Network Boundedness as Market Identity

Evidence from the Film Industry *

Amir Goldberg** and Anthony Vashevko**

**Stanford Graduate School of Business

July 30, 2013

Abstract

Network scholars have studied organizational creativity predominantly as a problem of acquiring and integrating information. In contrast, we re-conceptualize the structural tradeoff between brokerage and closure as an identity signal. With Hollywood as our empirical setting, we demonstrate that consumer perceptions of films are structured by a strong status hierarchy, and that boundedness – the extent to which a film's production team comprises an exclusive clique within the network of interpersonal collaborations in the film industry – serves as a signal of artistic quality. We draw on a uniquely detailed dataset of consumer preferences, and take advantage of the lag between film production and consumer evaluation as a means to demonstrate that production team members' career trajectories after a film had been produced have a bearing on audiences' evaluations. Films whose team members went on to collaborate in exclusive circles and thereby, we argue, establishing a high-status identity, tend to enjoy greater post-hoc artistic appreciation than at the time of their release.

*Please do not cite without permission. Direct all correspondence to Amir Goldberg, Stanford Graduate School of Business, Knight Management Center, 655 Knight Way, Stanford, CA 94305 (amirgo@stanford.edu).

Introduction

How do networks affect the quality of creative output in organizations? In answering this question, organizational scholars have predominantly focused on how networks afford access to diverse knowledge, and the means by which networks facilitate, or inhibit, the ability to integrate and act upon new information (Hansen 1999, Burt 2004, Uzzi & Spiro 2005, Vedres & Stark 2010, Aral & Alstyne 2011, Tortoriello & Krackhardt 2010). But network positions also confer identity upon those who occupy them. This is consequential for creativity because how actors are perceived by others has a bearing on whether or not their creations are interpreted as novel, or incompetent (Phillips & Zuckerman 2001, Rao, Monin & Durand 2005). It is one thing being creative, and another being acknowledged as such.

We argue that the extent to which an actor operates in a densely connected circle of relationships – a property which we refer to as boundedness – affects how its output is perceived. This is the case for two reasons. First, operating within a clearly bounded group makes it easier for outside observers to assign an unambiguous identity to an actor (White 1992, DiMaggio 2011). Second, group boundaries project exclusivity; an actor's ability to maintain them signals its quality (Gieryn 1983, Podolny 1994, Lamont & Molnár 2002).

Our empirical setting is the film industry, where products are judged on their creative quality, and produced by ad-hoc teams whose members often repeatedly collaborate with one another (Faulkner & Anderson 1987). We demonstrate that the tradeoff between network closure and brokerage – which animates a large body of literature on creativity in markets (Burt 2005) – is consequential for market identity. Whereas previous work has examined this tension exclusively through a material lens, whereby networks are understood as conduits upon which ideas travel and recombine, we posit that a production team's network boundedness can beneficially impact a film's reception by audiences irrespective of the material advantages (or disadvantages) it affords. Network boundedness, in other words, mediates the relationship between a film's attributes, and how audiences perceive them.

Measuring audience perceptions, however, can be tricky. Existing scholarship tends to infer what goes on in people's heads from aggregate measures such as price, market share or popularity. Yet such aggregations often obscure the diversity of preferences that underlie them. In contrast, this study relies on ratings of roughly three thousand films produced between 1920 and 2005, and provided by almost half a million users of the online video rental service *Netflix*. Rather than simply comparing films by their overall success or popularity, we employ a relational approach that distinguishes between films as a function of the different audience segments they appeal to. We find that audience evaluations are shaped by a highly structured hierarchy that ranges from sophisticated to common films. We demonstrate that the maintenance of identity and exclusivity through boundary work provides a production team, and the film it produces, with an aura of creative sophistication.

We draw on a unique feature of our data: that there exists high variability in the lag between a film's time of production and its evaluation by *Netflix* users. We use this lag as an instrument to demonstrate that the effect of boundedness on audience evaluation is, at least in part, attributable to the film's team members' career trajectories after the film had been released. Films whose team members went on to collaborate in exclusive circles and thereby, we argue, establishing a high-status identity, tend to enjoy greater post-hoc artistic appreciation than at the time of their release.

Our study integrates insights from two otherwise tangential organizational literatures on identity and networks, demonstrating how identity, and the status implications it affords, is rooted in network position. In Podolny's (2001) language, it highlights how networks function as prisms of the market. Whereas previous work in this vein has focused exclusively on the dyadic level – demonstrating that status is conferred through association with reputable others (Rossman, Esparza & Bonacich 2010, Stuart, Hoang & Hybels 1999) – we illustrate how network boundedness, or lack thereof, functions as a signal about an actor's underlying quality.

The remainder of the text proceeds as follows. In the following section, we discuss identity as an analytical construct, and theorize about its relationships with network position. We then introduce our method for inferring identity, and describe our data and analytical strategy in depth. We present the results in the following section, and end with a discussion on the implications and generalizability of our findings.

Theory

Where Does Identity Come From?

As an analytical concept, identity is fraught with ambiguity (Brubaker & Cooper 2000). We use the term to denote *identification*, the process by which an object is classified as an instantiation of a generic type. Whereas a large body of extant literature examines how organizational self-identity shapes members' understandings of their organization's objectives and normative ways of doing things (e.g. Whetten 2006, Gioia, Schultz & Corley 2000), we conceptualize identification as a process exerted on an actor by others (Hsu & Hannan 2005). In market contexts, these others – which are often referred to in the literature as "audience" – might be customers, partners, or other constituencies outside the organization. Such audiences command control over resources that are essential for an organization's success (Zuckerman 1999).

