offer a new contingent perspective on popular culture by specifying how feature-based associations organize the way in which audiences distinguish and evaluate products, compelling us to rethink some of the basic mechanisms behind cultural consumption and taste formation.

Cultural Preferences and the Similarity-Differentiation Tradeoff

Predicting how well a new product will fare in the marketplace for audience attention presents a difficult, if not impossible, challenge, 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), which tend to generate products and experiences whose evaluation involves significant subjectivity (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). The relative popularity of a cultural product is usually ascribed to prevailing tastes, which are largely considered a function of individuals' idiosyncratic preferences, past experiences, and exposure patterns, as well as the prevailing opinions of others. Moreover, different types of performance outcomes (e.g., mass appeal vs. critical acclaim) beget different varieties of explanation, and require audiences to consider distinct dimensions of evaluation that are often context specific.[4] Needless to say, our ability to explain what constitutes a hit versus a flop is limited.

Scholars interested in this question have traditionally taken 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, including their reputation (Bourdieu 1993), past performance outcomes (Peterson 1997), and the structure of their professional networks (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 collaborations between artists and producers display small world properties, their 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, Stark, and Vedres 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 through 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] We use the term "success" in this paper to connote mass or popular appeal, rather than critical acclaim or other legitimate measures of performance.

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 role in determining social reality—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.

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 variances 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; Salganik and Watts 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 a mystery.

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). 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), where classification systems provide the context through which producers and consumers structure their tastes, preferences, and identities (Bourdieu 1993; Peterson 1997), and determine how they 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, including 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), but this work has several important limitations. Although 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. We know, however, that categories and their boundaries are dynamic and eminently contested, signifying different meanings to different communities (Lena 2012; Sonnett 2004). Moreover, most research in this area highlights the social-symbolic labels attached to categories, ignoring the material features of the products that occupy them. While labels constitute socially constructed and symbolically ascribed descriptors for a given category, features provide considerably 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, Stark, and Vedres 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 dictum 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 countless features, including its duration (6:09), key (C Major), instrumentation (vocals, guitar, bass, electric organ, harmonica, tambourine), and thematic message (love, resentment). From our perspective, these features—the inherent, high-dimensional attributes that constitute the "DNA" of individual products—are culturally determined, grounding 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 sensemaking in distinct ways (Tversky 1977). Furthermore, empirical evidence from popular music 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 that there is a gap in the way we conceptualize features and their role in positioning products for success. Rather than influencing consumption independently, we believe that 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 cognitively structure and make sense of these representations (Brashears and Quintane 2015; Zerubavel et al. 2015), but we still know little about how this process unfolds.[5] In the context of consumption, we argue that consumers position products across some multi-dimensional feature space, whereby certain objects are perceived to be more (or less) similar depending on the features they do (or do not) share. These latent associations represent the world of products in which consumers are embedded, and exhibit a social life all their own (Carroll, Khessina, and McKendrick 2010; Douglas and

[5] While we do not explicitly invoke network terminology to describe our theory—in part because we do not use network measures to test it—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 employed 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, while the edges between them might represent varying degrees of feature overlap or similarity.

Isherwood 1996).[6] They also organize the relevant comparison set from which consumers select and evaluate products.

This argument is distinctive in several important ways. First, we highlight the consequentiality of the implicit relationships formed between products, rather than producers, consumers, or category labels. Audience evaluations of products are shaped not only by the characteristics of producers and consumers, or social influence pressures, but also by a product's position within a broader ecosystem of cultural production and consumption. The intuition behind this argument is relatively straightforward: while the choices consumers make are shaped by their individual preferences, relationships, and various other factors, they are also influenced by the feature-based similarity space within which products are embedded (Kovács and Hannan 2015). Put another way, a consumer's direct and indirect exposure to some set of related products plays a critical role in shaping his or her 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 across separate planes,

[6] Another familiar metaphor that approximates this idea is that of the cultural "milieu" or "fabric." This concept encompasses the population of cultural products that producers and/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 then be connected to one another, however distantly, based on their shared feature sets. While the theoretical and empirical implications of this idea extend beyond the scope of this paper, the imagery of a cultural fabric may help motivate our rationale for extending the concept of networks to cultural products and their constitutive features.

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 (e.g., Zuckerman 1999), but recent work at the intersection of culture, cognition, and strategy identifies the distinctive role of "product concepts," which form loose relational structures that shape consumer cognition beyond purely categorical classification (Kahl 2015). These insights reinforce our interest in feature-based associations, and suggest 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, but the study of product features and the associations they generate provides 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 (Podolny, Stuart, and Hannan 1996; Barroso et al. 2014), where the presence of too many competitors can saturate a consumer or product space (e.g., niche), making it increasingly difficult for new entrants to survive. Research across a number of empirical contexts suggests that the ability to differentiate oneself and develop a distinctive identity can help products, organizations, and other entities compete within or across niches (Hannan and Freeman 1977; Hsu and Hannan 2005; Swaminathan and Delacroix 1991).

Alternatively, some 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 or not they recognize having been previously exposed. In music, this means that the more a song sounds like something the listener has 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, but not so much as to make themselves unrelatable (Kaufman 2004). Research on consumer behavior suggests that audiences simultaneously pursue these competing goals as well, conforming on certain identity-signaling attributes (such as a product's brand or category) while distinguishing themselves on other product features (such as color or instrumentation; see Chan, Berger, and Van Boven 2012). This tension between differentiation and conformity is central to our understanding of social identities (Brewer 1991), category spanning (Zuckerman 1999; Hsu 2006), 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, we expect that songs able to strike a balance between "being recognizable" and "being different"—those that best manage the 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 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, which accounts for how features position a song relative to its musical neighbors. Controlling 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, while those 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, but 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 "rules" based on melody, harmony, and rhythm; listeners' tastes, on the other hand, do not have such concrete bounds. Although Salganik and colleagues (2006) showed that consumer choice in an artificial music market is driven both by social influence *and* a song's inherent quality, their measure of "quality" is simply audience preference in the absence of experimenter manipulation. Measuring quality "objectively" 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 employs 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 extremely difficult to apply and automate such complex codes at scale.

Fortunately, these difficulties have been partially attenuated by the application of digital data sources and new computational methods to the study of culture. Developed first by computer scientists and then adopted by mainstream social science, these technologies have begun to filter into the toolkits of cultural sociologists (Bail 2014), who have traditionally been criticized for being “methodologically impoverished” (DiMaggio, Nag, and Blei 2013). 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), fields that have developed new methods to reduce the high dimensionality of musical compositions to a set of discrete features, much like topic modeling has done for the study of large texts (Blei, Ng, and Jordan 2003). These developments have generated new research possibilities that were previously considered impractical. Using a novel dataset that includes discrete representations of musical features in the form of sonic features (a song’s “acoustic footprint”), we 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 have 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, providing 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, while 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*, making the charts an appropriate proxy for studying consumer evaluation and product performance in the field of popular music.[7] Furthermore, although the algorithm used to create the charts has evolved over the years—something we account for in our analysis—they remain the industry standard.[8] As such, they

---

[7] We recognize that by focusing on songs that appear in the Hot 100 our analysis may suffer from considerable selection bias. Nevertheless, 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? Further, it 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). While other factors such as artist popularity and marketing support play an important role in driving certain songs into the Hot 100, 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 both personal biases and external influence. In November 1991, *Billboard* replaced the 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]).

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), and are noted for their reliability as indicators of popular taste (e.g., Eastman and Pettijohn II 2014).

Dependent 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, leaving 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 both peak popularity and longevity as key variables in the measurement of cultural diffusion; we adapt them here as our dependent 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 are likely to remain on the charts longer, $R \approx .72$), we test both measures 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 also include a set of models that employ a third measure of success based on week-to-week

---

We run a series of supplementary analyses to test how this development influences our results, and find that the effect of SoundScan falls within the range of expected historical variation across our dataset (results available upon request).

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. While a third dependent variable complicates our analysis, 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 when compared to those at the top.[9]

Genre Data. The *Billboard* data require augmentation to capture more fully the multifaceted social and compositional elements of songs and artists. Although genre categories evolve and are potentially contentious (Lena and Pachucki 2013), they provide an important form of symbolic classification that organizes 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 (e.g., Peterson 1997; Phillips and Kim 2008), which in turn help to determine the listeners who seek out and are exposed to new music. Audiences consequently reinforce artist identities and genre structures (Negus 1992; Frith 1996), setting expectations for both producers and their products.

---

[9] While we could have included lagged chart position in our models to control for past performance instead of using change in position as our dependent variable, the resulting standard errors when running fixed effects analyses with a lagged DV are problematic (Nickell 1981). Findings are typically more robust when using change scores as a dependent variable in models with fixed effects (Morgan and Winship 2007).

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 sub-genre) attributions for each release (i.e., single, album, EP or LP). Although up to three genre and six style attributions are possible, 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 those songs that were not released as singles, we use the primary genre attributed to the album on which the song appears.[10] Based on these data, our sample covers fifteen genre categories—including Pop, Rock, Blues, Electronic, Jazz, and Hip Hop.[11]

---

[10] 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.

[11] 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. While 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, though genres appear and disappear over the course of our data, and those that persist have evolved, such changes have their provenance predominantly at the sub-genre or "style" level (e.g., "Hard Rock" and "Roots Rock" versus "Rock"). Employing primary, song-level genre assignments means that misattributions are unlikely or should be relegated to fringe cases.

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 necessitated 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 over 30 million songs and 3 million artists. Their data contains 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% of the songs (25,102 of 26,846 total songs) that appeared on the charts between 1958 and 2016, including 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. There are of course limitations associated with distilling complex cultural products into a handful of discrete features, but we believe that these features represent the best available approximation of what people hear when they listen to music. Nearly twenty 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). Our conversations 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.

Independent 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 construct 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 0-1 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.[12] For each song *i*, we pulled every other song that appeared on the charts during the year prior to song *i*'s debut, and 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.

---

[12] Each of these features is weighted equally to calculate our pairwise cosine similarity measure. While we considered prioritizing certain features over others (e.g., weighting "tempo" more heavily than "mode"), conversations with musicologists and computer scientists specializing in MIR provided no consistent rationale for using weights. Moreover, the eleven 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.

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 which "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 that 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 omnivorous consumption behavior is on the rise (e.g., 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" to be similar when seen as a vector of features, but perceived to "sound" quite different by listeners). Thus, we weight each song-pair's raw cosine similarity by the average similarity of those songs' parent 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 multi-dimensional 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 also recognize that for many listeners these genres and their boundaries are not absolute, particularly when it comes to the most mainstream music being 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 our *genre-weighted typicality (yearly)* measure, we calculated yearly within-genre averages for each sonic feature, and then again used a cosine similarity algorithm to measure the average proximity of each pair of genres in feature space. The resulting weights were then applied to the raw similarity measure summarized above for each song pair. For example, if one rock song and one folk song had a raw similarity of 0.75, and the average similarity between "rock" and "folk" in year *x* is 0.8, then that genre-weighted similarity between those two songs would be 0.75 * 0.8 = 0.6.[13] If both songs were categorized as "rock," then the weight would equal 1, and the genre-weighted similarity between songs would be 0.75. 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

[13] 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 the minimum weight (i.e., the largest distance between two genres' vectors of average features) is used as the weighting multiplier for that song's cosine similarity with every other song on the charts during the previous year.

high degree of similarity between songs across our dataset and in popular music more generally ($\mu = 0.81$; $\sigma = 0.06$; Range = 0.26–0.92; see Figure 3.1).[14]

---Insert Figure 3.1 here---

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 that a song appears on the charts. To do this, we first measured the cosine similarity between each song and the other songs with which it shared a chart. For each week, we created a matrix **A**, that has dimensions matching the number of songs on each week's charts (100x100),

[14] 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 week. 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 3-6 (results available upon request).

with cell $\boldsymbol{A}_{ijt}$ 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 those other songs with which it is directly competing on the charts.

Control 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 (e.g., 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 number 1 on the charts.

Second, we included a set of dummy variables in each of our models 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 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 fourth through tenth song on the charts, or (4) is she has had more than ten songs in the Hot 100. These dummies also help to 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 dummy variable called *multiple memberships* to account for artists who released songs under different names or band formations. For example, Annie Lennox appears on the charts both as a member of the Eurythmics and as a solo artist. As the Eurythmics represent Lennox's first appearance on the charts, every subsequent appearance of hers as a solo artist was coded as a 1 for *multiple memberships*. This was done for every artist who appeared with multiple bands (or with a band and as a solo artist) on the Hot 100 (roughly 6% of our dataset). For these artists, song counts were continued from previous chart incarnations—meaning that 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 to account for any potential benefits chart veterans receive.

Fourth, we included a dummy variable called *long song*, set to 1 if a song was unusually long. Historical recording formats, along with radio, have 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 impact 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. 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 any song with only a single genre classification. Using this method, roughly 24% 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 construct a dummy variable *reissue* for any song that was re-released and appeared on the charts for a second time. For example, Prince's track "1999" originally charted in 1982, reaching #12 and staying on the charts for 27 weeks. It was reissued around 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, are likely to 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 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. This was done 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 (see footnote 8). Employing 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 3.1 summarizes descriptive statistics and correlations for all the key variables in our analysis.

---Insert Table 3.1 around here---

Estimation Strategy

To demonstrate 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 & 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 the sonic features we collected 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 comprised of discrete whole numbers derived from ranks, we run a second set of models using an ordered logit specification (Models 3 & 4). We include various artist-level control variables (e.g., previous song and multiple band membership dummies) , as ordered logit models with fixed effects for artists can have inconsistent estimators (see Baetschmann, Staub, and Winkelmann 2015). Models estimating *weeks on charts* contain the same control variables, but use a truncated negative binomial specification, as the outcome is a count variable with a minimum value of one (Models 5 & 6).

The models described above reflect cross-sectional analyses that use 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 commensurately, 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' jockeying for position will be influenced by their sonic similarity to the weekly competition on the Hot 100 charts.. Thus, our final set of models estimates 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 & 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 3–6. All independent and control variables are lagged one week—both to match the one-week window used by *Billboard,* and because the constant churn within the charts would render longer lags substantively meaningless. These and all other models presented in the paper employ robust standard errors.

Results

Before presenting our main results, we first wanted to explore the historical relationship between song typicality and chart position. A descriptive analysis of this relationship is presented in Figures 3.2a and 3.2b, and indicates 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 over time (a) those songs that reached the top 40 with those that did not, and (b) those songs that reached number one with those that did not. Figure 3.2a indicates that the songs that peaked in the Top 40 are comparably typical to songs that failed to reach Top 40 status. In fact, in the early years of the charts, top 40 songs are slightly *more* typical than the songs that peaked in positions 41–100. Conversely, Figure 3.2b indicates that, aside from a few punctuated 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, although 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.

---Insert Figures 3.2a and b here---

It is also worth noting the general trends of song typicality across our dataset: the chart's early history was marked by more homogenous, "typical" songs, while more atypical songs tended to appear in the 1970s, '80s, and '90s. This trend toward greater atypicality has reversed itself in recent years, as songs appearing on the charts after 2000 seem to be growing more typical. While these trends tell us something interesting about *absolute* levels of feature-based typicality over time, the models that follow allow us to measure how a song's typicality *relative to its contemporaneous competition* affects its performance on the charts.

Results from our first pair of formal models are depicted in Figure 3.3, which graphically presents 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) are positively associated with peak chart position, while energy (intensity/noise), speechiness, and acousticness produce negative coefficients. Although we do not have space to theorize the interpretation of these individual results, they provide some face validity that product features matter for songs' chart performance.

---Insert Figure 3.3 here---

In addition to providing controls 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' often fail to produce a second song or album as popular as their first. Our results suggest that an artist's second and third "hit" songs do not perform as well as their first. The positive coefficient for songs released by artists with more than

[15] All control variables—including genre affiliation, dummies for each musical key (C through B), major label dummy, *long song*, *multiple memberships, crossover, reissue,* and half-decade time dummies—are included in these models, but are not shown in the figure.

10 previous hits similarly provides support for the “superstar” effects we anticipated, suggesting that artists receive additional advantage after they have achieved substantial popular 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 suggest a similar pattern of relationships, although one difference is worth noting: while we again 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 any subsequent hits will be more likely to experience shorter chart lives. Audiences may more quickly grow tired of music released by artists they already know.

Results from models 1 and 2 are instructive and provide *prima facie* evidence that sonic features are meaningfully correlated with chart performance. Nevertheless, as they appear independently in these models, the results reveal little about how bundles of features—i.e., songs—are similar to or different from each other *en masse*, 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.

Given the relatively high levels of typicality across our dataset, we argue that, while adhering to existing songwriting prescriptions is likely to help a song achieve some degree of success, those songs that effectively separate themselves from the pack tend to exhibit differentiation or novelty on one or more sonic features. For example, The Beatles’ 1969 hit “Come Together” reached the top of the charts on November 29, 1969, and featured a typicality score of 0.66 the week it debuted—over two standard deviations less typical than the average song released that year. When we dig a bit deeper 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 example does not statistically represent our entire dataset, it does speak to some of the factors that drive our typicality measure.

Table 3.2 summarizes the coefficients for our key independent and control variables from models 3–6. Recall that model 3 employs ordered logit regression to predict a song's peak position using its typicality relative to other songs that appeared on the charts in the previous 52 weeks. Results suggest a significant negative relationship between song typicality and peak position: 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 our hypothesized inverted U-shaped relationship between typicality and chart performance. Results support our 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.

---Insert Table 3.2 here---

Because second order terms in ordered logit models are difficult to interpret (Karaca-Mandic, Norton, and Dowd 2012), we created Figure 3.4 to visualize the marginal effects of songs' typicality on their peak chart position. For purposes of clarity and interpretability we partitioned peak chart position into meaningful "tranches" that are represented by the different lines in the figure. We then use 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 across the entirety of our dataset, but 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 distinct.

---Insert Figure 3.4 here---

Models 5 and 6 employ truncated negative binomial s 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, but when we include the squared term (model 6), we once more find an inverted-U-shaped relationship. That is, for the most novel songs in our dataset, higher levels of typicality would increase their odds of remaining on the charts, while 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 (0.75 vs. 0.81) 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). As the average song only lasts on the charts for about 11 weeks, half a week represents a substantially better outcome.

Across models 3–6, we find that songs are more likely to attract and maintain the attention of consumers if they are differentiated from other songs on the charts, but not so much that they fail to meet prevailing expectations. We also find consistent 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. 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 long, and spent a month at number 1; The Beatles' "Hey Jude" clocks in at 7:11 and spent 9 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 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 do as well as their chart debuts, while 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. These "superstar" effects on peak performance are further supported by the positive coefficient on *multiple memberships*, which suggests that veteran musicians who have 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 is surely benefitting those 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, where the chart's first several years (1958–1961) comprise the omitted reference group. Recall that the five years before (1987–1991) and after (1992–1996) the introduction of SoundScan are of particular interest (note that 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 the introduction of SoundScan, whereas 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 remain 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 explain these

---

[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 are benefitted by being optimally differentiated—but not for crossovers. However, crossovers do comprise a higher proportion of #1 songs (30%) than their overall chart presence would suggest (24% of all songs are crossovers by our measure).

week-to-week dynamics, we turn to the results of our fixed effects models, presented in Table 3.3.

---Insert Table 3.4 here---

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* 0.6 positions each week—which is substantial given the relatively low debut position of most songs on the charts (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.

Finally, in model 8 we add a quadratic term for song typicality and find that, all else equal, more typical songs tend to fare worse on subsequent weeks' charts than those 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. More practically, this means that 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 other, 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

These results provide evidence that the features of cultural products affect consumption behavior, both independently and in the way they structure how audiences compare and evaluate products (cf., de Vaan et al. 2015; Lena 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 content 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.

While we believe that these findings provide important insights into the consumption dynamics of a multi-billion-dollar industry, we also recognize several important limitations. Although the data we use to measure sonic features is relatively comprehensive and sophisticated, it represents a substantial distillation of a song's musical complexity. Reducing such a high dimensional object into eleven fixed features inevitably simplifies its cultural fingerprint and alters its relationships with other like-objects. As MIR tools improve, so too will our ability to map the connections between songs. Our data also does not allow us to account for listeners' idiosyncratic interpretations of features or lyric similarity between songs. Moreover, the bounded nature of the Hot 100—which includes only those 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. 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 encounter are likely to extend beyond our sample in the music industry and to the other creative industries as well.

The analyses in this paper allow us to explain why, 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 dynamic analyses to better 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 in this paper 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 in this paper do not account for the external consumption environment, 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 those 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 are likely to 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 study of the concepts and measures developed in this paper can be generative in a variety of empirical contexts, and serve 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.