Identity, therefore, is consequential because it inheres in audiences' perceptions, rather than an organization's attributes per-se. It is the taken for granted assumptions that audiences make about an organization, and the products it produces (Hannan, Pólos & Carroll 2007). Identity comes into play when audience members make inferences about an organization's characteristics in the absence of clear-cut information about these characteristics. It is therefore especially significant in contexts that are characterized by what Podolny (2001) calls "altercentric uncertainty:" the uncertainty an evaluator has about the quality of the object being evaluated. Under such circumstances, identity functions as a signal that substitutes for missing information about the objective quality or worth of an actor's output.

The market for consumer electronics, for example, is rife with altercentric uncertainty. Consider the iPhone, Apple Inc's highly popular smartphone introduced in 2007. Consumers assessing the relative advantages and disadvantages of purchasing an iPhone need to take into consideration its functional quality as well as the social status associated with displaying it in public. Apple's identity therefore presumably plays a role in shaping these audiences' assessments. Understood primarily as a computer hardware manufacturer, potential customers might be skeptical about Apple's ability to successfully foray into the domain of cellular telephony. Yet if Apple's most salient characteristic in the eyes of consumers is its flare for innovation, they might – as they actually have in droves – be inclined to interpret its technology as groundbreaking, and its public consumption as carrying a cachét, and concomitantly elect to buy an iPhone.

Similar problems present themselves in the market for feature films. Not only is there rarely consensus, whether among experts or the lay public, on the formal features that distinguish valuable from low quality art (Yogev 2010); like the iPhone, film is often conspicuously consumed as a social status signal (DiMaggio 1987, Baumann 2007). A film like the 2009 *Inglourious Basterds* – a revenge fantasy depicting an alternate history of World War II – might be interpreted as a distastefully inane war film, or a subversively profound observation about the character of human violence. The categorical prism through which the film is perceived has a bearing on its reception by audiences.

We draw on Brubaker & Cooper (2000) to make a distinction between two primary sources of identity: *categorical* and *structural*.[1] By categorical identity, we refer to an organization's identification as a function of its attributional sameness with other organizations that constitute an institutionalized category. Structural identity, in contrast, inheres in the set of relationships an organization is embedded in, and the social meanings these relationships imply to the audience in question.

Most organizational scholarship conceptualizes identity through a categorical lens (e.g. Bielby & Bielby 1994). In particular, a variety of studies demonstrate that straddling multiple categories comes at a significant cost (Zuckerman 1999, Ruef & Patterson 2009, Hsu, Hannan & Koçak 2009). Actors who send mixed signals about what they do, are not

[1]Brubaker & Cooper (2000) originally refer to the latter as 'relational' identity. Because this term connotes dyadic relationships, and because we are focusing on an actor's location within a network topology, we prefer using the term 'structural.'

easy to make sense of. Consequently, they are overlooked by potential audience members or, at worst, perceived as inferior to actors with more specialized identities. A war film that does not adhere to particular thematic and aesthetic conventions, to apply this logic, is likely to be devalued by audiences (Hsu 2006).

Though usually less explicit about its use of the term, a tangential strand of scholarship conceptualizes identity in structural terms, as an actor's position within a set of relationships. This line of work shifts focus from what an actor does, to whom it does it with (Roberts, Khaire & Rider 2011). Studies in this vein have been predominantly concerned with demonstrating how status rubs off through networks. Examining a variety of empirical settings, ranging from the biotechnology industry (Stuart, Hoang & Hybels 1999) to Hollywood (Rossman, Esparza & Bonacich 2010), these studies show that actors who are associated with prominent others tend to reap rewards they would otherwise have little access to (for a review, see Sauder, Lynn & Podolny 2012).

Boundedness as Identity

Identity, in other words, leaks through network ties. If that is the case, then the overall set of relationships an actor is embedded in should have an impact on how this actor is perceived by others. We argue that whether or not an actor is located within a cohesive group of relationships is particularly consequential for identity. The idea that network cohesion facilitates social identity has a long sociological history, and has been demonstrated in a variety of contexts, especially, though not exclusively, within formal organizations (Podolny & Baron 1997, Krackhardt 1999). Whereas such studies have mostly focused on how dense networks generate a sense of belongingness and commitment among their members, we suggest that network closure also establishes identity in the eyes of outside observers.

Key to this process is the demarcation of a boundary. First, actors that are enclosed within a densely connected network of relationships are more easily identifiable as belonging to a distinct group (Emirbayer 1997, White 1992). Such groupness makes them recognizable to others. This is because people identify objects through typification. Because the cognitive process of classification is inherently fuzzy, differences between types of objects are not always readily observable (Rosch 1978). Perceived boundaries therefore make it easier for people to discriminate between the objects they observe. A group of film professionals repeatedly collaborating with one another are more evidently classifiable as belonging to the same group than individuals who collaborate with different people each time.

Second, meaningful relationships are costly to engage in. It is not enough for a film producer to want to collaborate with a celebrated director. This desire needs to be reciprocated in order for the relationship to materialize. Because the decision to collaborate depends on the mutual assessment of each others' capabilities, collaborative relationships also function as "gestures of approval," and as such convey information about parties' underlying qualities (Gould 2002). Actors consequently have strong incentives to associate with well-regarded others and distance themselves from those of more dubious reputations (Podolny 1994, Bothner, han Kang & Stuart 2007). This maintenance of exclusivity is what Gieryn (1983) calls "boundary-work:" actors' attempts to reify their sameness by creating and preserving a perceived boundary between in-group members and out-group others. Such boundaries often serve to sustain social distance between the groups they distinguish, and reinforce the material and symbolic inequalities this distance entails. It is therefore in the interest of high status actors not only to distance themselves from those lower on the food chain but also to objectify this distance as an impermeable boundary.

These two mechanisms – the identifiability of groupness, and the maintenance of affiliational exclusivity – inform our main argument: that network *boundedness* functions as a marker of quality in creative markets. We prefer describing network cliqueness as boundedness, rather than using more conventional terms such as closure or cohesion, because we want to emphasize how it serves to create perceived boundaries between groups of actors (be them individuals or organizations). We expect to find that, in the film industry, production teams that bridge structural holes are at a disadvantage. This disadvantage is the structural analog to the problem of multiple category membership: actors who are embedded in structurally diverse and non-exclusive relationships emit incoherent signals about who they are, thereby undercutting their perceived quality.

Boundedness in the Feature Film Market

The film industry is a fitting empirical setting for testing the relationship between boundedness and perceived quality for several reasons. First, films are produced by teams comprised of multiple individuals who occupy clearly defined roles. In essence, each film is an ad-hoc organization that is formed, and later dissolved, merely for the purpose of creating one product. Far from random, however, repeat collaboration is very common in these temporary project-based teams (Faulkner & Anderson 1987). Over time, they create an ecology of makeshift organizations in which individual careers crisscross one another. These intersecting career paths form an evolving network structure. Thus the film industry provides an ideal setting for tracing organizations, their outputs, and the social networks in which they are embedded.[2]

Second, because of their ad-hoc nature, film-based organizations do not have a singular history. Almost by definition, audience expectations for novelty and creativity preclude the exact same production team from regrouping. Though repeat collaborations are very common, even sequels rarely feature identical teams. Thus, unlike with conventional firms, audiences rarely have information on a newly formed film crew's past performance as a team. They are therefore pushed to rely on identity as a signal about the potential quality of this production team's output.

Finally, the boundary separating authentic artwork from mere craft is particularly potent in art markets (Becker 1982, Bourdieu 1979). Though a relatively young form of creative expression – roughly a century old – American cinema underwent an interpretative transformation during the 1960s. The consecration of certain cinematic works as "art"

[2] The majority of films are produced under the umbrella of a film studio which is an organizational entity that outlives the duration of one cinematic project. Yet the core creative team, comprising such functionaries as the director, producer, actors and screenwriters, tends to vary in composition between one film and the other. During its first four decades, the film industry was dominated by the studio system, whereby films were produced almost exclusively by major studios, and team members were employed by these studios under long term contracts. During that period, studios exercised greater command over the production process than they do today. Nevertheless, production team composition was never fixed, and increasingly so after the studio system's demise in the early 1960s (Turner 1999).

effectively created a status hierarchy whereby some films were hailed as masterpieces and their creators as “auteurs” (Baumann 2007). Ever since, the common distinction between art house films on the one hand, and lowest common denominator blockbusters on the other, is popular among industry and audience members alike (King 2009).[3] It reinforces a perception whereby films are divided into two broad qualitative categories.

We conceptualize boundedness as the extent to which the individuals who comprise a film production team have collaborated with one another on other films. Highly bounded films are those whose members collaborated with one another on multiple films, whereas low bounded films are those whose members participated in multiple films, but who have collaborated with one another only on one. Our argument rests on the assumption that boundedness is related to how audiences infer identity from their knowledge about team members' past collaborations.

Yet boundedness might also be related to the conditions under which a film is produced, regardless of eventual audience perceptions. Existing literature generally offers two competing hypotheses about how boundedness might indirectly shape audience reception through its effects on production. The first argues that structural consistency is favorably related to performance (Hannan & Freeman 1984) and that network closure in particular facilitates innovation (Ahuja 2000). Repeat interaction is especially conducive to high quality artistic outcomes through the formation of what Becker (1982) calls “artworlds:” communities of artists, industry professionals and avid consumers who collectively negotiate the content of different forms of artistic expression. These forms cohere through the process of interpersonal interaction, such that tight-knit circles of creative collaboration and exchange facilitate the emergence of distinct styles (Crossley 2009, Lena 2012).

Yet a variety of other studies find that boundedness can constrain artistic performance. These rest on the idea that artistic creativity is highly valued by audiences, and that creative novelty emerges at the intersections of otherwise disconnected parts of a production network (Burt 2004). In a study of Italian TV production teams, for example, Zaheer &

[3]Consider, for example, celebrated auteur Robert Altman, who, in lamenting the state of American film quips that “the artists have left it, and it is run by bookkeepers and insurance people” (Turner 1999, p. 11).

Soda (2009) find that teams whose members spanned structural holes were more likely to succeed commercially. Uzzi & Spiro (2005) similarly find that creativity in the Broadway musical industry peaks in networks that are, on the one hand, sufficiently bounded to facilitate collaborative exchange but at the same time also porous enough to allow new and varied ideas to flow from outside these closed circles.

We argue that boundedness has a an effect on film audience perceptions net of its production-side effects on the quality and characteristics of cinematic output. As we discuss below, our research design enables us to separate production from consumption processes, and isolate the effect of boundedness on audience perceptions.

Measuring Identity

An analytical focus on identity calls for an empirical focus on audience perceptions. Surprisingly, however, the vast literature on identity in markets, almost without exception, uses aggregate measures as proxies for audience interpretations. Often, price or market share are assumed to represent consumers' revealed preferences. A film's gross income is understood as a product of its categorical coherence (Hsu, Hannan & Koçak 2009) or an investment bank's profit margin a result of its status (Podolny 1993). While these studies make strong arguments about what goes on in audiences' heads, these aggregate measure may be obfuscating, rather than elucidating, audience interpretations.