Conclusion

Without denying the important role of social dynamics, we remain convinced of the influence product features have on popular success. Both independently and in concert, the content of cultural products needs to be considered more seriously when investigating success in cultural markets. We have demonstrated 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 that this paper, 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.

To that end, we believe that the ideas presented in this paper make several 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. While these tools necessarily simplify the intrinsic high-dimensionality of culture, they also empower us to generate new insights in historically opaque contexts. Although many new cultural 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 (Lena 2006; Jones et al. 2012), but there has been no systematic attempt to theorize and measure how product features influence the emergence and diffusion of consumption patterns. In this paper, we introduce and exploit a rich dataset capable of exploring these dynamics, generating new insights into the world of popular music and cultural markets more broadly.

Second, we measure and test the effects of product features and the associations they generate among audiences. Our conception of feature-similarity space can serve both as a tool to map ecosystems of cultural products, and as a means to understand selection dynamics in markets that require subjective evaluation. We argue that the system of associations between products is theoretically and analytically distinct from—though integrally connected to and mediated through—networks of producers and consumers. In so doing, we raise the possibility that cultural content asserts its own autonomous influence over evaluation outcomes through product crowding and differentiation. This conceptualization of culture is dynamic and will we hope 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 the way in which 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 (e.g., 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 are likely to 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. The most successful producers, to paraphrase a well-known saying, will be aiming 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). While research in this area has explored the origins and consequences of categorical classification on firms and products, our results suggest a more grounded approach may be necessary to fully understand how markets are structured. Combinations of features likely play an integral role in the way products, organizations, and even individuals are perceived and evaluated. In our analysis, we include both product features (sonic attributes) and category labels (genres) to ensure that a computer-driven reduction in complexity did not cause inappropriate interpretation. In future work, we intend to dive even deeper into the interrelationship between features and labels. For example, how do product features help create the categorical structure of musicological genres? To draw an analogy, while research has 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 (Rao, Monin, and Durand 2003; Kovács and Johnson 2013), integrating these perspectives to explore the relationship between ingredients and the way that 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), and artists have more tools and data at their disposal than ever before, providing them with incredibly detailed information about the elements of popular songs, which might in turn help them to craft their own hits (Thompson 2014). Nevertheless, while 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 importantly, the characteristics of contemporaneous songs will have a significant impact on that song's success. Context matters. 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.

Tables

Table 1.1: Spheres of Influence

Type	**Tie Defined By**	**Example**
Collaborative	Direct interpersonal contact	Co-writing a song; playing in a band together
Cultural	Cultural similarity or category membership	Exposure to other artists via shared genre affiliation
Organizational	Shared organizational affiliation	Exposure to other artists via shared record label
Geographic	Physical proximity or co-location	Exposure to other artists via shared city or country

Table 1.2: Operationalization of Genre Diversity and Alter Creativity Variables

Sphere	Genre Diversity	Alter Creativity
Collaborative	Scaled count of unique genres across an artist's collaboration network	Average song novelty across an artist's collaboration network
Cultural	Scaled count of unique genres with which an artist is affiliated	Average song novelty across an artist's home genre(s)
Organizational	Scaled count of unique genres across an artist's record label	Average song novelty across an artist's record label
Geographic	Scaled count of unique genres across an artist's home city (or country)	Average song novelty across an artist's home city (or country)

Table 1.3: Descriptive Statistics and Pearson Correlations for Variables of Interest

Variable	Mean	Std. Dev.	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]
[1] Song Novelty	20.5	5.57	1									
[2] Female	0.09	0.29	0.03	1								
[3] Male	0.58	0.49	-0.07	-0.38	1							
[4] US	0.43	0.5	0.03	0.03	-0.07	1						
[5] Classical	0.04	0.19	0.14	-0.02	-0.15	-0.09	1					
[6] Jazz	0.12	0.32	0.07	0.09	-0.22	0.15	0	1				
[7] Past Creativity	30	7.77	0.17	0	-0.02	-0.03	0.15	0.02	1			
[8] Artist Productivity	35.9	61.3	0.06	0.03	-0.2	0.03	0.22	0.07	0.12	1		
[9] Artist Tenure	2,565	3,580	0.02	0.06	-0.3	0.08	0.11	0.17	0.08	0.57	1	
[10] Charting Artist	0.21	0.4	-0.07	0.11	-0.13	0.27	-0.06	-0.01	-0.08	0.16	0.23	1
[11] Major Label	0.2	0.4	-0.04	0.09	-0.14	-0.01	0.07	0.01	-0.03	0.11	0.1	0.22
[12] # of Alters (Collaborative Sphere)	4.4	9.22	0.04	-0.06	0.02	0.01	0.24	0.07	0.09	0.61	0.28	0.13
[13] Genre Diversity (Collaborative Sphere)	1.35	3.43	0.05	0	-0.1	0.04	0.18	0.1	0.08	0.52	0.28	0.18
[14] Alter Creativity (Collaborative Sphere)	20.23	2.2	0.21	0.02	-0.06	0	0.08	0.02	0.03	-0.03	-0.07	-0.09
[15] # of Alters (Cultural Sphere)	1,771	1,830	0.05	-0.03	0.16	0.05	-0.11	0	0.03	0.03	0.06	-0.01
[16] Genre Diversity (Cultural Sphere)	2.3	2.05	0	-0.06	0.18	0.06	-0.04	0	0.02	-0.02	-0.06	0.04
[17] Alter Creativity (Cultural Sphere)	19.22	2.69	0.43	0	-0.09	0.02	0.19	0.1	0.05	0.03	-0.06	0.12
[18] # of Alters (Organizational Sphere)	58.51	110.9	-0.04	0.07	-0.09	0.06	-0.01	0.03	-0.01	0.09	0.12	0.2
[19] Genre Diversity (Organizational Sphere)	35.24	47.19	-0.06	0.07	-0.07	0.06	-0.02	0	-0.01	0.08	0.12	0.21
[20] Alter Creativity (Organizational Sphere)	19.62	2.3	0.39	0	-0.05	0.03	0.11	0.06	0.06	0	-0.1	-0.02
[21] # of Alters (Geographic Sphere)	123	221.43	0.09	-0.03	0.08	0.14	0	0.02	0.06	-0.01	-0.01	-0.06
[22] Genre Diversity (Geographic Sphere)	61.36	78.81	0.08	-0.03	0.1	0.1	-0.03	0.02	0.06	-0.03	-0.01	-0.08
[23] Alter Creativity (Geographic Sphere)	19.51	2.58	0.29	0.02	-0.05	0.01	0.01	0.03	-0.02	0.02	-0.05	0.04

[11]	[12]	[13]	[14]	[15]	[16]	[17]	[18]	[19]	[20]	[21]	[22]	[23]
1												
0.1	1											
0.1	0.7	1										
-0.05	-0.09	-0.09	1									
-0.04	0.04	0.04	-0.02	1								
-0.05	0.05	0.07	-0.02	0.45	1							
0	0.03	0.05	0.43	-0.06	-0.04	1						
0.56	0.06	0.07	-0.08	0.07	0.07	-0.8	1					
0.58	0.07	0.09	-0.1	0.08	0.09	-0.1	0.96	1				
-0.07	0.01	0.02	0.4	-0.01	-0.02	0.78	-0.16	-0.18	1			
-0.07	0.06	0.04	0.03	0.17	-0.02	0.06	0	-0.1	0.06	1		
-0.08	-0.01	0.03	0	0.16	0	-0.01	0	0	0.01	0.8	1	
0.01	0	0.01	0.36	0.03	0	0.71	-0.05	-0.06	0.59	-0.01	-0.09	1

Table 1.4: OLS Regressions Predicting Song Novelty

Song Novelty

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
1960s	4.722***	4.551***	2.843***	3.791***	4.179***	4.018***
	(0.111)	(0.111)	(0.112)	(0.113)	(0.113)	(0.116)
1970s	5.478***	4.972***	0.162***	2.382***	3.884***	0.703***
	(0.040)	(0.041)	(0.049)	(0.045)	(0.044)	(0.053)
1980s	-0.652***	-1.028***	-3.785***	-2.352***	-1.586***	-3.687***
	(0.025)	(0.025)	(0.028)	(0.026)	(0.026)	(0.031)
1990s	-5.039***	-4.812***	-3.389***	-4.016***	-4.311***	-3.519***
	(0.015)	(0.015)	(0.016)	(0.015)	(0.016)	(0.018)
Female	0.277***	0.268***	0.387***	0.314***	0.252***	0.415***
	(0.025)	(0.025)	(0.024)	(0.025)	(0.025)	(0.024)
Male	-0.284***	-0.230***	-0.297***	-0.289***	-0.344***	-0.290***
	(0.015)	(0.016)	(0.015)	(0.015)	(0.015)	(0.015)
US	0.064***	0.070***	0.155***	0.062***	0.033*	0.096***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Classical	3.442***	3.152***	1.318***	2.570***	3.398***	0.876***
	(0.046)	(0.046)	(0.046)	(0.045)	(0.046)	(0.049)
Jazz	0.680***	0,642***	0.280***	0.541***	0.681***	0.183***
	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Past Creativity	0.113***	0.109***	0.087***	0.092***	0.107***	0.080***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Artist Productivity	0.001***	0.000	0.000	0.001*	0.001**	-0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Artist Tenure	-0.001***	-0.001***	-0.001***	-0.001***	-0.001***	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Artist Tenure (squared	0.0001***	0.0001***	0.0001***	0.0001***	0.0001***	0.0001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Charting Artist	-0.783***	-0.684***	-0.438***	-0.544***	-0.658***	-0.382***
	(0.017)	(0.017)	(0.016)	(0.016)	(0.016)	(0.016)
Major Label	-0.384***	-0.324***	-0.124***	-0.274***	-0.309***	-0.136***
	(0.016)	(0.016)	(0.015)	(0.019)	(0.016)	(0.019)
# of Alters		-0.001				-0.001
(Collaborative Sphere)		(0.001)				(0.001)
Genre Diversity		0.053***				0.010***
(Collaborative Sphere)		(0.003)				(0.003)
Alter Creativity		0.270***				0.035***
(Collaborative Sphere)		(0.003)				(0.003)

# of Alters			-0.001***			0.001*
(Cultural Sphere)			(0.000)			(0.000)
Genre Diversity			0.047***			0.030***
(Cultural Sphere)			(0.003)			(0.003)
Alter Creativity			0.797***			0.791***
(Cultural Sphere)			(0.004)			(0.007)
# of Alters				-0.001***		0.000
(Organizational Sphere)				(0.000)		(0.000)
Genre Diversity				0.006***		0.001
(Organizational Sphere)				(0.000)		(0.000)
Alter Creativity				0.527***		0.220***
(Organizational Sphere)				(0.003)		(0.004)
# of Alters					0.001***	-0.001***
(Geographic Sphere)					(0.000)	(0.000)
Genre Diversity					0.003***	0.002***
(Geographic Sphere)					(0.000)	(0.000)
Alter Creativity					0.369***	0.005
(Geographic Sphere)					(0.003)	(0.004)
Constant	17.998***	12.546***	3.447***	8.129***	10.737***	2.701***
	(0.036)	(0.072)	(0.070)	(0.064)	(0.067)	(0.095)
Observations	613,361	613,361	613,361	613,361	613,361	613,353
R-squared	0.200	0.233	0.292	0.271	0.246	0.300

Note: All Genre Diversity and Alter Creativity measures use imputed values. Reference groups for Female and Male = Group and for Decade = 1950s. Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table 1.5: OLS Regressions with Interactions

Song Novelty

	Model 7	Model 8	Model 9	Model 10	Model 11
# of Alters	-0.001	-0.001	-0.001	-0.001	0.002
(Collaborative Sphere)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Genre Diversity	0.130***	0.010***	0.010***	0.010***	0.064***
(Collaborative Sphere)	(0.017)	(0.003)	(0.003)	(0.003)	(0.017)
Alter Creativity	0.047***	0.036***	0.036***	0.035***	0.028
(Collaborative Sphere)	(0.004)	(0.003)	(0.003)	(0.003)	(0.019)
# of Alters	0.001*	0.001*	0.001*	0.001**	0.001*
(Cultural Sphere)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.030***	0.189***	0.029***	0.029***	0.170***
(Cultural Sphere)	(0.003)	(0.023)	(0.003)	(0.003)	(0.023)
Alter Creativity	0.792***	0.804***	0.793***	0.794***	1.010***
(Cultural Sphere)	(0.007)	(0.008)	(0.007)	(0.007)	(0.044)
# of Alters	0.000	0.000	0.000	0.000	0.000
(Organizational Sphere)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.001*	0.001*	0.007***	0.001	0.005***
(Organizational Sphere)	(0.000)	(0.000)	(0.001)	(0.000)	(0.001)
Alter Creativity	0.221***	0.221***	0.226***	0.221***	-0.122***
(Organizational Sphere)	(0.004)	(0.004)	(0.004)	(0.004)	(0.035)
# of Alters	-0.001***	-0.001***	-0.001***	-0.001***	-0.001
(Geographic Sphere)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.002***	0.002**	0.002***	0.007***	0.006***
(Geographic Sphere)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Alter Creativity	0.004	0.006	0.005	0.011**	0.139***
(Geographic Sphere)	(0.004)	(0.004)	(0.004)	(0.004)	(0.030)
Genre Diversity X Alter Creativity	-0.006***				-0.005***
(Collaborative Sphere)	(0.001)				(0.001)
Genre Diversity X Alter Creativity		-0.008***			-0.008***
(Cultural Sphere)		(0.001)			(0.001)
Genre Diversity X Alter Creativity			-0.001***		-0.001***
(Organizational Sphere)			(0.000)		(0.000)
Genre Diversity X Alter Creativity				-0.001***	-0.001***
(Geographic Sphere)				(0.000)	(0.000)
Constant	2.399***	2.302***	2.464***	2.424***	2.121***
	(0.102)	(0.111)	(0.105)	(0.108)	(0.409)
Observations	613,353	613,353	613,353	613,353	613,353
R-squared	0.300	0.300	0.300	0.300	0.300

Note: All Genre Diversity and Alter Creativity measures use imputed values. Reference groups for Female and Male = Group and for Decade = 1950s. All models in this table also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared). Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table A1.1: Robustness Checks

	Song Novelty			
	Full sample	Without Zeroes	First Songs Only	Max Alter Creativity
# of Alters	-0.002*	-0.001	0.030***	-0.001
(Collaborative Sphere)	(0.001)	(0.002)	(0.004)	(0.001)
Genre Diversity	0.000	0.010	-0.016	0.009***
(Collaborative Sphere)	(0.003)	(0.005)	(0.009)	(0.003)
Alter Creativity	0.014***	0.049***	0.045***	0.012***
(Collaborative Sphere)	(0.001)	(0.012)	(0.007)	(0.001)
# of Alters	-0.001***	0.001***	0.000	-0.001***
(Cultural Sphere)	(0.000)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.025***	-0.027*	0.017**	0.064***
(Cultural Sphere)	(0.003)	(0.013)	(0.006)	(0.003)
Alter Creativity	0.632***	0.424***	0.789***	0.428***
(Cultural Sphere)	(0.010)	(0.031)	(0.012)	(0.006)
# of Alters	0.003***	0.002***	0.001**	0.002***
(Organizational Sphere)	(0.000)	(0.001)	(0.000)	(0.000)
Genre Diversity	-0.007***	-0.002	-0.002**	-0.007***
(Organizational Sphere)	(0.001)	(0.001)	(0.001)	(0.001)
Alter Creativity	0.025***	0.255***	0.221***	0.016***
(Organizational Sphere)	(0.001)	(0.017)	(0.006)	(0.001)
# of Alters	0.001	-0.001	-0.001*	-0.001*
(Geographic Sphere)	(0.000)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.001***	0.004***	0.001***	0.002***
(Geographic Sphere)	(0.000)	(0.001)	(0.000)	(0.000)
Alter Creativity	0.000	0.073***	0.015*	0.003***
(Geographic Sphere)	(0.000)	(0.017)	(0.007)	(0.001)
Constant	4.527***	1.778***	3.358***	3.995***
	(0.085)	(0.344)	(0.192)	(0.086)
Observations	607,787	57,511	191,555	613,353
R-squared	0.287	0.286	0.289	0.288

Note: All Genre Diversity and Alter Creativity measures use imputed values, other than the first two model. All models also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared).
Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table A1.2: OLS Regression with Fixed Effects

Song Novelty

	Year Fixed Effects	Artist Fixed Effects	Year & Artist Fixed Effects
# of Alters	-0.003**	0.015***	0.006*
(Collaborative Sphere)	(0.001)	(0.003)	(0.003)
Genre Diversity	0.019***	-0.008	0.002
(Collaborative Sphere)	(0.002)	(0.005)	(0.005)
Alter Creativity	0.045***	0.000	0.023***
(Collaborative Sphere)	(0.003)	(0.005)	(0.005)
# of Alters	-0.001***	-0.001	-0.001***
(Cultural Sphere)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.034***	NA	NA
(Cultural Sphere)	(0.003)	(NA)	(NA)
Alter Creativity	0.811***	0.569***	0.261***
(Cultural Sphere)	(0.007)	(0.013)	(0.023)
# of Alters	0.001*	0.001***	0.001**
(Organizational Sphere)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.000	-0.002**	-0.002**
(Organizational Sphere)	(0.000)	(0.001)	(0.001)
Alter Creativity	0.229***	0.055***	0.050***
(Organizational Sphere)	(0.004)	(0.005)	(0.005)
# of Alters	-0.001	-0.001*	-0.001
(Geographic Sphere)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.001***	0.003***	0.000
(Geographic Sphere)	(0.000)	(0.000)	(0.000)
Alter Creativity	0.011**	0.043***	0.038***
(Geographic Sphere)	(0.004)	(0.006)	(0.006)
Year FE	YES	NO	YES
Artist FE	NO	YES	YES
Observations	620,174	620,174	620,174
R-squared	0.315	0.395	0.409

Note: All Genre Diversity and Alter Creativity measures use imputed values.
Robust standard errors reported in parentheses.
*p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table A1.3: Alternative Specifications for Song Novelty

	Song Novelty		
	No Window	Five-Year Window	Two-Year Window
# of Alters	0.003	0.004***	0.009***
(Collaborative Sphere)	(0.002)	(0.001)	(0.000)
Genre Diversity	0.008*	0.007***	0.009***
(Collaborative Sphere)	(0.004)	(0.002)	(0.001)
Alter Creativity	0.043***	0.012***	-0.001
(Collaborative Sphere)	(0.005)	(0.002)	(0.001)
# of Alters	-0.001***	-0.001***	0.001***
(Cultural Sphere)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.085***	0.034***	0.019***
(Cultural Sphere)	(0.005)	(0.002)	(0.001)
Alter Creativity	0.898***	0.428***	0.210***
(Cultural Sphere)	(0.011)	(0.005)	(0.002)
# of Alters	0.001*	-0.001***	-0.001***
(Organizational Sphere)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.000	0.002***	0.003***
(Organizational Sphere)	(0.000)	(0.000)	(0.000)
Alter Creativity	0.306***	0.134***	0.063***
(Organizational Sphere)	(0.006)	(0.003)	(0.001)
# of Alters	-0.001	-0.001***	-0.001***
(Geographic Sphere)	(0.000)	(0.000)	(0.000)
Genre Diversity	0.002***	0.002***	0.001***
(Geographic Sphere)	(0.000)	(0.000)	(0.000)
Alter Creativity	-0.037***	-0.024***	-0.013***
(Geographic Sphere)	(0.005)	(0.003)	(0.001)
Constant	21.192***	8.115***	5.097***
	(0.141)	(0.065)	(0.033)
Observations	613,353	613,353	613,353
R-squared	0.122	0.366	0.382

Note: All Genre Diversity and Alter Creativity measures use imputed values. All models in this table also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared).
Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table 2.1: Descriptive Statistics (By Gender)

	Female Artists (N = 56,911)				Male Artists (N = 198,244)			
Variable	Mean	Std. Dev.	Minimum	Maximum	Mean	Std. Dev.	Minimum	Maximum
[1] Song Novelty	21.58	6.09	11.39	54.26	21.58	6.07	11.47	61.23
[2] US	0.48	0.5	0	1	0.47	0.5	0	1
[3] Classical	0.03	0.16	0	1	0.08	0.28	0	1
[4] Jazz	0.21	0.41	0	1	0.2	0.4	0	1
[5] Past Creativity	30.12	7.19	0	68.6	30.41	7.23	0	76.76
[6] Artist Productivity	18.7	28.72	0	268	23.82	39.71	0	592
[7] Artist Tenure	3214.6	4072.24	1	16601	4094.2	4432.01	1	16687
[8] Charting Artist	0.34	0.48	0	1	0.25	0.43	0	1
[9] Major Label	0.31	0.46	0	1	0.26	0.44	0	1
[10] # of Genres Spanned	2.12	1.36	1	12	2.13	1.37	1	14
[11] Network Size	8.14	14.3	1	97	12.72	28.58	1	257
[12] Network Composition	0.52	0.35	0.03	1	0.07	0.12	0	0.75
[13] Genre Composition	0.34	0.11	0.03	1	0.23	0.11	0	0.67