Consider the following thought experiment as illustration. Imagine that a market contains only two similarly priced products, a and b, with 30% and 70% market share respectively. Imagine furthermore that we have two competing hypotheses about the sources of these products' appeal to consumers. The first hypothesis is, simply, that product b is objectively better than a, and that audiences are in agreement about this valuation, but are 30% likely to make an erroneous evaluation. Thus we should expect to find that all consumers are universally 30% likely to purchase product a and 70% likely to purchase product b. The second hypothesis, in contrast, is that each product appeals to different audience segments. Thus, 30% of customers are inclined to purchase product a whereas 70% are inclined to purchase b. By examining aggregate sales we cannot adjudicate between these two theories, as both conditions result in exactly the same aggregate market shares for each product.

This problem can be generalized to our setting. Film theater tickets, as well as video and DVD rental or internet streaming costs, are more or less constant within geographical regions. Different films therefore do not compete with one another on cost; rather, their likelihood of success rests almost exclusively on their perceived quality. Yet, examining films' aggregate appeal or market share, for the reasons discussed above, provides no insight as to whether audiences are in consensus about their qualitative rankings, or whether different audience segments specialize in different types of films. It is therefore impossible to assert how identity is affecting audience valuations.

Consequently, previous scholarship relies on strong assumptions about how audience members perceive market products. Studies that conceptualize identity as categorical implicitly assume that audiences draw on the standard classificatory schemes that are used by institutional actors; that consumers, for example, specialize in particular film genres.[4] Work on structural identity assumes that status is interpreted as quality, and that audiences are in agreement about this interpretation.

An alternative approach is to infer inductively, rather than presuppose, how audiences typify objects, by examining their explicit or revealed preferences individually. Identity, as we discuss earlier, inheres in an observer's attribution of sameness between objects. Two organizations or products have the same 'identity' because, in the eyes of the audience, they are instantiations of the same type. We can use individual audience members' preferences as a means to infer this perceived sameness. Such an approach conceptualizes these perceived types as latent semantic anchors in a structure of relationships between objects. These anchors serve as the prototypes against which identity is measured (Rosch 1978, Murphy 2004). An object's identity is therefore a function of its location in this web of associations

[4] Recent work in this vein questions the assumption that audiences' perceptions comport to these standard classifications. In a study of technology startups, Pontikes (2012), for example, infers categorical identities from these organizations' descriptions. Though far more nuanced than traditional reliance on formal taxonomies, this approach nevertheless still assumes that organizations' identity claims shape audience perceptions.

(Lévi-Strauss 1963, Emirbayer 1997, Mohr 1998).

Consider *Inglourious Basterds* again. Critics' initial reactions to this fictional portrayal of World War II were mixed. Whereas some saw it as a hyper-violent war film, an inane and "embarrassing ... revenge fantasy," others found it profound and audacious in a manner that "resists categorization."[5] Imagine that we were to ask viewers to name other films that are like *Inglourious Basterds*. The former interpretation would likely result in an association with similarly violent war films; the latter, on the other hand, with a thematically diverse set of movies that ostensibly possess a similar perceived cinematic profundity. If identity relates to the salient dimensions along which people discriminate between different objects, then we should expect audiences evaluations of these objects to be correlated with one another as a function of their assignment of identity.

Our data include *Netflix* users' express opinions about the various films they watched, ranging on a 5-point Likert scale. This dataset differs from traditional survey based data in two important ways: (1) individuals vary significantly in the number of films they rate,[6] and (2) in the subsets of films they rate. Consequently, whereas some films were rated by many users, the majority were rated by a relative few. We construct a measure of similarity between films as a function of the correlation between the ratings of the users who rated both films. This correlation measures the extent to which two films were liked, and disliked, by the same people. We argue that this web of inter-film correlations reflects an underlying intersubjectively shared attribution of identity. A film's location within this network corresponds to its perceived identity.

Formally, let R^i denote the set of users who have ranked film i, $R^{i \cap j}$ denote the set of users who have ranked both films i and j, and $R_i^{i \cap j}$ denote the rankings received by film i from users who have ranked both films i and j. The measure of association between two films i and j is calculated as the correlation coefficient between their two corresponding sets of rankings, formally:

[5] see http://www.rogerebert.com/reviews/inglourious-basterds-2009

[6] The number of films rated follows a heavy-tailed distribution, with the median user rating 96 films, and the average user 209 films.

$$A_{ij} = \rho(R_i^{i\cap j}, R_j^{j\cap j}) \quad (1)$$

A is an NxN (N = number of films) matrix. This matrix can be thought of as a non-directed, weighted network. It can be transformed into a sparse network by removing all non-significant correlations.

This approach has several advantages. First, rather than looking at consumer behaviors in the aggregate, it examines consumers' tastes as they relate to one another through individual preferences; namely, it does not merely take into account whether two films are similarly appealing on average, but rather whether they are appealing (or unappealing) to the same people. Second, unlike most work on the topic, it does not essentialize categorical boundaries a-priori but rather assumes these boundaries are implicit in the structure of relationships between films. Finally, this approach examines how films relate to one another, and allows for instances of negative correlation, where there exists significant identity dissimilarity between films.

Data and Analytical Strategy

Data

In 2006, *Netflix* released a large dataset of user ratings as part of a competition titled the *Netflix Prize*.[7] The data include raw rankings of 17,770 titles produced by 480,189 users. Users are identified by random unique numbers; no additional user-related information is provided. Each observation includes a rating, timestamp, user ID, free text title, and year of production. Overall the data include roughly 100 million unique user-title pairs.