Table 2.2: Pearson Correlations

Variable	[0]	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]	[12]	[13]
[0] Artist Gender	1													
[1] Song Novelty	0	1												
[2] US	0.01	0	1											
[3] Classical	-0.1	0.17	-0.16	1										
[4] Jazz	0.01	0.08	0.22	-0.07	1									
[5] Past Creativity	-0.02	0.2	-0.07	0.23	0.04	1								
[6] Artist Productivity	-0.06	0.06	0.02	0.25	0.05	0.14	1							
[7] Artist Tenure	-0.08	0.01	0.11	0.09	0.15	0.08	0.52	1						
[8] Charting Artist	0.09	-0.09	0.35	-0.11	-0.03	-0.12	0.12	0.17	1					
[9] Major Label	0.06	-0.04	-0.02	0.07	-0.04	-0.03	0.1	0.04	0.16	1				
[10] # of Genres Spanned	0	-0.02	0.1	-0.02	0.07	-0.02	0.04	-0.02	0.15	-0.02	1			
[11] Network Size	-0.07	0.05	0.03	0.36	0.13	0.11	0.64	0.3	0.1	0.12	0.04	1		
[12] Network Composition	0.69	-0.01	-0.02	-0.04	-0.05	-0.01	-0.08	-0.11	0.03	0.04	-0.03	-0.08	1	
[13] Genre Composition	0.39	-0.08	-0.07	-0.19	-0.17	-0.1	-0.03	0	0.17	0.14	-0.03	-0.13	0.31	1

Table 2.3: OLS Regression Results (Main Models)

	Song Novelty (logged)					
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Female	0.000	0.309***	0.299***	0.374***	0.496***	0.533***
	(0.001)	(0.026)	(0.026)	(0.036)	(0.028)	(0.037)
1970s		-6.879***	-6.886***	-6.875***	-6.902***	-6.909***
		(0.062)	(0.062)	(0.062)	(0.062)	(0.062)
1980s		-11.153***	-11.148***	-11.148***	-11.201	-11.196***
		(0.057)	(0.057)	(0.057)	(0.057)	(0.057)
1990s		-6.664***	-6.694***	-6.658***	-6.737***	-6.772***
		(0.057)	(0.055)	(0.055)	(0.055)	(0.056)
US		0.035	0.041	0.035	-0.016	-0.012
		(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
Classical		3.441***	3.591***	3.447***	3.300***	3.463***
		(0.055)	(0.058)	(0.055)	(0.056)	(0.059)
Jazz		0.828***	0.877***	0.824***	0.744***	0.793***
		(0.029)	(0.030)	(0.029)	(0.030)	(0.030)
Past Creativity		0.155***	0.155***	0.155***	0.154***	0.154***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Artist Productivity		0.005	0.003***	0.001	0.001	0.004***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Artist Tenure		-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Artist Tenure (squared)		0.000***	0.000***	0.000***	0.000***	0.000***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Charting Artist		-0.955***	-0.941***	-0.956***	-.902***	-0.884***
		(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Major Label		-0.398***	-0.386***	-0.397***	-0.356***	-0.340***
		(0.024)	(0.024)	(0.024)	(0.025)	(0.025)
# of Genres Spanned		-0.334***	-0.339***	-0.335***	-0.356***	-0.355***
		(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
# of Genres Spanned (squared)		0.042***	0.043***	0.042***	0.043***	0.044***
		(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Network Size			-0.006***			-0.007***
			(0.001)			(0.001)
Network Composition				-0.143**		-0.084
				(0.054)		0.054
Genre Composition					-1.796***	-1.901***
					(0.104)	(0.105)
Constant	3.036***	24.254***	24.280***		24.781***	24.848***
	(0.001)	(0.082)	(0.082)		(0.088)	(0.087)
Observations	255,155	249,445	249,428	249,445	249,445	249,445
R-squared	0.00	0.23	0.24	0.24	0.24	0.25

Note: Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table 2.4: OLS Regression Results (Interactions X Female)

	Song Novelty (logged)			
	Model 7	Model 8	Model 9	Model 10
Female	0.132***	0.471***	-0.018	-0.321**
	(0.030)	(0.044)	(0.087)	(0.098)
Network Size	-0.007***			-0.009***
	(0.001)			(0.001)
Female X Network Size	0.020***			0.023**
	(0.002)			(0.002)
Network Composition		0.194		0.547***
		(0.100)		(0.101)
Female X Network Composition		-4.790***		-0.438***
		(0.121)		(0.126)
Genre Composition			-2.131***	-2.406***
			(0.113)	(0.115)
Female X Genre Composition			1.621***	1.952***
			(0.252)	(0.253)
Constant	24.354***	24.262***	24.863***	25.033***
	(0.082)	(0.082)	(0.088)	(0.088)
Observations	249,445	249,445	249,445	249,445
R-squared	0.25	0.25	0.25	0.25

Note: All models in this table also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared), and # of Genres Spanned (squared). Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table A2.1: OLS Regression (By Gender Subsample)

	Song Novelty (logged)		
	Full Sample	Female Only	Male Only
Female	0.533***		
	(0.037)		
1970s	-6.909***	-7.300***	-6.747***
	(0.062)	(0.138)	(0.069)
1980s	-11.196***	-11.824***	-10.992***
	(0.057)	(0.123)	(0.064)
1990s	-6.772***	-7.465***	-6.543***
	(0.056)	(0.119)	(0.063)
US	-0.012	-0.086	0.036
	(0.024)	0.052)	(0.027)
Classical	3.463***	3.849***	3.452***
	(0.059)	(0.185)	(0.063)
Jazz	0.793***	1.561***	0.561***
	(0.030)	(0.071)	(0.034)
Past Creativity	0.154***	0.125***	0.160***
	(0.002)	(0.004)	(0.002)
Artist Productivity	0.004***	0.009***	0.002***
	(0.000)	(0.001)	(0.000)
Artist Tenure	-0.000***	-0.000***	-0.000***
	(0.000)	(0.000)	(0.000)
Artist Tenure (squared)	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
Charting Artist	-0.884***	-1.056***	-0.889***
	(0.026)	(0.057)	(0.030)
Major Label	-0.340***	-0.403***	-0.333***
	(0.025)	(0.049)	(0.028)
# of Genres Spanned	-0.355***	-0.361***	-0.405***
	(0.023)	(0.057)	(0.026)
# of Genres Spanned (squared)	0.044***	0.048***	0.047***
	(0.003)	(0.002)	(0.004)
Network Size	-0.007***	0.000	-0.007***
	(0.001)	(0.002)	(0.001)
Network Composition	-0.084	0.177*	0.518***
	0.054	(0.076)	(0.101)
Genre Composition	-1.901***	0.307	-2.380***
	(0.105)	(0.247)	(0.116)
Constant	24.848***	25.676***	24.719***
	(0.087)	(0.200)	(0.099)
Observations	249,445	55,580	193,865
R-squared	0.25	0.25	0.24

Note: Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-taile

Table A2.2: Robustness Check (Coarsened Exact Matching)

	Song Novelty (logged)			
	Regular Sample	Matched Sample	Regular Sample	Matched Sample
Female	0.533***	0.440***	-0.321**	-0.497***
	(0.037)	(0.039)	(0.098)	(0.103)
Network Size	-0.007***	-0.002	-0.009***	-0.007***
	(0.001)	(0.001)	(0.001)	(0.001)
Female X Network Size			0.023**	0.025***
			(0.002)	(0.003)
Network Composition	-0.084	0.034	0.547***	0.515***
	0.054	(0.057)	(0.101)	(0.111)
Female X Network Composition			-0.438***	-0.287*
			(0.126)	(0.136)
Genre Composition	-1.901***	-1.891***	-2.406***	-2.479***
	(0.105)	(0.110)	(0.115)	(0.122)
Female X Genre Composition			1.952***	2.224***
			(0.253)	(0.260)
Constant	24.848***	25.617***	25.033***	25.794***
	(0.087)	(0.097)	(0.088)	(0.097)
Observations	249,445	202,498	249,445	202,517
R-squared	0.25	0.24	0.25	0.24

Note: All models in this table also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared), and # of Genres Spanned (squared). Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table A2.3: Robustness Check (First Song vs. Tenure > Median)

	Song Novelty (logged)			
	First Song	Tenure > Median	First Song	Tenure > Median
Female	0.254**	0.277***	-0.207	0.037
	(0.090)	(0.057)	(0.194)	(0.019)
Network Size	-0.005	-0.008***	-0.007*	-0.009***
	(0.003)	(0.001)	(0.003)	(0.001)
Female X Network Size			0.019*	0.004
			(0.009)	(0.003)
Network Composition	0.139	0.154	0.001	0.590***
	(0.111)	(0.096)	(0.227)	(0.157)
Female X Network Composition			0.351	-0.581**
			(0.272)	(0.213)
Genre Composition	-1.412***	-0.028	-1.634***	-0.238
	(0.198)	(0.179)	(0.221)	(0.192)
Female X Genre Composition			0.795	0.890
			(0.458)	(0.480)
Constant	27.374***	11.123***	27.448***	11.218***
	(0.135)	(0.262)	(0.137)	(0.263)
Observations	60,683	97,327	60,683	97,327
R-squared	0.24	0.28	0.25	0.28

Note: All models in this table also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared), and # of Genres Spanned (squared). Robust standard errors reported in parentheses. *p<0.05; **p<0.01; ***p<0.001 (two-tailed tests).

Table A2.4: Logistic Regression ("Optimal" vs. "Radical" Song Novelty)

	Optimal Song Novelty		*Radical Song Novelty*	
Female	0.05**	-0.11*	0.18**	-0.04
	(0.02)	(0.04)	(0.02)	(0.06)
Network Size	0.001***	0.001***	-0.001**	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Female X Network Size		0.00		0.01***
		(0.00)		(0.00)
Network Composition	-0.02	0.09*	0.05	0.22**
	(0.03)	(0.05)	(0.04)	(0.07)
Female X Network Composition		-0.15*		-0.09
		(0.06)		(0.08)
Genre Composition	-0.65***	-0.78***	-0.95***	-1.08***
	(0.05)	(0.05)	(0.07)	(0.08)
Female X Genre Composition		0.56***		0.44**
		(0.11)		(0.16)
Constant	-1.00***	-0.97***	-3.13***	-3.06***
	(0.03)	(0.04)	(0.06)	(0.06)
Observations	249,445	249,445	249,445	249,445
Pseudo R-squared (Cragg-Uhler)	0.06	0.06	0.17	0.17

Note: "Optimal Novelty" is defined as those songs with Song Novelty scores between the 70th and 90th percentiles, while "Radical Novelty" is defined as those songs with Song Novelty scores above the 90th percentile. All models in this table also include dummies for Decades, US, Classical, Jazz, Charting Artist, and Major Label, along with controls for Past Creativity, Artist Productivity, Artist Tenure (squared), and # of Genres Spanned (squared).
*p<0.05; **p<0.01; ***p<0.001 (two-tailed tests)

Table 3.1: Descriptive Statistics and Pearson Correlations for Variables of Interest

	[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	0.72	1																			
[3] Genre-weighted typicality (yearly)	-0.04	-0.09	1																		
[4] Genre-weighted typicality (weekly)	-0.03	-0.08	0.99	1																	
[5] Major label dummy	0.07	0.10	-0.06	-0.06	1																
[6] Long song	0.03	0.00	-0.06	-0.06	0.02	1															
[7] Crossover track	0.02	0.03	-0.04	-0.04	0.04	0.00	1														
[8] Multiple memberships	0.05	0.01	-0.01	-0.01	0.03	0.02	0.00	1													
[9] 2-3 Previously charting songs	-0.08	-0.02	0.00	0.00	-0.04	-0.03	-0.01	-0.11	1												
[10] 4-10 Previously charting songs	0.01	-0.01	0.01	0.01	0.05	0.00	-0.01	-0.01	-0.32	1											
[11] > 10 Previously charting songs	0.05	-0.06	0.01	0.01	0.06	0.03	0.02	0.26	-0.32	-0.40	1										
[12] Song tempo	-0.01	-0.02	0.07	0.07	-0.01	-0.02	-0.01	0.00	0.00	-0.01	0.01	1									
[13] Song energy	-0.01	0.05	0.16	0.16	0.01	-0.03	0.00	0.03	0.00	0.00	-0.02	0.16	1								
[14] Song speechiness	-0.01	0.05	-0.10	-0.10	0.00	0.02	0.00	-0.02	0.02	-0.02	-0.04	0.01	0.09	1							
[15] Song acousticness	-0.04	-0.19	0.00	-0.01	-0.08	0.00	-0.02	-0.05	-0.01	0.01	0.01	-0.08	-0.56	-0.11	1						
[16] Minor/Major Mode (0 or 1)	-0.01	-0.06	0.66	0.65	-0.01	-0.02	-0.05	0.00	0.00	0.00	0.01	0.03	-0.07	-0.10	0.13	1					
[17] Song danceability	0.03	0.13	0.12	0.13	-0.02	-0.06	0.07	0.02	0.03	-0.03	-0.05	-0.14	0.16	0.19	-0.28	-0.14	1				
[18] Song valence	0.00	-0.05	0.30	0.30	-0.11	-0.10	0.00	-0.01	0.04	-0.02	-0.06	0.09	0.30	0.03	-0.11	-0.04	0.47	1			
[19] Song instrumentalness	0.00	-0.03	-0.30	-0.30	-0.05	0.02	0.01	0.00	0.02	-0.03	-0.05	0.01	-0.08	-0.07	0.12	-0.02	-0.03	0.03	1		
[20] Song liveness	0.05	0.00	-0.11	-0.11	0.01	0.02	-0.04	0.01	-0.01	0.03	0.00	0.01	0.15	0.08	0.02	0.02	-0.22	-0.07	-0.02	1	
[21] Song time singature = 4/4	0.05	0.08	0.09	0.10	0.03	-0.03	0.00	0.03	0.01	0.00	0.00	-0.03	0.27	0.01	-0.27	-0.05	0.26	0.18	-0.07	0.01	1
Mean	56.27	11.57	0.81	0.81	0.67	0.04	0.24	0.08	0.21	0.29	0.28	119.09	0.59	0.07	0.34	0.74	0.58	0.62	0.08	0.24	0.90
Standard Deviation	30.45	7.79	0.06	0.06	0.47	0.19	0.42	0.27	0.40	0.45	0.45	27.70	0.22	0.08	0.31	0.44	0.16	0.24	0.21	0.22	0.29
Minimum	1.00	1.00	0.26	0.26	0	0	0	0	0	0	0	0	0	0.02	0	0	0	0	0	0.01	0
Maximum	100.00	87.00	0.92	0.92	1.00	1.00	1.00	1.00	1.00	1.00	1.00	242.51	1.00	0.96	1.00	1.00	0.99	1.00	1.00	1.00	1.00

Table 3.2: Models Predicting *Billboard* Hot 100 Chart Performance

MODEL:	3	4	5	6
	Ordered Logit	Ordered Logit	Negative Binomial	Negative Binomial
OUTCOME VARIABLE:	Peak Position (inverted)	Peak Position (inverted)	Weeks on Charts	Weeks on Charts
Genre-weighted typicality (yearly)	-2.419**	7.672*	-0.538**	1.791
	(0.429)	(2.987)	(0.150)	(1.051)
Genre-weighted typicality (yearly)^2		-6.805**		-1.570*
		(2.004)		(0.698)
Major label dummy	0.145**	0.145**	0.0246**	0.0245**
	(0.0255)	(0.0255)	(0.00883)	(0.00882)
Long song	0.262**	0.265**	0.0291	0.0290
	(0.0609)	(0.0608)	(0.0193)	(0.0193)
2-3 previously charting songs	-0.306**	-0.306**	-0.138**	-0.138**
	(0.0353)	(0.0353)	(0.0119)	(0.0119)
4-10 previously charting songs	-0.0305	-0.0298	-0.118**	-0.118**
	(0.0331)	(0.0331)	(0.0108)	(0.0108)
10+ previously charting songs	0.0874*	0.0878*	-0.168**	-0.168**
	(0.0347)	(0.0347)	(0.0115)	(0.0115)
Crossover track	0.151**	0.149**	-0.00556	-0.00590
	(0.0303)	(0.0303)	(0.0107)	(0.0107)
Multiple memberships	0.146**	0.147**	0.0554**	0.0559**
	(0.0417)	(0.0417)	(0.0133)	(0.0133)
Reissued track	-0.204*	-0.204*	-0.0812*	-0.0814*
	(0.0923)	(0.0921)	(0.0409)	(0.0409)
Half-Decade Dummies				
1987-1991	0.265**	0.232**	0.440**	0.432**
	(0.0697)	(0.0702)	(0.0217)	(0.0218)
1992-1996	-0.282**	-0.328**	0.567**	0.557**
	(0.0701)	(0.0714)	(0.0239)	(0.0241)
Observations	25,077	25,077	25,077	25,077

Robust standard errors in parentheses
Reference categories for dummy variables: Pop (genre), Independent label, 1st charting song (previously charting songs), Key of E-Flat, and all non-4/4 time signatures.
** p<0.01, * p<0.05

Table 3.3: Fixed Effect Models Predicting Change in Chart Position (Weekly)

MODEL:	7	8
OUTCOME VARIABLE:	Change in (Inverted) Chart Position	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)
Week (on charts)	-1.941**	-1.941**
	(0.0129)	(0.0129)
Week (on charts)^2	0.0334**	0.0334**
	(0.000478)	(0.000477)
Constant	21.95**	3.805
	(1.873)	(5.848)
Observations	263,715	263,715
R-squared	0.432	0.432

Robust standard errors in parentheses
** p<0.01, * p<0.05, two-tailed test

Figures

Figure 1.1: Distribution of Song Novelty

Note: Dashed vertical lines in all plots represent median (as opposed to mean) values.

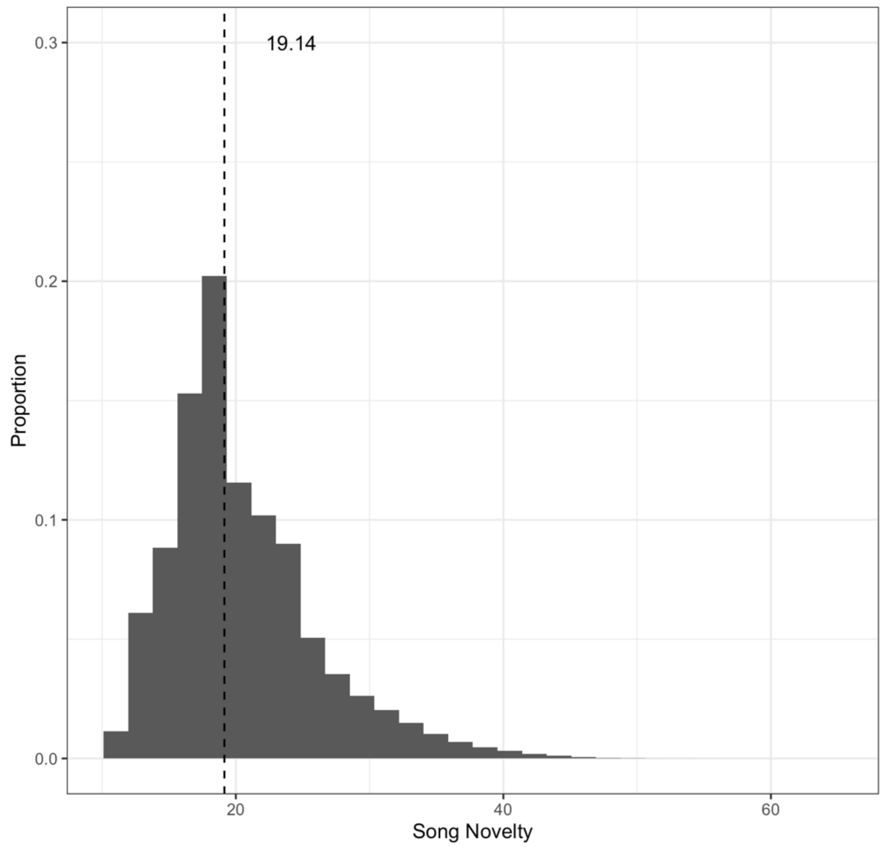

Figure 1.2a-d: Histograms of # of Alters by Sphere (Full and Without Zeroes)

Collaborative Sphere

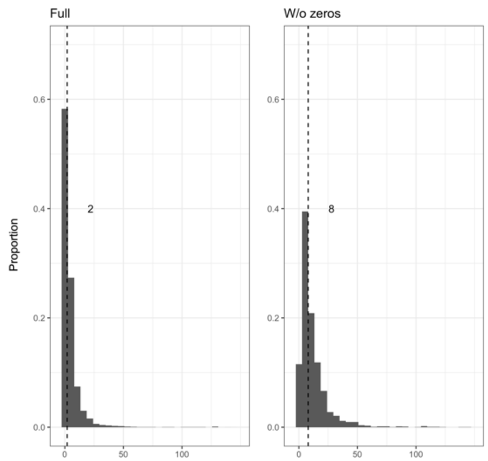

Cultural Sphere

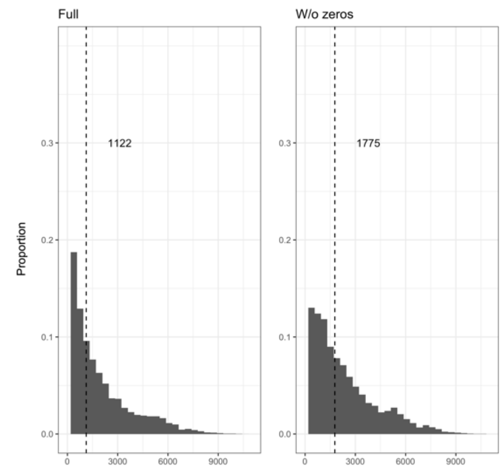

Organizational Sphere

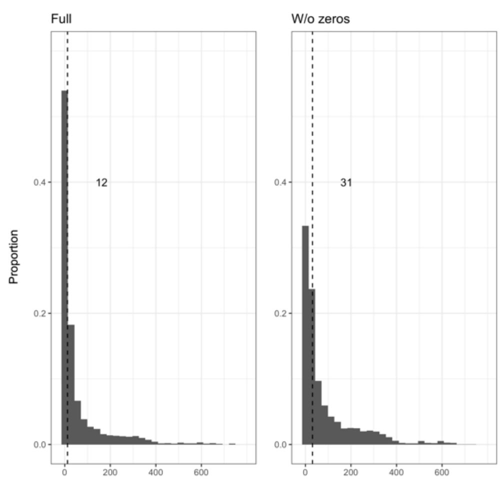

Geographic Sphere

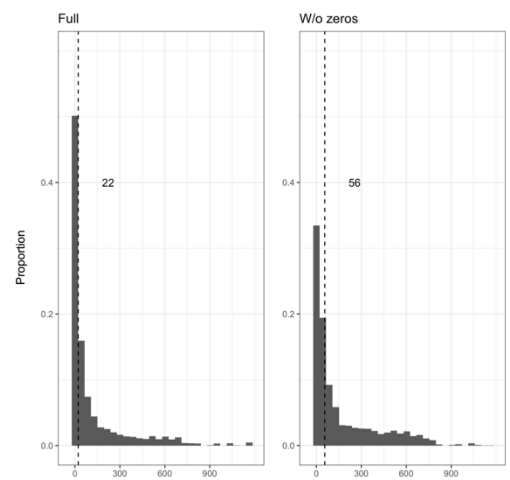

Figure 1.3a-d: Histograms of Genre Diversity by Sphere (Full, Imputed, Without Zeroes)

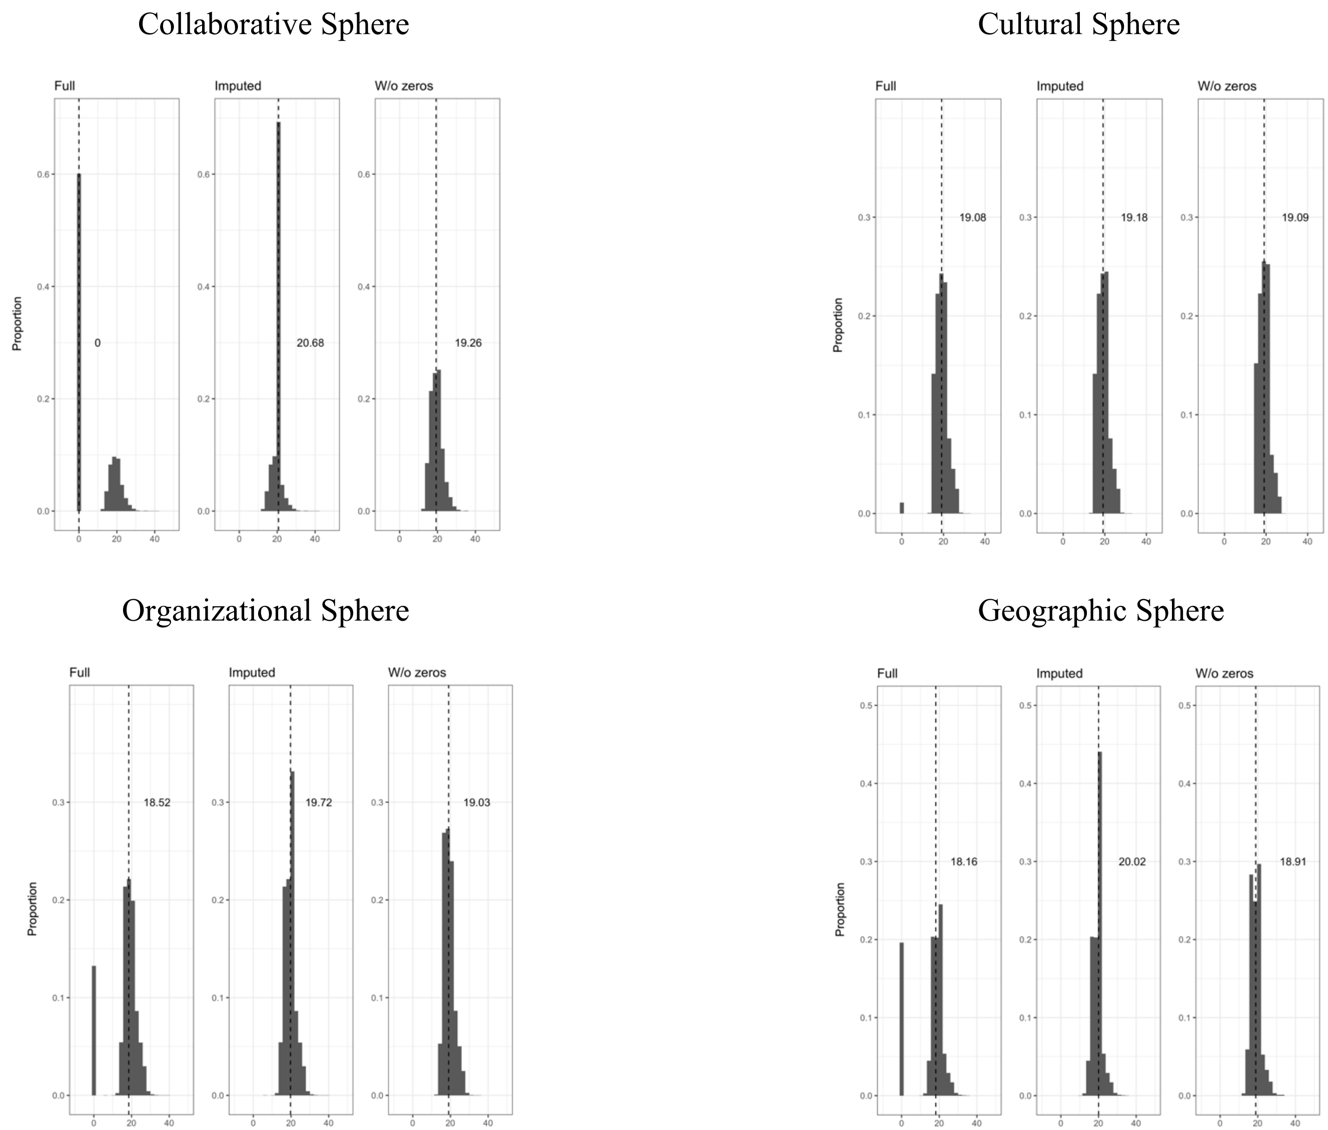

Figure 1.4a-d: Histograms of Alter Creativity by Sphere (Full, Imputed, Without Zeroes)

Collaborative Sphere

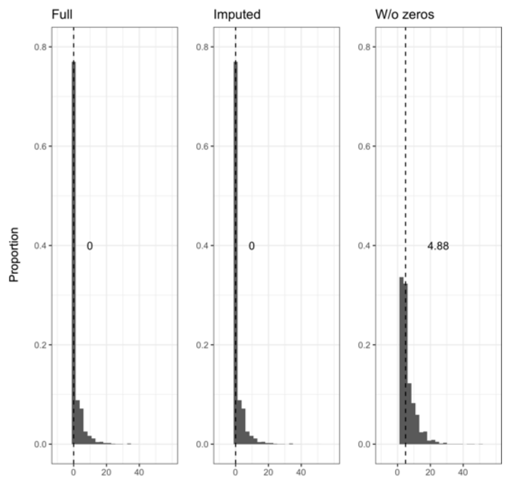

Cultural Sphere

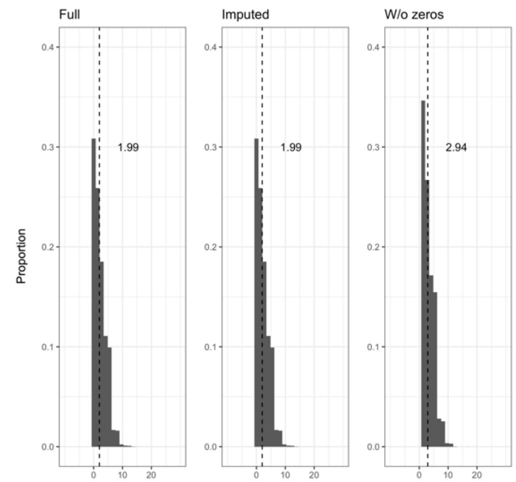

Organizational Sphere

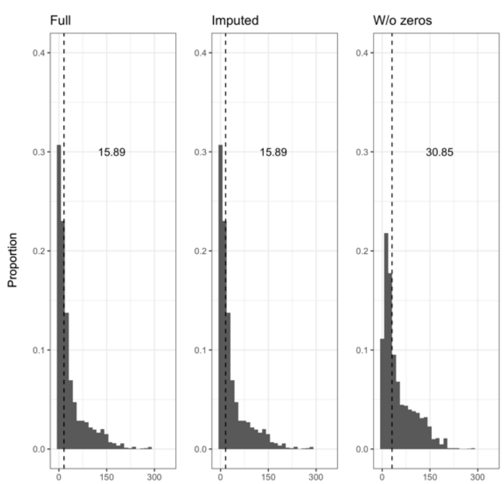

Geographic Sphere

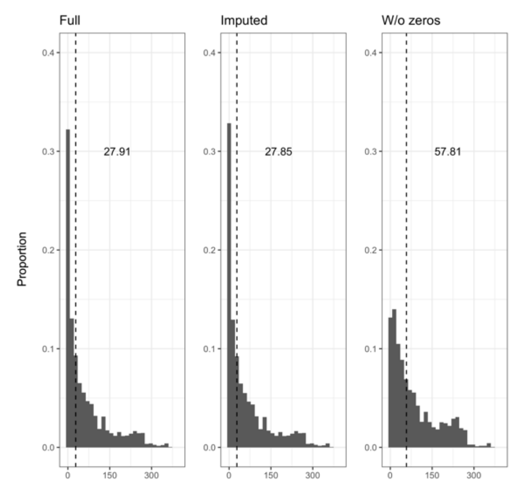

Figure 1.5. Musical Creativity Over Time

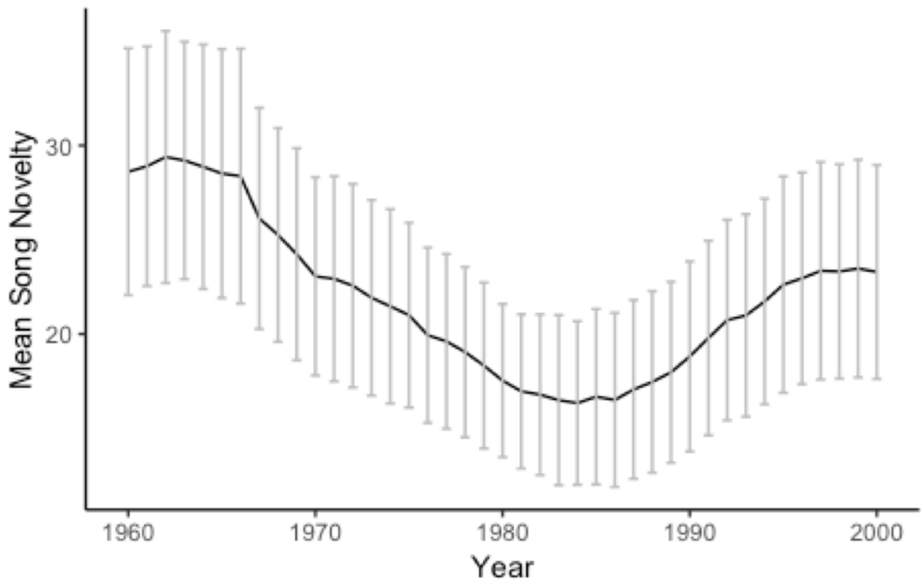

Figure 1.6: Standardized Regression Results (Table 1.4, Model 6)

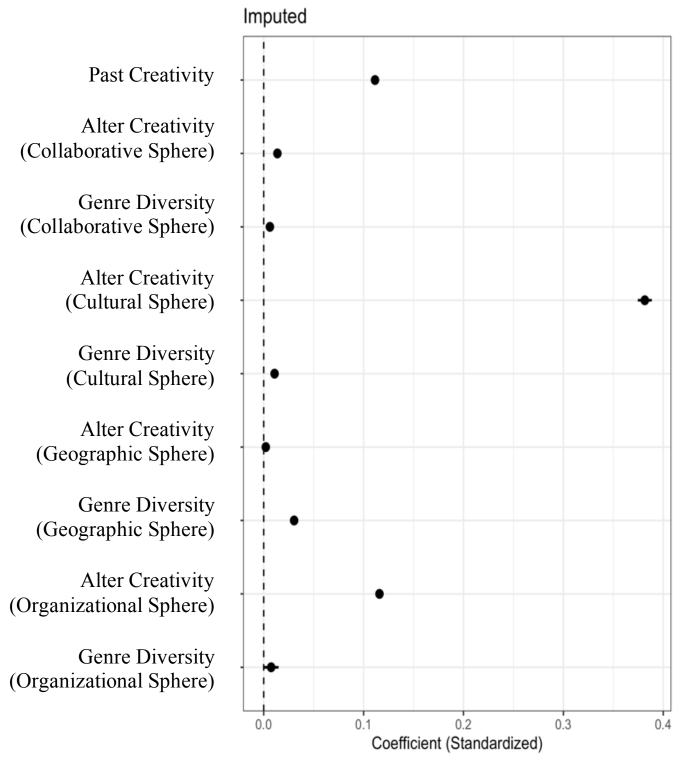

Figure 1.7a-d: Interaction Plots (Table 1.5, Model 5)

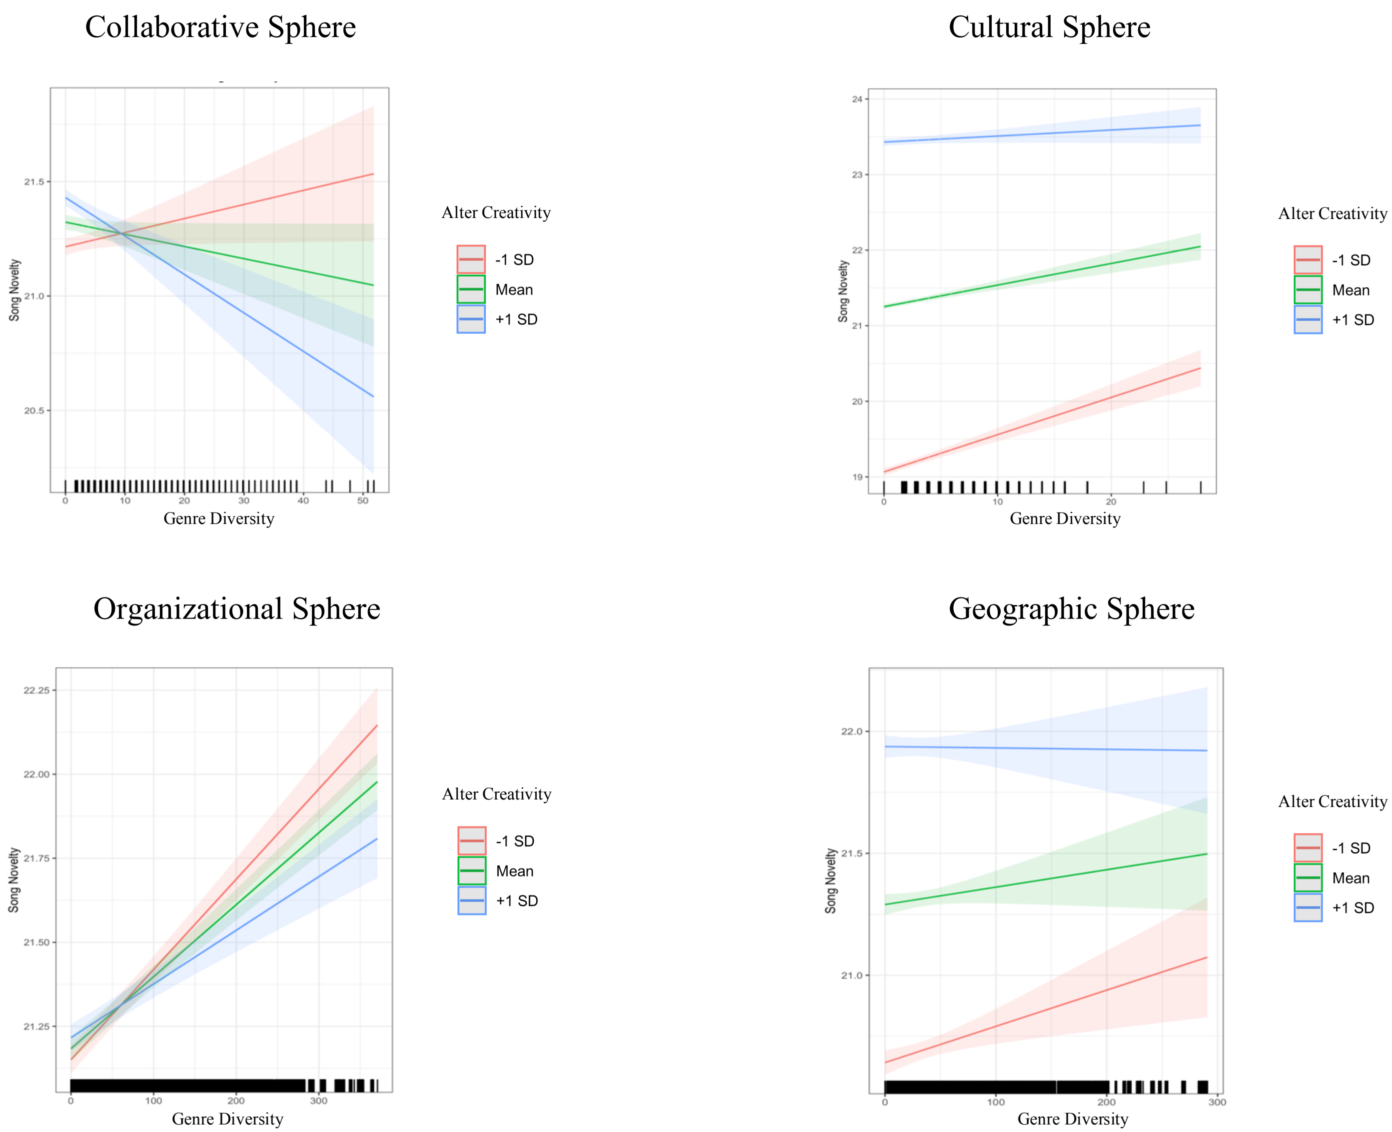

Figure A1.1: Standardized Regression Results (with Raw Measures and Without Zeroes)

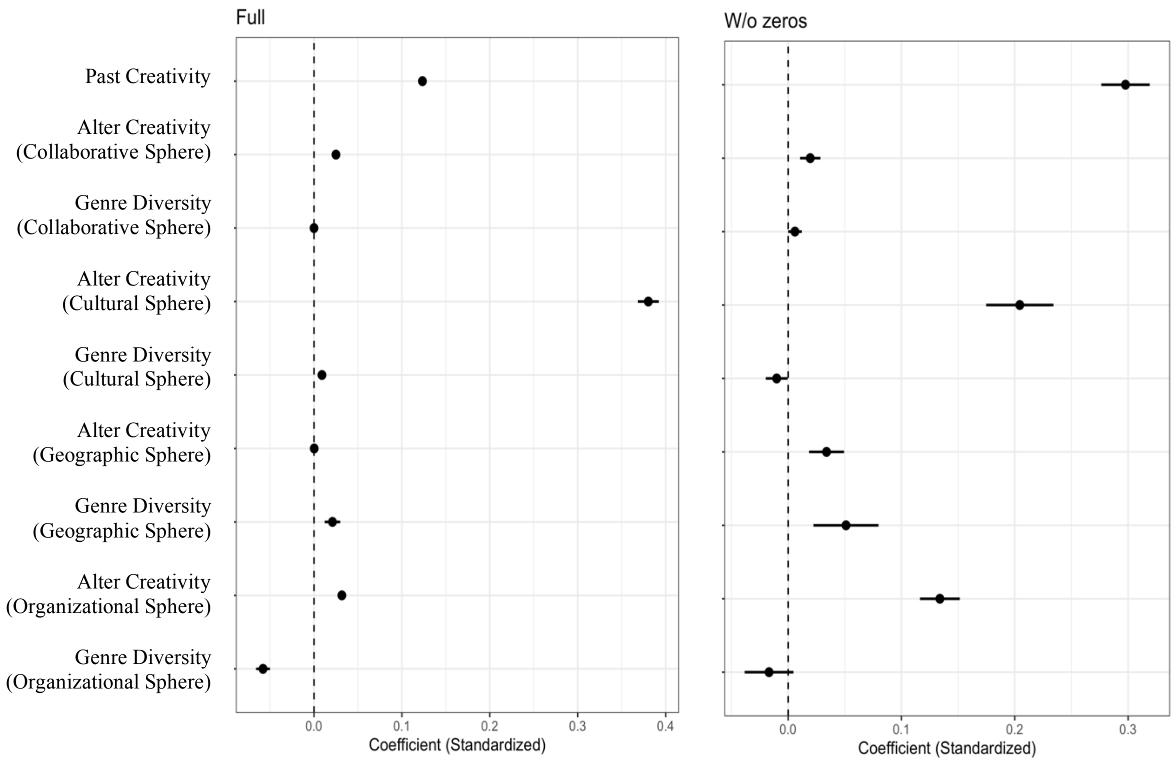

Figure 2.1: Histograms of Song Novelty (By Gender)

Note: Dashed vertical lines in all plots represent median (as opposed to mean) values.