The almost eighteen thousand unique titles were linked with their corresponding entries on *Netflix*'s website, as well as the publicly available *Internet Movie Database* (IMDB), which includes rich production- and distribution-related information as well as genre assignments and aggregate user ratings.[8] Relying on IMDB's *type* classification, the dataset

[7] The purpose of the competition, bearing a $1 million prize, was for contestants to improve Netflix's movie recommendation algorithm by at least 10% of accuracy.

[8] Of the 17,770 titles, 81 could not be found on Netflix's website. These account for 0.03% of all obser-

was cleaned to include only motion pictures (as opposed to TV shows, video movies and other non-film products). Duplicate titles (such as different editions of the same film) were also removed, resulting in 9,845 films, ranked by 479,578 unique users and overall encompassing more than 91 million rankings. The data were also linked with the movie review aggregation website *Rotten Tomatoes*, where each film is assigned a score on a scale of 0 to 100 which is based on the aggregate reviews of U.S. film critics. Films with high scores are those that enjoy great critical acclaim.

To make the data manageable, the dataset was reduced to include the films that were rated by the highest number of users, overall corresponding to 95% of all movie ratings included in the dataset.[9] The resulting dataset, which was used in the following analyses, comprises 2,876 movie titles, rated by 479,087 unique users. It includes films released between 1920 and 2005 (the majority after 1995), mostly produced in the U.S. but not exclusively. The movie with the largest number of user ratings, *Miss Congeniality* (released in 2000), was rated by more than 230,000 users, whereas the least rated film in the dataset, the 1954 biographic drama *The Glenn Miller Story* received 4,808 ratings. Year is significantly correlated with number of ratings (yet explains less than 5% of variance) and with a film's average rating.[10] The most liked film in the dataset is the 2001 fantasy action movie *The Lord of the Rings* (the first of a highly successful trilogy), with an average rating of 4.72, and the least liked is the 2003 *Gigli*, with an average rating of 1.94. With more than $871 million in gross earnings, *The Lord of the Rings* is the tenth highest grossing film released before 2006, whereas *Gigli* was a box-office failure, netting almost $70 million in losses. Overall, Netflix users' ratings, both in volume and average score, correspond to movies' financial success and do not seem to suggest that this sample represents an audi-

vations in the dataset, and were removed from the dataset. For more information about the data retrieval process, and particularly how Netflix and IMDB titles were matched, please consult Appendix A.

[9] Two reasons motivate this decision. First, because movie rating frequencies follow a heavy-tailed distribution, a vast majority of films were rated by a relative small number of users. If all were taken into account in generating the film association matrix A, many cell values would have been estimated using a small or even empty set of intersecting users, resulting in a biased measure of film identity. Second, because we are interested in public perceptions, we purposefully include films with a relative high general visibility.

[10] Generally speaking, the older the film included in the dataset, the more likely it received more favorable ratings.

ence of film connoisseurs whose tastes are significantly different from the general moviegoing audience's.[11]

Identity

We infer identity by constructing the matrix of film associations A as described above. Figure 1 plots the 500 most popular films (i.e. those rated by the highest number of users) comprising this network.[12] As is clearly visible from the diagram, the network is divided into two cliques, with a third more loosely structured group of films in between the two main clusters. A formal partitioning of the network into clusters confirms this intuitive impression. The eigenvalue-based spectral partitioning algorithm for networks with positive and negative links (Traag & Bruggeman 2009) divides the network into three clusters comprising 47.15%, 15.61% and 37.24% of the nodes.[13] This division is color coded in Figure 1.

[—— Figure 1 about here ——]

The structure emerging from the network visualization in Figure 1 is one a of a uni-dimensional continuum. Most nodes are clustered around the two poles of this continuum, but some are in between the poles. A principal component analysis (PCA) of the full covariance matrix between films confirms this intuition. PCA is a dimensionality-reduction technique that reduces a set of observations onto components that account for as much variability as possible in a descending order. The scree plot of the principle component eigenvectors is presented in Figure 2. PCA eigenvalues correspond to the amount of variance that each component explains. As is visible from the scree plot, the first dimension produced by PCA explains a significantly greater amount of variance than the other components

[11] Number of ratings is strongly correlated with a film's adjusted gross income, at 0.50. Average rating is correlated with a film's yield (gross income divided by budget) at 0.29.

[12] To make the image informative, we restricted the network size to 500. Negative edge weights are not visualized. Nodes are spatially positioned using the Fruchtman-Reingold algorithm.

[13] To reduce noise, we bootstrap with resampling do determine which edge weights are significant at the $\alpha = 0.05$ level. Insignificant edge weights were reduced to zero before applying the partitioning algorithm. For more details about the bootstrapping method, see Goldberg (2011).

explain. This component corresponds to the underlying dimension of variability along which the various films are positioned in the network diagram in figure 1. With a Spearman rank correlation coefficient of 0.894, the first PCA component and the partitioning assignments are almost perfectly overlapping. Both correspond to the same axis of structure, but in different ways. The PCA scale provides a continuous measure of placement on this axis. The partitioning assignment, on the other hand, draws boundaries around three sections of this axis. If network A corresponds to an underlying classificatory scheme, as we argue, then these different groups delineate three latent categories that structure people's sense-making of films.[14]

[—— Figure 2 about here ——]

Boundedness

Our central independent variable of interest is boundedness: the extent to which a film's team members tend to collaborate with one another on other films. To measure boundedness, we construct a network, whereby each node represents a film, and edges connecting films to one another are weighted as a function of the number of individuals who were involved in both films' production. An edge weighing 4, for example, implies that four individuals were team members on both of the films it connects. Because we are interested in the effects of boundedness on audience perceptions, we construct this network relying exclusively on functionaries that are significantly visible to audiences: directors(s), producer(s) and actor(s).[15]

Let T_i denote the set of n individuals $\{t_1, t_2 ... t_n\}$ who are members of film i's production team. Network S is defined as $S_{ij} = |T_i \cap T_j|$. Network S represents the compositional similarity between films' production teams. For each year y contained in our dataset, we

[14] In Appendix B, we discuss at length why our measure of identity is not biased by *Netflix*'s website, and particularly its recommendation engine.

[15] These different production roles are often highlighted in marketing materials, and tend to be conspicuously positioned in the opening credit sequence. We also experimented with limiting this measure to include only the first ten actors, by credit order, as a means of focusing exclusively on cast members that are saliently visible to audiences. The results reported below are robust to these different ways of measuring boundedness.

produce two networks: (1) $\overleftarrow{S^y}$ includes all films produced up until and including year y, and (2) $\overrightarrow{S^y}$ which includes all films produced on or after year y. Whereas $\overleftarrow{S^y}$ represents the structure of inter-film collaborations at time of production, $\overrightarrow{S^y}$ represents the structure of collaboration post-production.

Socially bounded films are those films that are enclosed in tight-knit circles in network S, that is, whose team members tend to repeatedly collaborate with one another. A common measure for network closure is the clustering coefficient. The clustering coefficient measures the extent to which an ego node's network of alters are tied with one another. In a binary network, the clustering coefficient is simply the proportion of realized ties among ego's neighbors, which is calculated as the number of closed triads divided by the number of potential triads (often referred to as "network density"). In a weighted network, the clustering coefficient also takes into account the intensity of each triad, and is calculated as follows (Onnela, Saramäki, Kertész & Kaski 2005):

$$C_i = \frac{2}{k_i(k_i-1)} \sum_{j,k} (\widetilde{S}_{ij}\widetilde{S}_{jk}\widetilde{S}_{ki})^{\frac{1}{3}} \tag{2}$$

where edge weights are scaled by the largest weight in i's neighborhood, $\widetilde{S}_{ij} = S_{ij}/max(S_{ij})$, and k_i is node i's degree, or the number of nodes with which it is connected with an edge of $weight > 0$.[16] Substantively, the clustering coefficient measures the extent to which members of a film's production team also tend to collaborate with one another on other films. For each film i we construct two measures of boundedness: (1) $\overleftarrow{C_i}$, which is calculated over $\overleftarrow{S^y}$ (where y is the film's year of production) and represents its *production boundedness*, and (2) $\overrightarrow{C_i}$, which is calculated over $\overrightarrow{S^y}$ and represents the film's *post-production boundedness*. As we explain below, this distinction is central to our research design. Because boundedness follows a positively skewed distribution, we log transform it (and because it ranges from 0 to 1, we multiply it by 100 and add 1 before applying the log transformation). Moreover, though C_i is formulated such that it is adjusted by its network neighborhood size, generally

[16] $C_i = 0$ for nodes that have no neighbors.

it tends to decrease as the number of individuals making up a production team increases.[17] We therefore include $Ln(|T_i|)$ as a control measure in the models below, as a means to account for variation in boundedness that is net of team size.

Additional Film Attributes

The IMDB database provides a variety of information about the films included in the dataset. Several types of these data are used in the following analysis. We use IMDB's genre assignments as a means to measure institutional categorical identity. Overall, the films included in the dataset were assigned 22 different and non-mutually exclusive genre labels. These vary significantly in prevalence (from almost 50% of films labeled Drama, to 1.5% labeled Documentary). Films also vary in the number of labels they were assigned. Only a minority are assigned less than two labels. We use the number of genre labels assigned to a film as a measure of its genre niche width, i.e. the extent of its multi-categorical membership (following Hsu (2006) we log transform this variable to account for its skewed distribution). Two additional variables that relate to film content are included in our models: a film's runtime, and whether or not it is a sequel.

We collect additional information from IMDB about the film's production process. To account for the organizational and capital resources available at time of production, we differentiate between films produced by major studios (and their subsidiaries) and independent studios. Historically, major studios have enjoyed disproportional market share throughout the industry's existence. Their unrivaled capital resources, as well as their access to exclusive relationships and distribution channels, provides them with significant commercial advantages over independent studios. At the same time, because major studios account for the vast majority of blockbusters, they are often perceived as willing to make compromises on the artistic quality of their products.[18] The production team size is also included in

[17] This is largely because the number of triads increases squarely with a linear increase in the number of individuals comprising a film.

[18] We do not include film budget in our models because of the prevalence of missing data, at roughly 35%. Because budget is highly correlated with whether or not a film was produced by a major studio (the latter explaining almost 37% of the variance in the former, adjusted for inflation), we use studio size as a measure of capital resources. Moreover, because major studios have access to exclusive distribution channels, studio

our models. As we explain earlier, we include this variable in order to correctly adjust our boundedness measure, yet in its own right, team size also measures the complexity and labor intensity of the production process.