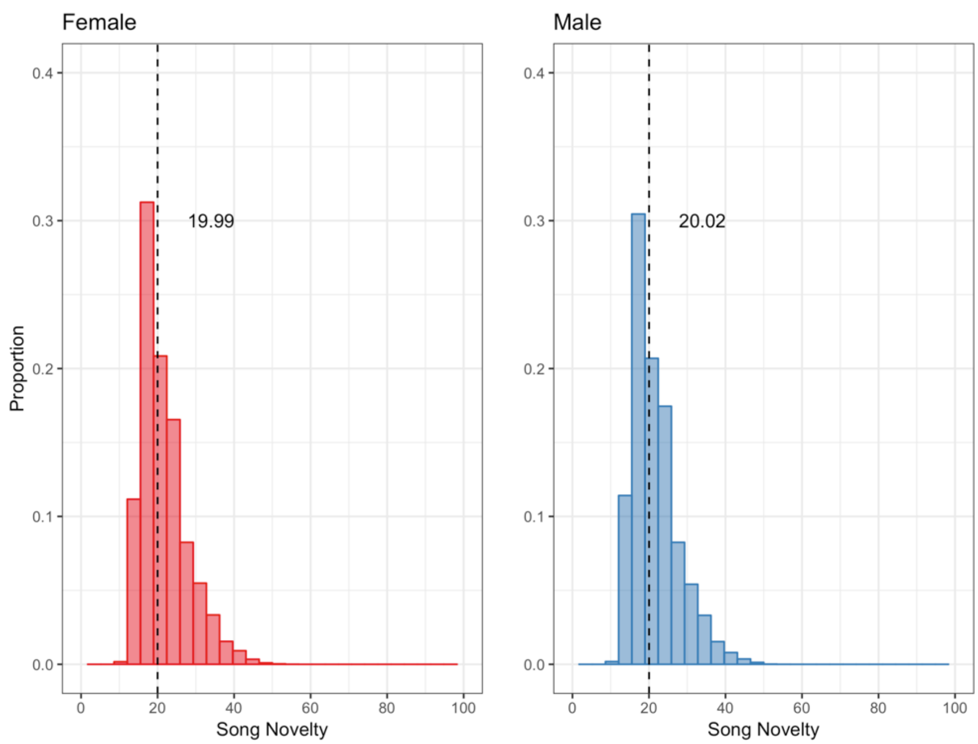

Figure 2.2: 2-D Visualization of Songs in Sonic Feature Space

Layout generated using t-Distributed Stochastic Neighbor Embedding (t-SNE) of 10,000 randomly selected songs.

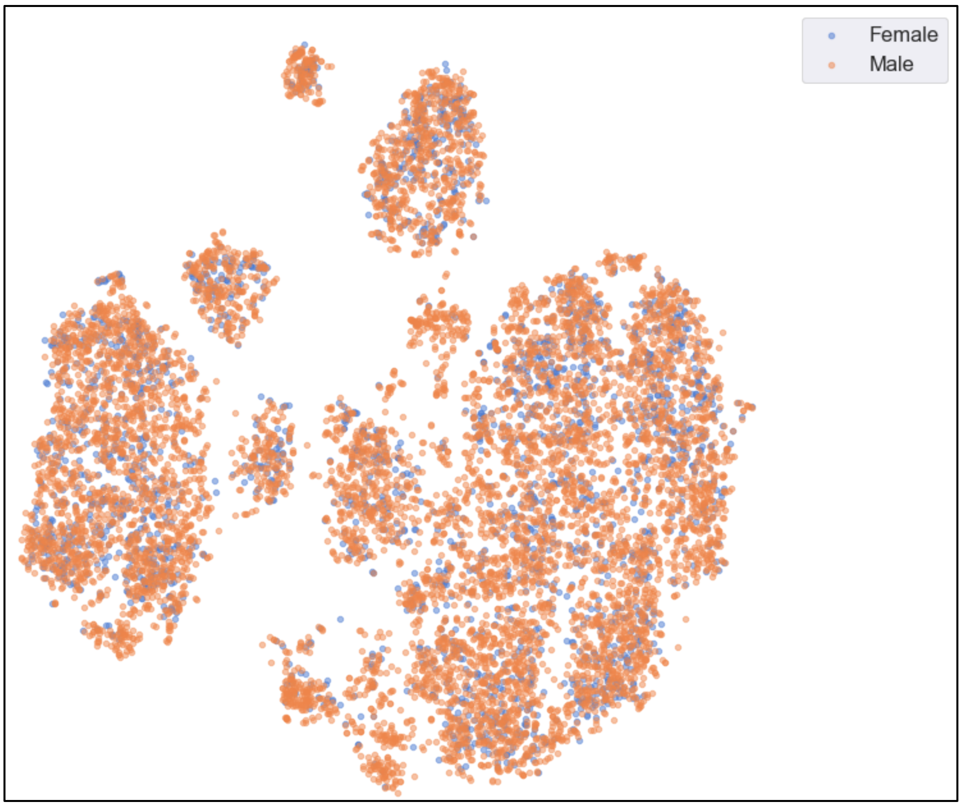

Figure 2.3: Standardized Regression Results (Table 2.3, Model 6)

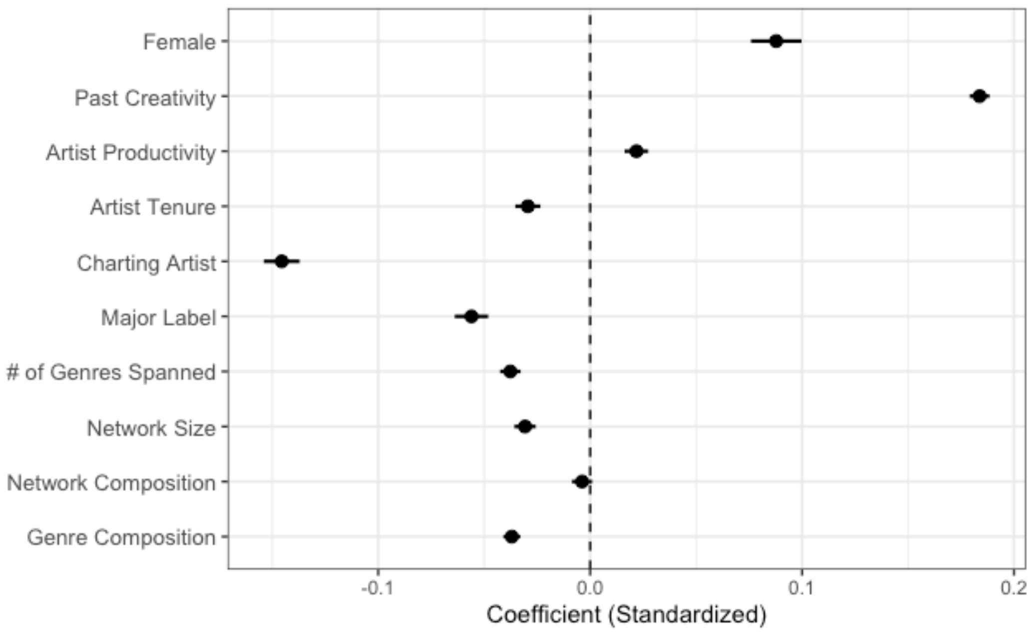

Figure 2.4a-c: Histograms and Marginal Effect Plots for Key Interactions (By Gender)

Predicted values plots generated from Table 2.4, Model 10 (full interaction model)
Note: Dashed vertical lines in all plots represent median (as opposed to mean) values.

2.4a: Number of Collaborators (Gender X Network Size)

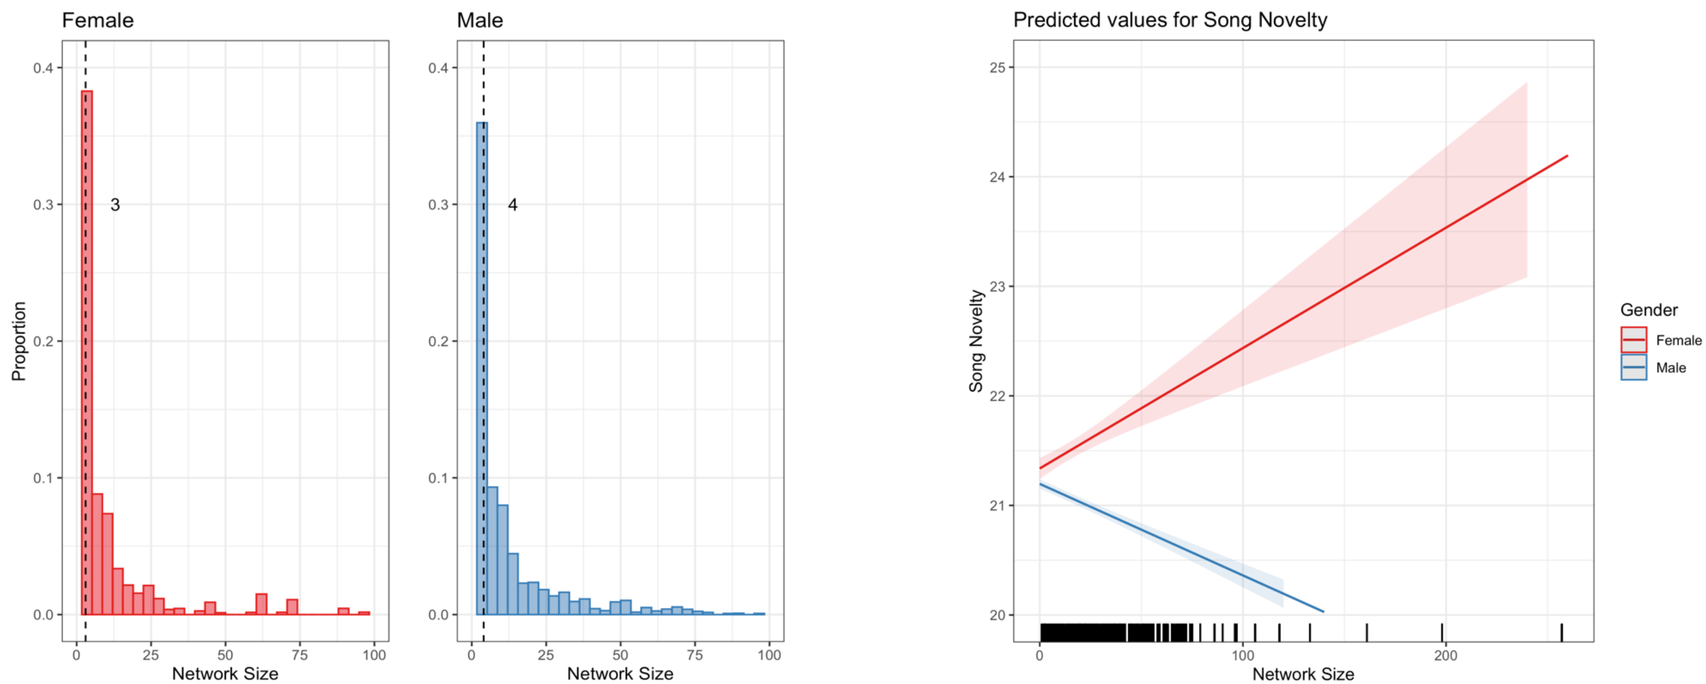

2.4b: Proportion of female artists in collaboration network (Gender X Network Composition)

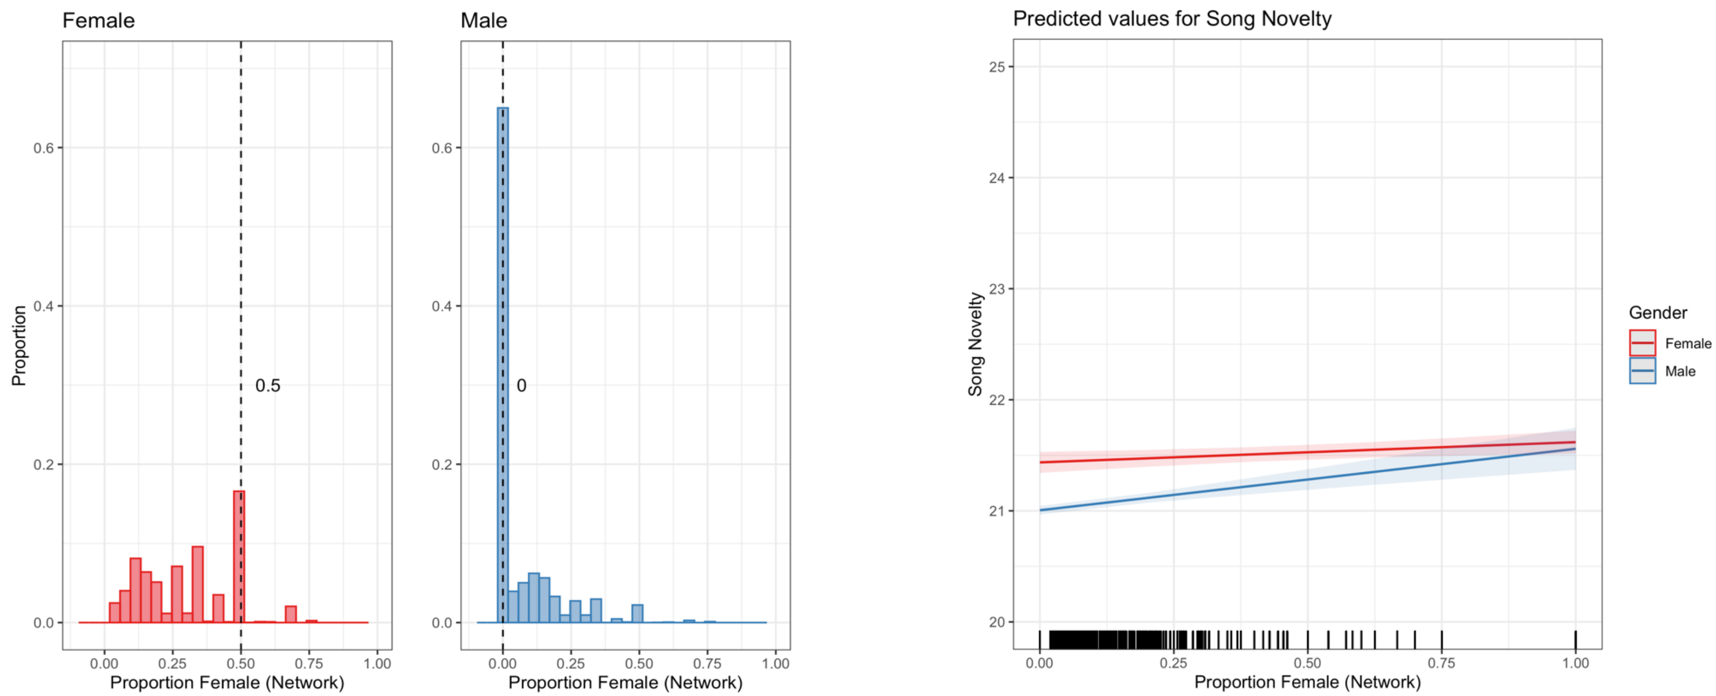

2.4c: Proportion of female artists in primary genre (Gender X Genre Composition)

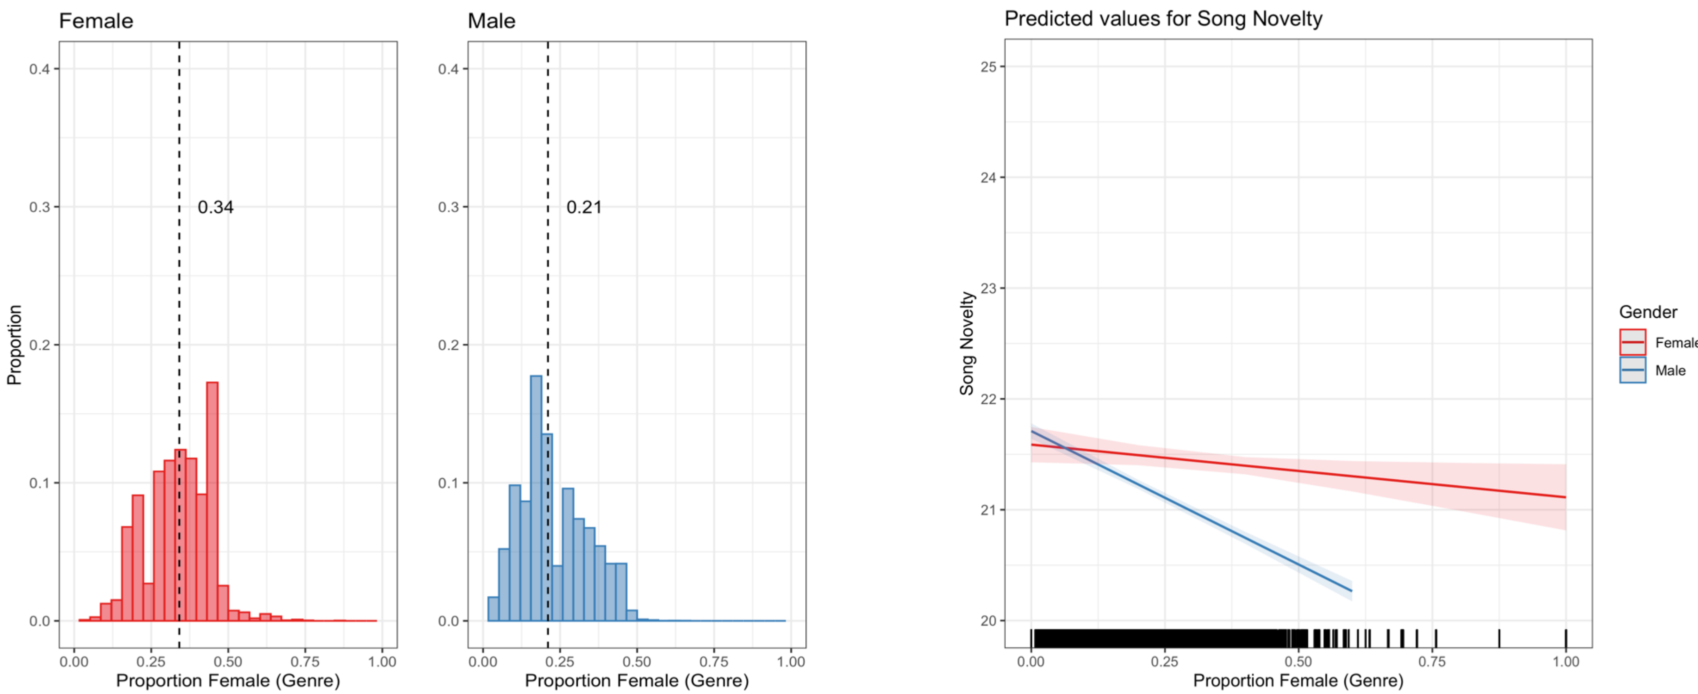

Figure 3.1: Distribution of Genre-Weighted Song Typicality (yearly)