We use information on all films included in the IMDB database (irrespective of whether or not these films are included in our dataset) to construct additional control variables that represent the cumulative experience of the individuals involved in a film's production process. To measure team member experience, we use the number of other films that the average team member participated in during the years preceding year of production.[19] We also determine team members' reputation by enumerating how many have been awarded major awards for their work on films produced in earlier years. As previous research demonstrates (Rossman, Esparza & Bonacich 2010), these forms of consecration significantly affect how individuals are perceived. Roughly 41% of the films in our dataset include no members who have received a major award for prior work.[20]

Analytical Strategy

Our goal is to show that boundedness affects audiences' perceptions. To do so, we need to isolate audience- from production-side processes, and demonstrate the former's implications on moviegoers' interpretations. By production-side processes, we refer to procedures that shape the product itself, and by audience-side we refer to processes that influence perceptions net of product properties. Though analytically separable, in reality the two are inherently intertwined: boundedness relates both to the organization of production, and to structural identity.

size can have an affect on eventual audience reach beyond mere availability of capital resources.

[19]We include only team members who participate in a creative capacity: producers, directors, actors and actresses, screenwriters, cinematographers and editors. For actors, we include only the first ten (or less) by order of credit. For each role, we calculate the average member's experience, and standardize it on a 0 to 1 scale (where 1 is the maximum per role in the dataset). We then construct the variable by averaging over all roles.

[20]We use the three major and established awarding institutions in three different categories of awards: industry awards (Academy Awards, Golden Globe Awards and BAFTA Awards), festival awards (Cannes Film Festival, Venice Film Festival and Berlin Film Festival) and critics awards (NSFC, NYFCC and LAFCA). The oldest awards in the film industry, these awards are also generally considered the most prestigious. In constructing this measure, we consider only awards awarded to an individual (eg. Best Actress or Best Achievement in Film Direction).

This is where the distinction between production and post-production boundedness comes handy. Production boundedness captures the social position of a film at the time of its production. It likely affects both the production process in and of itself, as well as audience reception. Post-production boundedness, in contrast, relates to production team members' career trajectories after a film had been completed and released. If it has an effect on how a film is received, this effect can only operate through its implications on audience perceptions.[21] The films contained in our sample are evaluated by consumers only after they had been released as DVDs, and in most cases, years after the film's production.[22] This lag allows us to meaningfully isolate post-production boundedness and estimate its effect on audience perceptions. Because boundedness affects perception at time of production as well, post-production boundedness is overall a conservative estimate of structural identity in the film market.

Ours is not a perfect difference-in-differences design, as we do not have information on audience perceptions at time of original release. We therefore cannot measure the full extent to which perceptions explained by post-production boundedness differ from perceptions when the film was first viewed by audiences. This introduces endogeneity bias if post-production boundedness results from an unobserved film characteristic (for example, if an unobserved artistic quality that affects how audiences perceive the film also makes it more likely for team members to collaborate again in the future). Though we have no information about audience perceptions at time of release, we do know how films were received by contemporaneous critics and award committees. We therefore include in our models a variable that gauges whether, and to what extent, a film was nominated for awards by established award-granting institutions immediately after its production.[23] Because

[21] As illustration of how films can be historically reinterpreted, consider Douglas Sirk's 1950s melodramas. Though commercially successful, these films were poorly received by contemporaneous critics. By the 1980s, however, a new generation of critics and scholars rediscovered Sirk's work as a staple of subtle irony and cinematic prowess under the strict limitations imposed by the studio production system of the early postwar era. His ingenuity was celebrated in homages by acclaimed art-house directors such as Pedro Almodovár (Klinger 1994).

[22] The median film is evaluated ten years after it was produced.

[23] For this variable, we include only award nominations that relate to the film as a whole, such as Best Picture, Best Director or Best Screenplay, relying only on the most prominent industry and festival awards (because critic awards do not nominate contenders, we exclude them from this measure). We calculate the

established awards tend to be bestowed upon mainstream productions, we also control for whether a film was nominated for the Grand Jury Prize at the independent Sundance Film Festival. Nomination for an award, as opposed to actually winning it, is a liberal measure of a film's recognition. As we demonstrate in the following section, these two variables serve as a reasonable approximation of audience perceptions at time of production, allowing us to rule out endogeneity bias.

Results

Identity as Sophistication

What structures the identity space depicted in figure 1? A naked-eye inspection of network *A* strongly echoes with popular distinctions between commercial and art films. Whereas one group of films – corresponding to one end of the PCA scale – includes critically acclaimed movies such as Francis Ford Coppola's 1972 *The Godfather* (ranked by the American Film Institute as the second greatest American movie of all time) or the 1996 *Breaking the Waves* (winner of the 1996 Cannes Festival Grand Prix), another group, overlapping with the opposite PCA scale end, includes blockbusters such as the 1998 action movie *Armageddon* or the sex comedy *Striptease*, the latter winner of the 1996 *Razzie Award* for worst picture. A third category of 'in-between' movies includes a variety of films that over the years have reached cult status, among them the 1980 horror film *Friday the 13th* and the satirical comedy *Airplane!*, also from 1980. A comparison with critics' opinions gives flesh to this impression.

Figure 3 plots critics' evaluations as a function of the division into latent categories (box-plot on left) and the PCA scale (right). These diagrams clearly illustrate that audience tastes' are structured by a hierarchy that strongly corresponds to critics' evaluations of films. The latent categories map onto a division between what can be described as sophisticated and common movies.[24] These two groups of films are wedged by crossover movies

proportion of institutions nominating a film for an award out of the overall number of institutions giving awards in the year following a film's release.