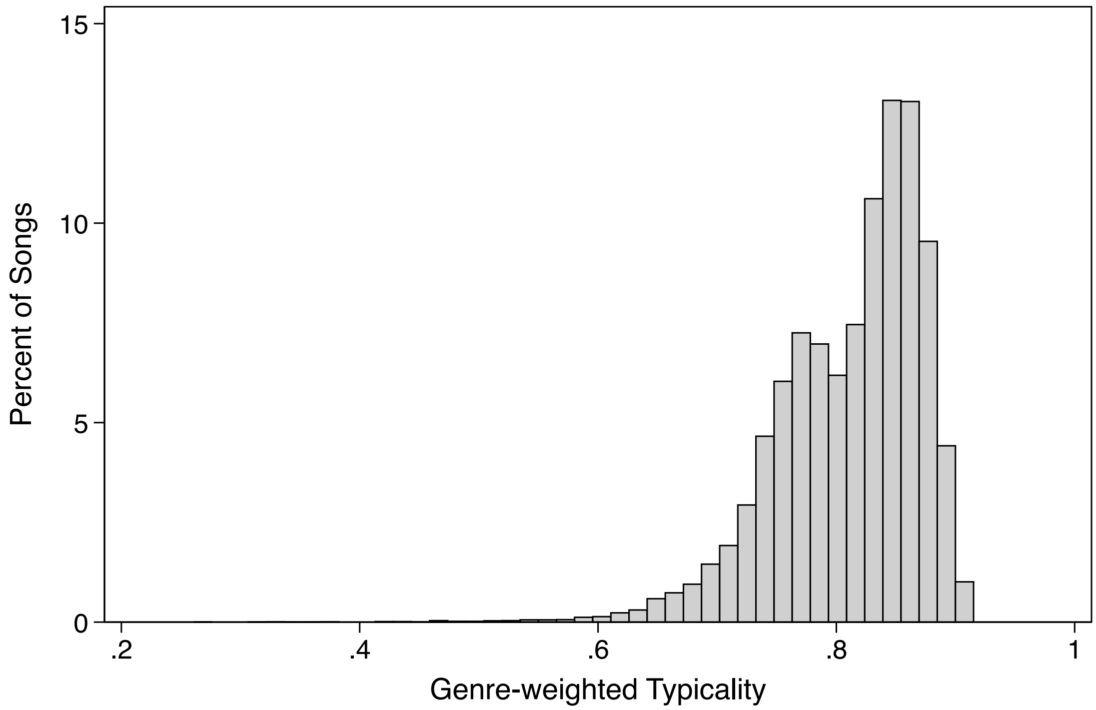

Figure 3.2a: Comparison of Typicality for Top 40 versus All Other Songs (By Year)

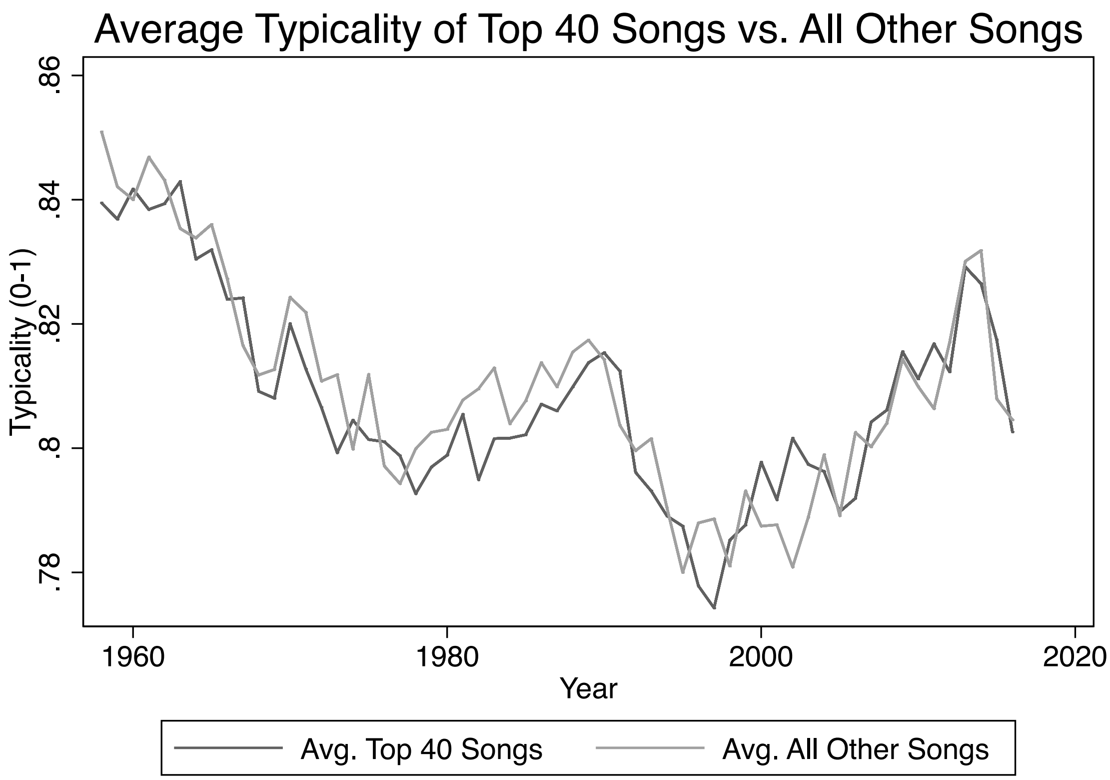

Figure 3.2b: Comparison of Typicality for #1 versus All Other Songs (By Year)

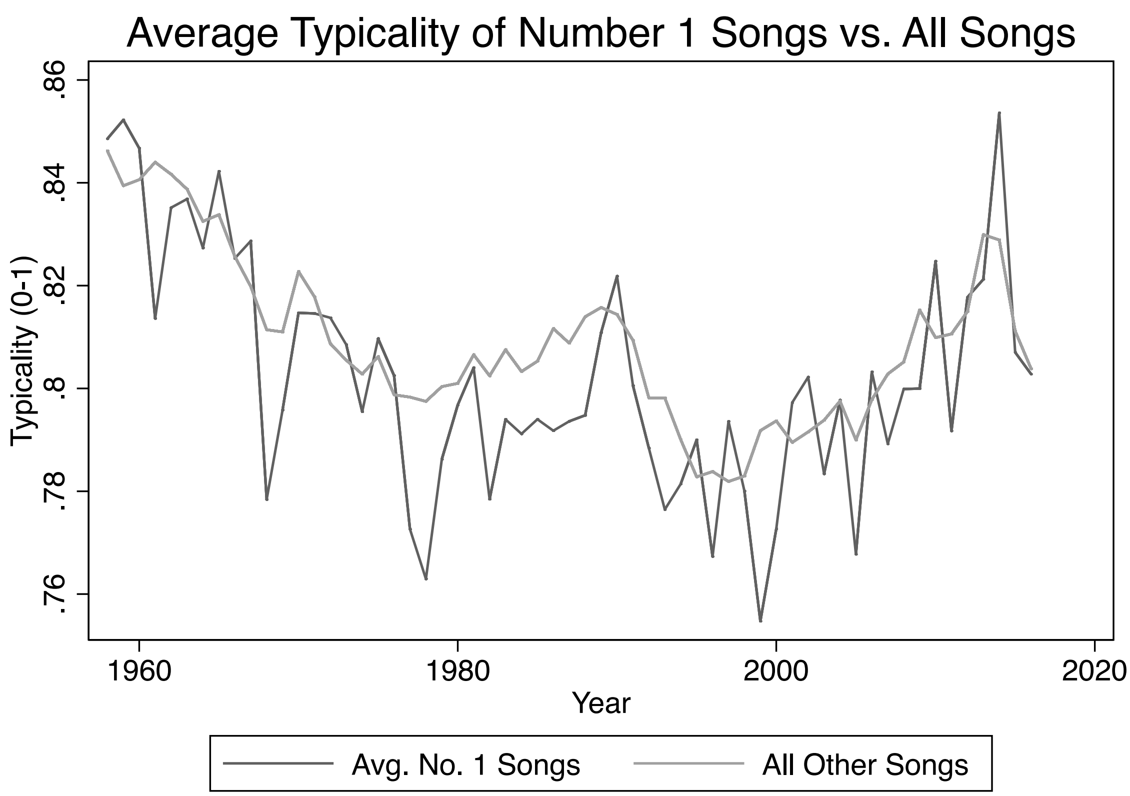

Figure 3.3. Standardized Coefficients from Models Predicting Chart Performance

Horizontal bars represent 95% CI
See Table 3.2 (Models 4 and 6) for full (unstandardized) results.

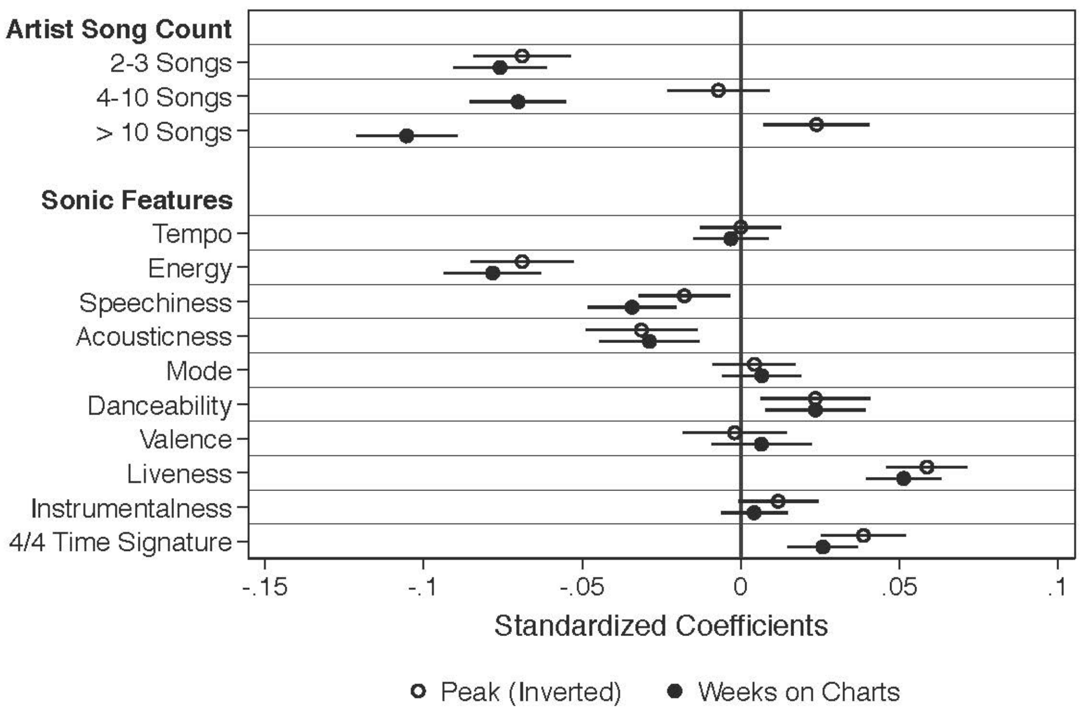

Figure 3.4: Predicted Marginal Probability of Achieving Selected Peak Position

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 below, the predicted positions are coded as they would be on the charts (i.e., #100 is the lowest, #1 the highest).

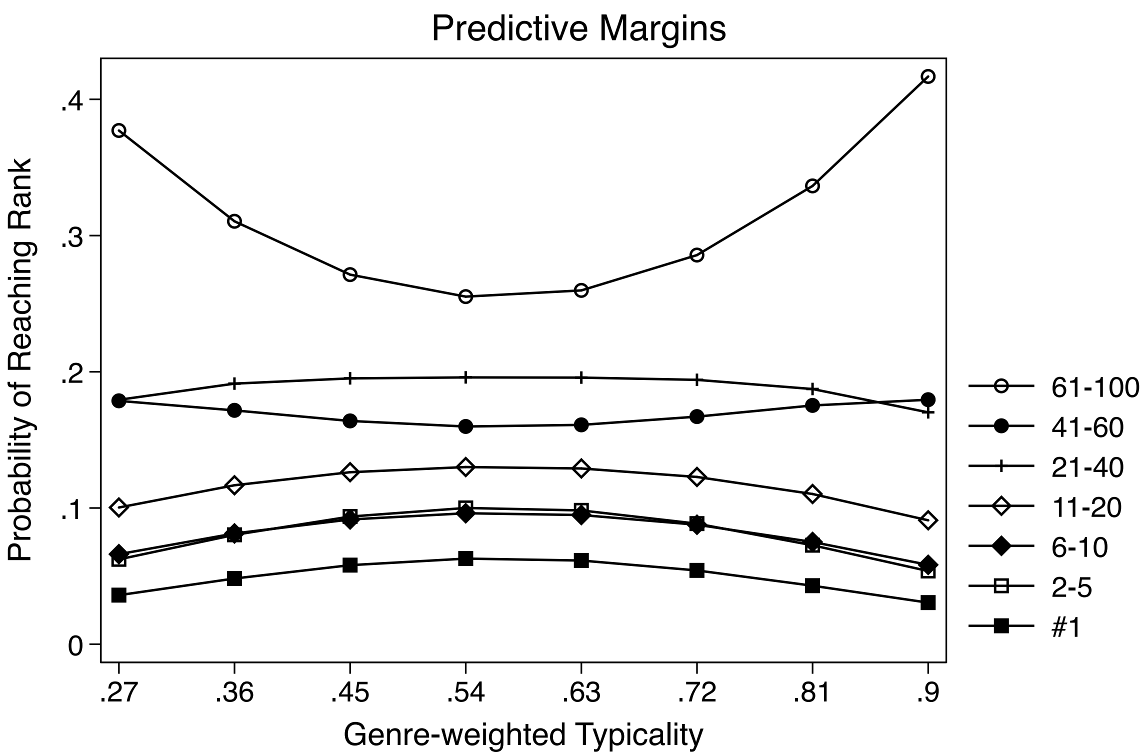

References

Chapter 1

Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: a longitudinal study. *Administrative Science Quarterly* 45: 425–455.

Amabile, T. M. 1996. *Creativity in context*. Boulder, CO: Westview Press.

Amabile, T. M. & Mueller, J. S. 2008. Studying creativity, its processes, and its. antecedents: An exploration of the componential theory of creativity. In J. Zhou & C. E. Shalley (Eds.), *Handbook of Organizational Creativity*, 33-64. New York: Lawrence Erlbaum.

Amabile, T.M. and Julianna Pillemer. 2012. Perspectives on the Social Psychology of Creativity. *The Journal of Creative Behavior* 46(1):3-15.

Askin, N. and M. Mauskapf. 2017. What Makes Popular Culture Popular?: Product Features and Optimal Differentiation in Music. *American Sociological Review 82(5):* 910–944.

Baker, T. and R.E. Nelson. 2005. Creating Something from Nothing: Resource Construction through Entrepreneurial Bricolage. *Administrative Science Quarterly* 50(3): 329–366.

Becker. H. 1982. *Art Worlds.* Berkeley, CA: University of California Press.

Benner, M.J. and J. Waldfogel. 2016. The Song Remains the Same?: Technological Change and Positioning in the Recorded Music Industry. *Strategy Science* 1(3): 129–147.

Bandura, A. and R.H. Walters. 1977. *Social Learning Theory.* New York: Prentice Hall.

Boccaletti, S., G. Bianconi, R. Criado, C. del Genio, J. Gomez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin. The Structure and Dynamics of Multilayer Networks. *Physics Reports* 544(1): 1–122.

Boschma, R.A. 2005. Proximity and Innovation: A Critical Assessment. *Regional Studies* 39(1): 61–74.

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

Breiger, R.L. and K. Puetz. 2015. Culture and Networks. *The International Encyclopedia of the Social and Behavioral Sciences.* Pp. 557–62 in International Encyclopedia of Social and Behavioral Sciences, vol. 5, edited by J. D. Wright. Oxford, UK: Elsevier.

Burt, R.S. 1992. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

_______. 2004. Structural Holes and Good Ideas. *American Journal of Sociology* 110(2): 349–399.

_______. 2005. *Brokerage and Closure: An Introduction to Social Capital.* New York: Oxford University Press.

Carnabuci, G. and E. Operti. 2013. Where do Firm's Recombinant Capabilities Come From?: Intraorganizational Networks, Knowledge, and Firm's Ability to Innovate through Technological Recombination. *Strategic Management Journal* 34(13): 1591–1613.

Cattani, F., S. Ferriani, and P.D. Allison. 2014. Insiders, Outsiders, and the Struggle for Consecration in Cultural Fields: A Core-Periphery Perspective. *American Sociological Review* 79(2): 258–281.

Cerulo, K. 1988. Analyzing Cultural Products: A New Method of Measurement. *Social Science Research* 17: 317-352.

Cho, S.M. and M. Mauskapf. 2019. The Effects of Competitive Differentiation on Critical versus Popular Success in a Cultural Market. Working paper.

Collins, R. 2005. *Interaction Ritual Chains*. Princeton, NJ: Princeton University Press.

Cooper, A. and T. Folta. 2000. Entrepreneurship and High-technology Clusters. In D.L. Sexton and H. Landstrom (eds.), *The Blackwell Handbook of Entrepreneurship*, pp. 348–367. Oxford: Blackwell Business.

Crossley, N. 2009. The Man Whose Web Expanded: Network Dynamics in Manchester's Post/punk Music Scene, 1976–1980. *Poetics* 37: 24–49.

de Vaan, M., D. Stark, and B. Vedres. 2015. Game Changer: The Topology of Creativity. *American Journal of Sociology* 120:1144-1194.

DiMaggio, P. 1997. Culture and Cognition. *Annual Review of Sociology* 23: 263–287.

DiMaggio, P. and W.W. Powell. 1983. The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields. *American Sociological Review* 48(2): 147–160.

Dowd, T.J. 2004. Concentration and Diversity Revisited: Production Logics and the U.S. Mainstream Recording Market, 1940–1990. *Social Forces* 82(4): 1411–1455.

Durkheim, E. 1912/1965. *The Elementary Forms of Religious Life*. London: Allen & Unwin.

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

Fleming, L. and D. Waguespack. 2007. Brokerage, Boundary Spanning, and Leadership in Open Innovation Communities. *Organization Science* 18(2): 165–180.

Fleming, L., S. Mingo, and D. Chen. 2007. Collaborative Brokerage, Generative Creativity, and Creative Success. *Administrative Science Quarterly* 52:443-475.

Fligstein, N. and L. Dauter. 2007. The Sociology of Markets. *Annual Review of Sociology* 33(1):105–28.

Florida, R., P. Adler, and C. Mellander. 2017. The City as Innovation Machine. *Regional Studies* 51(1): 86 – 96.

Friberg, A., W. Schoonderwaldt, A. Hedblad, M. Fabiani, and A. Elowsson. 2014. Using Listener-based Perceptual Features as Intermediate Representations in Music Information Retrieval. *Journal of the Acoustical Society of America* 136(4): 1951–1963.

Friedlander, J.P. 2014. *News and Notes on 2014 RIAA Music Industry Shipment and Revenue Statistics.* Recording Industry Association of America.

Funk, R.J. 2014. Making the Most of Where You Are: Geography, Networks, and Innovation in Organizations. *Academy of Management Journal* 57(1): 193–222.

Goldberg, A., M.T. Hannan, and B. Kovács. 2016. What Does It Mean to Span Cultural Boundaries? Variety and Atypicality in Cultural Consumption. *American Sociological Review* 81:215-241.

Goldberg, A., S.B. Srivastava, V.G. Manian, W. Monroe, and C. Potts. 2016. Fitting in or Standing Out?: The Tradeoffs of Structural and Cultural Embeddedness. *American Sociological Review* 81(6): 1190–1222.

Granovetter, M.S. 1985. Economic Action and Social Structure: The Problem of Embeddedness. *American Journal of Sociology* 91:481-510.

Greenberg, D.M., M. Kosinski, D.J. Stillwell, B.L. Monteiro, D.J. Levitin, and P.J. Rentfrow. 2016. The Song is You: Preferences for Musical Attribute Dimensions Reflect Personality. *Social Psychological and Personality Science.*

Guzman, J. and S. Stern. 2015. Where is Silicon Valley? *Science* 347(6222): 606–609.

Heaney, M.T. 2014. Multiplex Networks and Interest Group Influence Reputation: An Exponential Random Graph Model. *Social Networks* 36: 66–81.

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

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

Hsu, G. 2006. Jack of All Trades and Masters of None: Audiences' Reactions to Spanning Genres in Feature Film Production. *Administrative Science Quarterly* 51: 420–450.

Hsu, G. and M.T. Hannan. 2005. Identities, Genres, and Organizational Forms." Organization Science 16:474-490.

Jones, B.F. 2010. Age and Great Invention. *Review of Economics and Statistics* 92(1): 1–14.

Jones, C., M. Lorenzen, and J. Sapsed (eds.). 2015. *The Oxford Handbook of Creative Industries.* New York: Oxford University Press.

Karas, M. 2016. REDEF Music Newsletter, May 2, 2016. <https://redef.com/mix/music-mix-1462167047636>

Kovacs, B. and M.T. Hannan. 2015. Conceptual Spaces and the Consequences of Category Spanning. *Sociological Science* 2: 252–286.

Lena, J.C. and R.A. Peterson 2008. Classification as Culture: Types and Trajectories of Music Genres. *American Sociological Review* 73(5): 697-718.

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

Lena, J.C. 2015. Genre: Relational Approaches to the Sociology of Music. In *Routledge International Handbook of the Sociology of Art and Culture*, edited by Laurie Hanquinet and Mike Savage. Pp. 149–160. New York: Routledge.

Lippmann, S. and H.E. Aldrich. 2016. A Rolling Stone Gathers Momentum: Generational Units, Collective Memory, and Entrepreneurship. *Academy of Management Review* 41(4): 658–675.

Mauch, M., R.M. MacCallum, M. Levy, and A.M. Leroi. 2015. The Evolution of Popular Music: USA 1960–2010. *Royal Society Open Science* 2.

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

Obstfeld, D. 2005. Social networks, the tertius iungens orientation, and involvement in innovation. *Administrative Science Quarterly* 50(1): 100-130.

Padgett, J.F. and W.W. Powell. 2012. *The Emergence of Organizations and Markets*. Princeton: Princeton University Press.

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

Peterson, R.A. and N. Anand. 2004. The Production of Culture Perspective. *Annual Review of Sociology* 30: 311–334.

Peterson, R.A. and D.G. Berger. 1971. Entrepreneurship in Organizations: Evidence from the Popular Music Industry. *Administrative Science Quarterly* 16(1): 97–106.

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

Polanyi, K. 1944*. The Great Transformation.* New York: Academic Press.

Ratliff, B. 2016. *Every Song Ever: Twenty Ways to Listen in an Age of Musical Plenty.* New York: Farrar, Straus and Giroux.

Reagans, R. and B. McEvily. 2003. Network Structure and Knowledge Transfer: The Effects of Cohesion and Range. *Administrative Science Quarterly* 48(2): 240–267.

Rubio, F.D. 2012. The Material Production of the Spiral Jetty: A Study of Culture in the Making. *Cultural Sociology* 6(2): 143-161.

Ruef, M. 2000. The Emergence of Organizational Forms: A Community Ecology Approach. *American Journal of Sociology* 106(3): 658–714.

______. 2002. Strong Ties, Weak Ties and Islands: Structural and Cultural Predictors of Organizational Innovation. *Industrial and Corporate Change* 11(3): 427–449.

______. 2010. *The entrepreneurial group: Social identities, relations, and collective action*. Princeton, NJ: Princeton University Press.

Runtagh, J. 2016. Songs on Trial: 12 Landmark Music Copyright Cases. *Rolling Stone Magazine*. June 8, 2016. < https://www.rollingstone.com/politics/politics-lists/songs-on-trial-12-landmark-music-copyright-cases-166396/led-zeppelin-vs-willie-dixon-1972-64390/>

Schein, E.H. 1990. Organizational Culture. *American Psychologist* 45(2): 109–119.

Schumpeter, J.A. 1934. *The Theory of Economic Development.* Cambridge, MA: Harvard Economic Studies.

______________ 1942. *Capitalism, Socialism and Democracy*. New York, NY: Harper.

Seidel, M-D.L. and H.R. Greve. Forthcoming. Emergence: How Novelty, Growth, and Formation Shape Organizations and their Ecosystems. *Research in the Sociology of Organizations.*

Simmel, G. 1950. *The Sociology of Georg Simmel*. New York: Free Press.

Simonton, D.K. 1984. Artistic Creativity and Interpersonal Relationships Across and Within Generations. *Journal of Personality and Social Psychology* 46(6): 1273–1286.

Soda, G., D. Stea, and T. Pedersen. 2019. Network Structure, Collaborative Context, and Individual Creativity. *Journal of Management* 45(4): 1739–1765.

Thornton, P.H., and K.H. Flynn. 2003. Entrepreneurship, networks, and geographies. In Z. J. Acs & D. B. Audretsch (Eds.), *Handbook of entrepreneurship research*: 401–433. London: Kluwer Law International.

Uzzi, B. 1996. The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *American Sociological Review* 61:674-698.

______. 1997. Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness. *Administrative Science Quarterly* 42:35-67.

Uzzi, B., S. Mukherjee, M. Stringer, and B. Jones. 2013. Atypical Combinations and Scientific Impact. *Science* 342:468-472.

Uzzi, B. and J. Spiro. 2005. Collaboration and Creativity: The Small World Problem. *American Journal of Sociology* 111:447-504.

Walsh, J.P. 1995. Managerial and Organizational Cognition: Notes from a Trip Down Memory Lane. *Organization Science* 6(3): 280–321.

Wang, D. 2015. Activating Brokerage: Interorganizational Knowledge Transfer through Skilled Return Migration. *Administrative Science Quarterly* 60(1): 133–176.

Welch, L. 1946. Recombination of Ideas in Creative Thinking. *Journal of Applied Psychology* 30(6): 638–643

White, D.R., W.W. Powell, and J. Owen-Smith. 2003. Embeddedness in Multiple Networks, Organization Theory, and Structural Cohesion Theory. *Computational and Mathematical Organization Theory*.

White, H.C. 1981. Where Do Markets Come From? *American Journal of Sociology*, 87: 517-547.

_________. 2002. *Markets from Networks: Socioeconomic Models of Production*. Princeton, N.J.: Princeton University Press.

Zaheer, A. and G.G. Bell. 2005. Benefiting from Network Position: Firm Capabilities, Structural Holes, and Performance. *Strategic Management Journal* 26(9): 809–825.

Zukin, S. and P. DiMaggio. 1990. *Structures of Capital: The Social Organization of the Economy*. New York: Cambridge University Press.

Chapter 2

Amabile, T.M. 1982. Social Psychology of Creativity: A Consensual Assessment Technique. *Journal of Personality and Social Psychology* 43(5): 997.

Amabile, T.M. 1996. *Creativity in Context*. Boulder, CO: Westview Press.

Amabile, T.M. 2012. Componential Theory of Creativity. *Harvard Business School* 12(96): 1-10.

Amabile, T.M and J. Pillemer. 2012. Perspectives on the Social Psychology of Creativity. *The Journal of Creative Behavior* 46(1): 3-15.

Anzia, S.F and C.R Berry. 2011. The Jackie (and Jill) Robinson Effect: Why Do Congresswomen Outperform Congressmen?. *American Journal of Political Science* 55(3): 478-93.

Apesteguia, J., G. Azmat and N. Iriberri. 2012. The Impact of Gender Composition on Team Performance and Decision Making: Evidence from the Field. *Management Science* 58(1): 78-93.

Askin, N. and M. Mauskapf. 2017. What Makes Popular Culture Popular? Product Features and Optimal Differentiation in Music. *American Sociological Review* 82(5): 910-44.

Askin, N. and J. Mol. 2018. Institutionalizing Authenticity in the Digitized World of Music. Pp. 159-202 in *Frontiers of Creative Industries: Exploring Structural and Categorical Dynamics*: Emerald Publishing Limited.

Baer, J. and J.C. Kaufman. 2008. Gender Differences in Creativity. *The Journal of Creative Behavior* 42(2): 75-105.

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

Barron, F. and D.M Harrington. 1981. Creativity, Intelligence, and Personality. *Annual review of psychology* 32(1): 439-76.

Becker, G.S. 1994. Human Capital Revisited. Pp. 15-28 in *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education (3rd Edition)*: The university of Chicago press.

Becker, H.S. [1982] 2008. *Art Worlds: Updated and Expanded*: University of California Press.

Benner, M.J and J. Waldfogel. 2016. The Song Remains the Same? Technological Change and Positioning in the Recorded Music Industry. *Strategy Science* 1(3): 129-47.

Berger, J. and G. Packard. 2018. Are Atypical Things More Popular?. *Psychological Science* 29(7): 1178-84.

Bielby, D.D and W.T. Bielby. 1996. Women and Men in Film: Gender Inequality among Writers in a Culture Industry. *Gender & Society* 10(3): 248-70.

Bielby, W.T. and D.D. Bielby. 1992. Cumulative Versus Continuous Disadvantage in an Unstructured Labor Market: Gender Differences in the Careers of Television Writers. *Work and Occupations* 19(4): 366-86.

Blau, F.D. and L.M. Kahn. 2000. Gender Differences in Pay. *Journal of Economic Perspectives* 14(4): 75-99.

Blei, D.M, A.Y. Ng and M.I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of machine Learning research* 3(Jan): 993-1022.

Botelho, T. and M. Abraham. 2017. Pursuing Quality: How Search Costs and Uncertainty Magnify Gender-Based Double Standards in a Multistage Evaluation Process. *Administrative Science Quarterly* 62(4): 698–730.

Bourdieu, P. 1993. *The Field of Cultural Production: Essays on Art and Literature*, Edited by R. Johnson. New York: Columbia University Press.

Bourdieu, P. 1996. *The Rules of Art: Genesis and Structure of the Literary Field*. Palo Alto, CA: Stanford University Press.

Brass, D.J. 1985. Men's and Women's Networks: A Study of Interaction Patterns and Influence in an Organization. *Academy of Management Journal* 28(2): 327-43.

Brown, T.W. 2019. Quality or Inequality?: A Machine Learning Approach to the Study of Value and Sexism in Creative Profession. Duke University, Durham, NC.

Burt, R.S. 1992. *Structural Holes. The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, R.S. 1998. The Gender of Social Capital. *Rationality and society* 10(1): 5-46.

Burt, R.S. 2004. Structural Holes and Good Ideas. *American Journal of Sociology* 110(2): 349-99.

Butler, J. 1988. Performative Acts and Gender Constitution: An Essay in Phenomenology and Feminist Theory. *Theatre Journal* 40(4): 519-31.

Cardador, M.T. 2017. Promoted up but Also Out? The Unintended Consequences of Increasing Women's Representation in Managerial Roles in Engineering. *Organization science* 28(4): 597-617.

Carli, L.L. 1989. Gender Differences in Interaction Style and Influence. *Journal of Personality and Social Psychology* 56(4): 565.

Castilla, E.J. 2011. Bringing Managers Back In: Managerial Influences on Workplace Inequality. *American Sociological Review* 76(5): 667-94.

Cech, E.A., B. Rubineau, S. Silbey and C. Seron. 2011. Professional Role Confidence and Gendered Persistence in Engineering. *American Sociological Review* 76(5): 641-66.

Cech, E.A. 2013. Ideological Wage Inequalities? The Technical/Social Dualism and the Gender Wage Gap in Engineering. *Social Forces* 91(4): 1147-82.

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

Childress, C. 2017. *Under the Cover: The Creation, Production and Reception of a Novel*, Edited by P. DiMaggio, M. Lamont, R. J. Wuthnow and V. A. Zelizer. Princeton, New Jersey: Princeton University Press.

Childress, C. and J-F. Nault. 2019. Encultured Biases: The Role of Products in Pathways to Inequality. *American Sociological Review*.

Cho, S.M. and M. Mauskapf. 2019. Unpacking the Ecological Dynamics of Cultural Products: Evidence from Popular Music. Columbia University Working Paper.

Christin, A. 2012. Gender and Highbrow Cultural Participation in the United States. *Poetics* 40(5): 423-43.

Cohen, L.E., J.P. Broschak and H.A. Haveman. 1998. And Then There Were More? The Effect of Organizational Sex Composition on the Hiring and Promotion of Managers. *American Sociological Review*: 711-27.

Cole, S. 1979. Age and Scientific Performance. *American Journal of Sociology* 84(4): 958-77.

Connell, R.W. [1995] 2005. *Masculinities*. Berkeley, CA: University of California Press.

D., Charles. 1872. *The Descent of Man, and Selection in Relation to Sex*, Vol. 2: D. Appleton.

Davies-Netzley, S.A. 1998. Women above the Glass Ceiling Perceptions on Corporate Mobility and Strategies for Success. *Gender & Society* 12(3): 339-55.

Doering, L.and S. Thébaud. 2017. The Effects of Gendered Occupational Roles on Men's and Women's Workplace Authority: Evidence from Microfinance. *American Sociological Review* 82(3): 542-67.

Dowd, T. J. 2002. The Musical Structure and Social Context of Number One Songs, 1955 to 1988: An Exploratory Analysis. Pp. 138-65 in *Vocabularies of Public Life*: Routledge.

Elsbach, K.D. and F.J. Flynn. 2013. Creative Collaboration and the Self-Concept: A Study of Toy Designers."*Journal of Management Studies* 50(4): 515-44.

England, P. 1992. *Comparable Worth: Theories and Evidence*: Transaction Publishers.

England, P. 2010. The Gender Revolution Uneven and Stalled."*Gender & Society* 24(2): 149-66.

Fekadu, M. 2018. Rise Up: Female Voices Take Center Stage at Grammys. in *Associated Press*: Associated Press.

Foschi, M. 2000. Double Standards for Competence: Theory and Research. *Annual Review of Sociology*: 21-42.

Foster, J.G., A. Rzhetsky and J.A. Evans. 2015. Tradition and Innovation in Scientists' Research Strategies. *American Sociological Review* 80(5): 875-908.

Friberg, A., E. Schoonderwaldt, A. Hedblad, M. Fabiani and A. Elowsson. 2014. Using Listener-Based Perceptual Features as Intermediate Representations in Music Information Retrieval. *The Journal of the Acoustical Society of America* 136(4): 1951-63.

Frith, S. 1996. Music and Identity. *Questions of cultural identity* 1: 108-28.

Goldin, C. and C. Rouse. 2000. Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians. *The American Economic Review* 90(4): 715-41.

Gorman, E.H. 2005. Gender Stereotypes, Same-Gender Preferences, and Organizational Variation in the Hiring of Women: Evidence from Law Firms. *American Sociological Review* 70(4): 702-28.

Gorman, E.H. 2006. Work Uncertainty and the Promotion of Professional Women: The Case of Law Firm Partnership. *Social Forces* 85(2): 865-90.

Greenberg, D.M., M. Kosinski, D.J. Stillwell, B.L. Monteiro, D.J. Levitin and P.J. Rentfrow. 2016. The Song Is You: Preferences for Musical Attribute Dimensions Reflect Personality. *Social Psychological and Personality Science* 7(6):597-605.

Guetzkow, J., M. Lamont and G. Mallard. 2004. What Is Originality in the Humanities and the Social Sciences?. *American Sociological Review* 69(2):190-212.

Hargadon, A. B. and B. A. Bechky. 2006. When Collections of Creatives Become Creative Collectives: A Field Study of Problem Solving at Work. *Organization Science* 17(4): 484-500.

Harrison, S., L. Hagvedt and N. Askin. 2019. The Early Success Trap: Group Creativity, Ambition, and the Grammy for Best New Artist. Working paper.

Harrison, S. and E.D. Rouse. 2014. Let's Dance! Elastic Coordination in Creative Group Work: A Qualitative Study of Modern Dancers. *Academy of Management Journal* 57(5): 1256-83.

Heilman, M/E. 2012. Gender Stereotypes and Workplace Bias. *Research in organizational behavior* 32: 113-35.

Hengel, E. 2017. Publishing While Female. Are Women Held to Higher Standards? Evidence from Peer Review. Cambridge Working Paper in Economics 1753.

Hennessey, B.A and T.M Amabile. 2010. Creativity. *Annual review of psychology (Print)* 61: 569-98.

Hesmondhalgh, D. 1999. Indie: The Institutional Politics and Aesthetics of a Popular Music Genre. *Cultural Studies* 13(1):34-61.

Hsu, G. 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.

Iacus, S.M., G. King and G. Porro. 2012. Causal Inference without Balance Checking: Coarsened Exact Matching. *Political Analysis* 20(1): 1-24.

Ibarra, H. 1992. Homophily and Differential Returns: Sex Differences in Network Structure and Access in an Advertising Firm. *Administrative Science Quarterly* 37(3): 422-47.

Ibarra, H. 1997. Paving an Alternative Route: Gender Differences in Managerial Networks. *Social Psychology Quarterly*: 91-102.

Ibarra, H. and J. Petriglieri. 2017. Impossible Selves: Image Strategies and Identity Threat in Professional Women's Career Transitions. in *Advancing Women and Leadership: Moving the Needle through Applied Theory Building*, edited by J. Storberg-Walker and P. Haber-Curran: Information Age Publishing.

Jones, B.F. and B.A. Weinberg. 2011. Age Dynamics in Scientific Creativity. *Proceedings of the National Academy of Sciences* 108(47): 18910-14.

Kanter, R.M. 1977. Some Effects of Proportions on Group Life: Skewed Sex Ratios and Responses to Token Women. *American Journal of Sociology* 82(5): 965-90.

Kanze, D., L. Huang, M.A. Conley and E.T. Higgins. 2017. We Ask Men to Win & Women Not to Lose: Closing the Gender Gap in Startup Funding. *Academy of Management Journal*.

Kaufman, J.C., J. Baer and C.A. Gentile. 2004. Differences in Gender and Ethnicity as Measured by Ratings of Three Writing Tasks. *The Journal of Creative Behavior* 38(1): 56-69.

Kay, F.M. and J. Hagan. 1998. Raising the Bar: The Gender Stratification of Law-Firm Capital. *American Sociological Review* 63(5): 728-43.

Kogan, N. 1974. Creativity and Sex Differences. *The Journal of Creative Behavior* 8(1): 1-14.

Koppman, S. 2014. Making Art Work: Creative Assessment as Boundary Work. *Poetics* 46(5): 1-21.

Koppman, S. 2016. Different Like Me: Why Cultural Omnivores Get Creative Jobs. *Administrative Science Quarterly* 61(2): 291-331.

Lagaert, S., M.V. Houtte and H. Roose. 2017. Engendering Culture: The Relationship of Gender Identity and Pressure for Gender Conformity with Adolescents' Interests in the Arts and Literature. *Sex roles* 77(7-8):482-95.

Lazer, D. and J. Radford. 2017. Data Ex Machina: Introduction to Big Data. *Annual Review of Sociology* 43:19-39.

Leahey, E. 2006. Gender Differences in Productivity: Research Specialization as a Missing Link. *Gender & Society* 20(6): 754-80.

Leahey, E. 2007. Not by Productivity Alone: How Visibility and Specialization Contribute to Academic Earnings. *American Sociological Review* 72(4): 533-61.

Leahey, E., C.M. Beckman and T.L. Stanko. 2017. Prominent but Less Productive: The Impact of Interdisciplinarity on Scientists' Research. *Administrative Science Quarterly* 62(1): 105-39.

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

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

Leslie, S-J., A. Cimpian, M. Meyer and E. Freeland. 2015. Expectations of Brilliance Underlie Gender Distributions across Academic Disciplines. *Science* 347(6219): 262-65.

Leung, M.D. 2014. Dilettante or Renaissance Person? How the Order of Job Experiences Affects Hiring in an External Labor Market. *American Sociological Review* 79(1): 136-58.

Lincoln, A.E., S. Pincus, J.B. Koster and P.S. Leboy. 2012. The Matilda Effect in Science: Awards and Prizes in the Us, 1990s and 2000s. *Social Studies of Science* 42(2): 307-20.

Lindemann, D.J., C.A. Rush and S.J. Tepper. 2016. An Asymmetrical Portrait: Exploring Gendered Income Inequality in the Arts. *Social Currents* 3(4): 332-48.

Lizardo, O. 2006. The Puzzle of Women's 'Highbrow' Culture Consumption: Integrating Gender and Work into Bourdieu's Class Theory of Taste. *Poetics* 34(1): 1-23.

Lopes, P.D. 1992. Innovation and Diversity in the Popular Music Industry, 1969 to 1990. *American Sociological Review*: 56-71.

Lutter, M. 2015. Do Women Suffer from Network Closure? The Moderating Effect of Social Capital on Gender Inequality in a Project-Based Labor Market, 1929 to 2010. *American Sociological Review* 80(2): 329-58.

Ma, Y., D.F.M. Oliveira, T.K. Woodruff, and B. Uzzi. 2019. Women Who Win Prizes Get Less Money and Prestige. *Nature Publishing Group*.

Maaten, L.vd. and G. Hinton. 2008. Visualizing Data Using T-Sne. *Journal of machine Learning research* 9(Nov): 2579-605.

Mauskapf, M., Eric Quintane, N. Askin and J. Mol. 2019. The Social Organization of Creativity in Music. Columbia University Working Paper.

McPherson, M., L. Smith-Lovin and J.M. Cook. 2001. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*: 415-44.

Meyer, M., A. Cimpian and S.-J. Leslie. 2015. Women Are Underrepresented in Fields Where Success Is Believed to Require Brilliance. *Frontiers in psychology* 6.

Miller, D.L. 2014. Symbolic Capital and Gender: Evidence from Two Cultural Fields. *Cultural Sociology* 8(4): 462-82.

Miller, D. L. 2016. Gender and the Artist Archetype: Understanding Gender Inequality in Artistic Careers. *Sociology Compass* 10(2): 119-31.

Oliveira, D.F.M., Y. Ma, T.K. Woodruff and B. Uzzi. 2019. Comparison of National Institutes of Health Grant Amounts to First-Time Male and Female Principal Investigators. *Jama* 321(9): 898-900.

Perrow, C. 1967. A Framework for the Comparative Analysis of Organizations. *American Sociological Review*: 194-208.

Peterson, R.A. and D.G. Berger. 1971. Entrepreneurship in Organizations: Evidence from the Popular Music Industry. *Administrative Science Quarterly*: 97-106.

Peterson, R.A. 1990. Why 1955? Explaining the Advent of Rock Music. *Popular music* 9(1): 97-116.

Peterson, R.A. and N. Anand. 2004. The Production of Culture Perspective. *Annual Review of Sociology*: 311-34.

Phillips, D.J. and E.W. Zuckerman. 2001. Middle-Status Conformity: Theoretical Restatement and Empirical Demonstration in Two Markets. *American Journal of Sociology* 107(2): 379-429.

Proudfoot, D., A.C. Kay and C.Z. Koval. 2015. A Gender Bias in the Attribution of Creativity Archival and Experimental Evidence for the Perceived Association between Masculinity and Creative Thinking. *Psychological Science* 26(11): 1751-61.

Ragins, B.R., B. Townsend and M. Mattis. 1998. Gender Gap in the Executive Suite: Ceos and Female Executives Report on Breaking the Glass Ceiling. *Academy of Management Perspectives* 12(1): 28-42.

Reskin, B.F. and P.A. Roos. [1990] 2009. *Job Queues, Gender Queues: Explaining Women's Inroads into Male Occupations*. Philadelphia: Temple University Press.

Richards, C. 2019. The Highwomen Are Raising Their Voices. Nashville Should Listen. in *The Washington Post*. Washington, D.C. .

Ridgeway, C.L. 2011. *Framed by Gender: How Gender Inequality Persists in the Modern World*. New York: Oxford University Press.

Rosen, S. 1981. The Economics of Superstars. *The American Economic Review* 71(5): 845-58.

Rosenblum, Barbara. 1978. "Style as Social Process." *American Sociological Review* 43:422-38.

Roth, L.M. 2006. *Selling Women Short: Gender and Money on Wall Street*: Princeton University Press.

Schmutz, V. 2009. Social and Symbolic Boundaries in Newspaper Coverage of Music, 1955–2005: Gender and Genre in the Us, France, Germany, and the Netherlands. *Poetics* 37(4): 298-314.

Schmutz, V. and A. Faupel. 2010. Gender and Cultural Consecration in Popular Music. *Social Forces* 89(2): 685-707.

Schmutz, V., E. Stearns and E.J. Glennie. 2016. Cultural Capital Formation in Adolescence: High Schools and the Gender Gap in Arts Activity Participation. *Poetics* 57: 27-39.

Seidel, M-D.L. and H.R. Greve. 2017. Emergence: How Novelty, Growth, and Formation Shape Organizations and Their Ecosystems. Pp. 1-27 in *Emergence*: Emerald Publishing Limited.

Smith, S.L., M. Choueiti, K. Pieper, H. Clark, A. Case, and S. Villanueva. 2019. Inclusion in the Reporting Studio? Vol. Los Angeles, CA: University of Southern California, Annenberg.

Stokes, A. 2015. The Glass Runway How Gender and Sexuality Shape the Spotlight in Fashion Design. *Gender & Society* 29(2): 219-43.

Tak, E., SJ. Correll, and S.A. Soule. 2019. Gender Inequality in Product Markets: When and How Status Beliefs Transfer to Products. *Social Forces*.

Thébaud, S. 2015. Business as Plan B: Institutional Foundations of Gender Inequality in Entrepreneurship across 24 Industrialized Countries. *Administrative Science Quarterly* 60(4): 671-711.

Tversky, A. 1977. Features of Similarity. *Psychological Review* 84(4): 327.

U.S. Department of Education. 2017. Gender Distribution of Bachelor's Degrees in the Humanities. in *Human Indicators*: American Academy of Arts and Sciences.

Uzzi, B. and J. Spiro. 2005. Collaboration and Creativity: The Small World Problem. *American Journal of Sociology* 111(2): 447-504.

Uzzi, B., S. Mukherjee, M. Stringer, and B. Jones. 2013. Atypical Combinations and Scientific Impact. *Science* 342(6157): 468-72.

Watson, J. 2019. Gender on the Billboard Hot Country Songs Chart, 1996–2016. *Popular Music and Society* 42(5): 538-60.

Werde, B. 2018. Werde: On The Grammys, the Music Industry & Progress." *Billboard*. Retrieved February 7, 2018 (https://www.billboard.com/articles/news/8098506/bill-werde-grammys-neil-portnow-tom-windish-la-reid-timesup).

Whittington, K.B. and L. Smith-Doerr. 2008. Women Inventors in Context: Disparities in Patenting across Academia and Industry. *Gender & Society* 22(2): 194-218.

Wu, L., D. Wang and J.A. Evans. 2019. Large Teams Develop and Small Teams Disrupt Science and Technology. *Nature* 566(7744): 378.

Yahr, E. 2018. Grammys President Says He Regrets Comment That Women Need to 'Step up' to Win More Awards. in *Washington Post*. Washington DC.

Yang, Y., N.V. Chawla, and B. Uzzi. 2019. A Network's Gender Composition and Communication Pattern Predict Women's Leadership Success. *Proceedings of the National Academy of Sciences* 116(6): 2033-38.

Zuckerman, E.W., T-Y. Kim, K. Ukanwa, and J. von Rittmann. 2003. Robust Identities or Nonentities? Typecasting in the Feature-Film Labor Market. *American Journal of Sociology* 108(5):1018-73.

Zukin, Sharon and Paul DiMaggio. 1990. *Structures of Capital: The Social Organization of the Economy*: CUP Archive.

Chapter 3

Ahn, Y.-Y, S.E. Ahnert, J.P. Bagrow, and A.-L. Barabási. 2011. Flavor Network and the Principles of Food Pairing. *Scientific Reports* 1.

Alexander, J. C. and P. Smith. 2002. The Strong Program in Cultural Theory: Elements of a Structural Theory. Pp. 135-150 in *Handbook of Sociological Theory*, edited by J. Turner. New York: Kluwer Academic/Plenum Publishers.

Anderson, J. R. 1991. The Adaptive Nature of Human Categorization. Psychological Review 98(3): 409-429.

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

Anand, N., and R.A. Peterson. 2000. When Market Information Constitutes Fields: Sensemaking of Markets in the Commercial Music Industry. *Organization Science* 11(3): 270–284.

Baetschmann, G., K.E. Staub, and R. 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, C.A. 2014. The Cultural Environment: Measuring Culture with Big Data. *Theory and Society* 43: 465–482.

Barroso, A., M. S. Giarratana, S. Reis, and O. Sorenson. 2014. Crowding, Satiation, and Saturation: The Days of Television Series' Lives: Crowding, Satiation, and Saturation. *Strategic Management Journal*.

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

Berger, J. and G.L. Mens. 2009. How Adoption Speed Affects the Abandonment of Cultural Tastes. *Proceedings of the National Academy of Sciences* 106(20): 8146–8150.

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

Bielby, W.T. and D.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, D. M., A.Y. Ng, and M. I. Jordan. 2003. “Latent Dirichlet Allocation.” *The Journal of Machine Learning Research* 3: 993–1022.

Bothner, M. S., J. Kang, and T. E. Stuart. 2007. Competitive Crowding and Risk Taking in a Tournament: Evidence from NASCAR Racing. *Administrative Science Quarterly* 52(2): 208–47.

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

Brackett, D. 1994. The Politics and Practice of “Crossover” in American Popular Music, 1963 to 1965. *The Music Quarterly* 78(4): 774–797.

Brashears, M.E. and E. Quintane. 2015. The Microstructures of Network Recall: How Social Networks are Encoded and Represented in Human Memory. *Social Networks* 41: 113–126.

Breiger, R.L. and K. 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, M.B. 1991. The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin* 17: 475–482.

Carroll, G.R., O.M. Khessina, D.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, R.E. 2000. Creative Industries: Contracts Between Art and Commerce. Cambridge, MA: Harvard University Press.

Cerulo, K. 1988. Analyzing Cultural Products: A New Method of Measurement. *Social Science Research* 17: 317-352.

Chan, C., J. Berger, and L. Van Boven. 2012. Identifiable but Not Identical: Combining Social Identity and Uniqueness Motives in Choice. *Journal of Consumer Research* 39(3): 561–573.

DellaPosta, D., Y. Shi, and M.W Macy. 2015. Why Do Liberals Drink Lattes? *American Journal of Sociology* 120(5): 1473–1511.

de Vaan, M., B. Vedres, and D. Stark. 2015. Game Changer: The Topology of Creativity. *American Journal of Sociology* 120(4):1144–94.

Dhanaraj, R. and B. Logan. 2005. Automatic Prediction of Hit Pop Songs. *Proceedings of the International Society for Music Information Retrieval*: 488-491.

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

DiMaggio, P., M. Nag, and D. 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, M. and B.C. Isherwood. 1996. *The World of Goods: Towards an Anthropology of Consumption*. New York: Routledge.

Dowd, T.J. 2004. Concentration and Diversity Revisited: Production Logics and the U.S. Mainstream Recording Market, 1940-1990. *Social Forces* 82(4): 1411–1455.

Eastman, J.T. and T.F. Pettijohn II. 2014. 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*.

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

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

Fleming, L., S. Mingo, and D. Chen. 2007. Collaborative Brokerage, Generative Creativity, and Creative Success. *Administrative Quarterly* 52: 443–475.

Fligstein, N. and L. Dauter. 2007. The Sociology of Markets. *Annual Review of Sociology* 33(1): 105–28.

Foster, J.G., A. Rzhetsky, and J.A. Evans. 2015. Tradition and Innovation in Scientists' Research Strategies. *American Sociological Review*.

Friberg, A., W. Schoonderwaldt, A. Hedblad, M. Fabiani, and A. Elowsson. 2014. Using Listener-based Perceptual Features as Intermediate Representations in Music Information Retrieval. *Journal of the Acoustical Society of America* 136(4): 1951–1963.

Frith, S. 1996. Music and Identity. In S. Hall and P. Du Gay (eds.), *Questions of Cultural Identity*. Pp. 108–125. London: Sage.

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

________. 2006. *Genre: The New Critical Idiom*. New York: Taylor and Francis.

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

Goldberg, A.., M.T. Hannan, and B. Kovács. 2015. What Does It Mean to Span Cultural Boundaries. *American Sociological Review*.

Hadida, A. 2015. Performance in the Creative Industries. In *The Oxford Handbook of Creative Industries*, edited by C. Jones, M. Lorenzen, and J. Sapsed. New York: Oxford University Press.

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

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

Heckman, J. 1979. Sample selection bias as a specification error. *Econometrica* 47 (1): 153–61.

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

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

Hsu, G. 2006. Jacks of all Trades and Masters of None: Audiences' Reactions to Spanning Genres in Feature Film Production. *Administrative Science Quarterly* 51(3): 420–450.

Hsu, G. and M.T. Hannan. 2005. Identities, genres, and organizational forms. *Organization Science* 16(5): 474–490.

Hsu, G., G. Negro, and F. Perretti. 2012. Hybrids in Hollywood: A Study of the Production and Performance of Genre Spanning Films. *Industrial and Corporate Change* 21: 1427-1450.

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

Jones, C., M. Maoret, F.G. Massa, and S. Svejenova. 2012. Rebels with a Cause: Formation, Contestation, and Expansion of the De Novo Category "Modern Architecture," 1870–1975. *Organization Science* 23(6): 1523–1545.

Kahl, S.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. Emerald Group Publishing Limited.

Karaca-Mandic, P, E.C. Norton, and B. Dowd. 2012. Interaction Terms in Nonlinear Models. *Health services research* 47(1pt1):255–74.

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

Kaufman, J. 2004. Endogenous Explanation in the Sociology of Culture. *Annual Review of Sociology* 30: 335–357.

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

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

Kovács, B. and R. Johnson. 2013. Contrasting Alternative Explanations for the Consequences of Category Spanning: A Study of Restaurant Reviews and Menus in San Francisco. *Strategic Organization*.

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

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

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

La Rue, J. 2001. Fundamental Considerations in Style Analysis. *The Journal of Musicology* 18(2): 295-312.

Lena, J.C. 2006. Social Context and Musical Content of Rap Music, 1979-1995. *Social Forces* 85(1): 479–495.

_________. 2012. *Banding Together: How Communities Create Genres in Popular Music*. Princeton: Princeton University Press.

_________. 2015. Genre: Relational Approaches to the Sociology of Music. In *Routledge International Handbook of the Sociology of Art and Culture*, edited by Laurie Hanquinet and Mike Savage. Pp. 149–160. New York: Routledge.

Lena, J.C. and R.A. Peterson 2008. Classification as Culture: Types and Trajectories of Music Genres. *American Sociological Review* 73(5): 697-718.

Lena, J.C. and M.C. Pachucki. 2013. The Sincerest Form of Flattery: Innovation, Repetition, and Status in an Art Movement. *Poetics* 41: 236–264.

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

Lizardo, O. 2006. How cultural tastes shape personal networks. *American Sociological Review* 71: 778-807.

________. 2014. Omnivorousness as the Bridging of Cultural Holes: A Measurement Strategy. *Theory and Society* 43: 395–419.

Lopes, P.D. 1992. Innovation and Diversity in the Popular Music Industry, 1969 to 1990. *American Sociological Review* 57(1): 56–71.

Lounsbury, M. and M.A. Glynn. 2001. Cultural Entrepreneurship: Stories, Legitimacy, and the Acquisition of Resources. *Strategic management journal* 22(6-7): 545–64.

Mark, N. 1998. Birds of a Feather Sing Together. *Social Forces* 77(2): 453–485.

_______. 2003. Culture and Competition: Homophily and Distancing Explanations for Cultural Niches. *American Sociological Review* 68(3): 319–345.

Meng, T. and M-L Shyu. 2012. Leveraging Concept Association Network for Multimedia Rare Concept Mining and Retrieval. *IEEE International Conference on Multimedia and Expo (ICME) Proceedings*: 860–865.

Mohr, J.W. 1994. Soldiers, Mothers, Tramps, and Others: Discourse Roles in the 1907 New York City Charity Directory. *Poetics* 22: 327–357.

________. 1998. Measuring Meaning Structures. *Annual Review of Sociology* 24: 345–370.

Mood, C. 2010. Logistic Regression: Why We Cannot Do What We Think We Can Do, and What We Can Do About It. *European Sociological Review* 26(1):67–82.

Morgan, S.L. and C. Winship. 2007. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge, UK: Cambridge University Press.

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

Nickell, S. 1981. Biases in Dynamic Models with Fixed Effects. *Econometrica: Journal of the Econometric Society*: 1417–26.

Nunes, J.C. and A. Ordanini. 2014. I Like the Way it Sounds: The Influence of Instrumentation on Pop Song's Place in the Charts. *Musicae Scientiae*: 1–18.

Pachucki, M.A. and R.L. Breiger. 2010. Cultural Holes: Beyond Relationality in Social Networks and Culture. *Annual Review of Sociology* 36: 205-224.

Pachet, F. and P. Roy 2008. Hit Song Science is Not Yet a Science. *Proceedings of the International Society for Music Information Retrieval*: 355-360.

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

___________. 1997. *Creating Country Music: Fabricating Authenticity*. Chicago: University of Chicago Press.

Peterson, R.A. and D.C. Berger. 1975. Cycles in Symbol Production: The Case of Popular Music. *American Sociological Review* 40(2): 158–173.

Phillips, D.J. 2011. Jazz and the Disconnected: City Structural Disconnectedness and the Emergence of a Jazz Canon, 1897–1933. *American Journal of Sociology* 117(2): 420–483.

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

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

Podolny, J.M., T.E. Stuart, and M.T. Hannan.1996. Networks, Knowledge, and Niches: Competition in the Worldwide Semiconductor Industry. *American Journal of Sociology* 102(3): 659–689.

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

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

Rao, H., P. Monin, and R. 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, S. 1981. The Economics of Superstars. *The American Economic Review* 71(5): 845-858.

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

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

Rubio, F.D. 2012. The Material Production of the Spiral Jetty: A Study of Culture in the Making. *Cultural Sociology* 6(2): 143-161.

Salganik, M.J., P.S. Dodds, and D.J. Watts. 2006. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311: 854-856.

Salganik, M.J. and D.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.

________________. 2009. Web-Based Experiments for the Study of Collective Social Dynamics in Cultural Markets. *Topics in Cognitive Science* 1(3): 439–68.

Schafer, J. and K. 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).

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

Serrà, J., Á. Corral, M. Boguñá, M. Haro, and J.L. Arcos. 2012. Measuring the Evolution of Contemporary Western Popular Music. *Scientific Reports* 2.

Shi, Y., O. Sorenson, and D.M. Waguespack. The Stability of Centrality-Based Competitive Advantage. *Sociological Science* 4: 107–122.

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

Sonnett, J. 2004. Musical Boundaries: Intersections of Form and Content. *Poetics* 32(3-4): 247-264.

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

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

Teng, C.-Y., Y.-R. Lin, and L.A. Adamic. 2012. Recipe Recommendation Using Ingredient Networks. Pp. 298–307 in *Proceedings of the 3rd Annual ACM Web Science Conference*. ACM.

Thompson, D. 2014. The Shazam Effect. *The Atlantic Magazine*. December.

Tversky, A.. 1977. Features of Similarity. Psychological Review 84(4): 327-352.

Uzzi, B. and J. Spiro. 2005. Collaboration and Creativity: The Small World Problem. *American Journal of Sociology* 111: 447-504.

Uzzi, B., S. Mukherjee, M. Stringer, and B. Jones. 2013. Atypical Combinations and Scientific Impact. *Science* 342: 468–472.

Vaisey, S. and O. Lizardo. 2010. Can Cultural Worldviews Influence Network Composition? *Social Forces* 88(4): 1595-1618.

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

_________. 1991. *Joel Whitburn's Top Pop Singles, 1955-1990: Compiled from Billboard's Pop Singles Charts, 1955-1990*. Record Research.

Wijnberg, N.M. and G. Gemser. 2000. Adding Value to Innovation: Impressionism and the Transformation of the Selection System in Visual Arts. *Organization Science* 11(3): 323-329.

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

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

Zerubavel, N., P.S. Bearman, J. Weber, and J.N. N. Ochsner. 2015. Neural Mechanisms Tracking Popularity in Real-World Social Networks. *Proceedings of the National Academy of Sciences (PNAS)* 112(49): 15072–15077.

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

_________. 2004. Structural Incoherence and Stock Market Activity. *American Sociological Review* 69: 405–32.

_________. 2016. Optimal Distinctiveness Revisited: An Integrative Framework for Understanding the Balance between Differentiation and Conformity in Individual and Organizational Identities. In *Handbook of Organizational Identity*. Oxford, UK: Oxford University Press.