[24]By using the term 'sophistication' we do not mean to imply that these films are objectively more complex or refined. Rather, we use the term to denote the received wisdom about these films' qualities. Our usage

that, like sophisticated films, tend to be highly evaluated by critics. The PCA scale, as illustrated on the right-hand panel, similarly delineates a hierarchy of sophistication. It explains more than 59% of the variance in critics' judgements. It appears that, in the eyes of the audiences perceiving them, films are identified primarily as a function of their perceived artistic quality.[25]

[—— Figure 3 about here ——]

The sophistication scale crosscuts conventional genre labels. The category Comedy, for example, includes such celebrated films as Stanley Kubrick's nuclear-scare satire *Dr. Strangelove*, and, on the other hand, the box-office flop *For Richer or Poorer*, described by one critic as a "bottom of the barrel comedy tripe." The vast majority of genres include films that span a wide range of the scale. This is not say that adherence to genre conventions does not affect audience perceptions. As recent scholarship demonstrates, it does (Hsu 2006, Hsu, Hannan & Koçak 2009). Rather, it suggests that other social mechanisms are at play.

The Structural Antecedents of Sophistication

What might these mechanisms be? Of course, a natural explanation is that something inherent about the products themselves, the ineffable yet presumably recognizable quality that makes certain films artistic, stands at the core of this consensus. Yet, if our hypothesis about the structural origins of identity is correct, we should find that boundedness, at least in some part, explains films' location on the sophistication scale.

Table 1 reports several OLS models, where the dependent variable is the sophistication scale (reported as a standardized score). Our independent variable of interest in Model 1, production boundedness, is positively associated with perceived sophistication. Because the coefficient is difficult to interpret, we plot the effect in Figure 4. As it illustrates,

is therefore consonant with the common distinction between high and low brow art (Bourdieu 1979).

[25]It is important to distinguish between two types of audiences: critics and regular consumers. It may very well be the case, and in fact highly consistent with theories about identity in markets, that critics' appraisals of films have a significant impact on consumers' evaluations. Yet this distinction is outside the scope of this study. Rather, we use critics' reviews as a means to measure a film's institutional status. Like consumers, critics also make inferences about quality from categorical and structural identity signals.

boundedness overall translates to almost a full standard deviation increase in sophistication.

[—— Table 1 about here ——]

[—— Figure 4 about here ——]

Model 1 includes a variety of control variables that account for various film characteristics. We include a complete set of 22 genre dummies, as a means to capture film content. Consistent with previous literature, films' categorical niche widths, or the extent of their categorical incoherence, is negatively associated with perceived sophistication. Additional attributes are also related to audience perceptions. As the number of previously awarded individuals in the production team increases, so does the film's perceived sophistication, suggesting that attributions of quality originate, to some extent, from knowledge about acknowledged talent.[26] In contrast, sophistication generally decreases with production team size and studio size, implying that the large cinematic productions pursued by major studios do not tend to be associated by audiences with artistic quality. Sophisticated audiences appear to be similarly averse to the repetitiveness of sequels, and, all other things being equal, more appreciative of newcomers. They also interpret an unusually long cinematic production as a marker of quality. Finally, because the older films included in the dataset tend to be those that have already been established as 'classics,' we include fixed year effects in the model.

Since we cannot separate production- from audience-side mechanisms by examining production boundedness alone, in Model 2 we include post-production boundedness as an additional variable. Both coefficients are positive and significant. Though we cannot rule out that production boundedness affects sophistication through the material advantages it affords to the production team, we can safely say that post-production boundedness can only affect sophistication through audience perceptions, as it occurs after the film was

[26]It might also be the case that the objective qualities valued by award committees are similarly valued by audiences who have a preference for sophisticated films.

produced.

Yet, as we discuss earlier, distinguishing between production and post-production boundedness does not fully mitigate problems of endogeneity. Because the feature film market is rife with uncertainty, professionals often prefer teaming with those they have collaborated successfully with in the past (Faulkner & Anderson 1987). Post-production boundedness might therefore be the result, rather than the cause, of perceived sophistication. To control for this possibility, we include two variables that measure how films were received by expert audiences immediately after their production: nomination to established and independent film awards.

Before moving on to report the remaining two models, it is important to clarify why contemporaneous recognition allows us to address endogeneity. First, as we saw earlier, sophistication is highly consistent with critics' evaluations. Whether this is because award committee members and the general audience rely on similar appraisal criteria, or whether audiences adjust their evaluations in response to accolades received by these venerable institutions, is beyond the scope of our study. Nevertheless, it is reasonable to assume that a film's recognition by award committees also reflects its interpretation by the general public at the time. In fact, these two variables, when modeled separately, explain more than 29% of the variance in sophistication.

Second, whether or not a film was nominated for an award would have provided its production team members with strong contemporaneous information as to whether it was successful. If the reason for concern about endogeneity relates to production team members' decision to regroup in response to their film's perceived sophistication at time of release, then award nomination allows us to control for cases in which team members decided to collaborate again because of their impression that their film was greeted by audiences with recognition. Any changes in perceived sophistication due to team members' subsequent career trajectories, specifically as they affect post-production boundedness, are not endogenous to film characteristics at time of production.

Our central empirical findings are presented in Models 3 and 4 where we control for contemporaneous institutional recognition. All the variables included in the model predate the audience evaluations that were used to generate the sophistication measure; reverse causality, therefore, can be ruled out. And as we argue above, the distinction between production and post-production boundedness, as well as the inclusion of institutional recognition variables, enable us to identify the effect of boundedness on perceived sophistication that is exclusively attributable to changes in the network positions of the people making up a film's production team after they completed their work on the film, and it was released to the general audience. As we hypothesized, boundedness marks quality, and is translated, in the eyes of audiences, into sophistication. At 0.09, the post-production coefficient represents an almost half standard deviation increase in sophistication for films whose post-production boundedness is 1. This is, most likely, a very conservative estimate of the overall effect of boundedness on audience perceptions. It excludes any effects that boundedness at time of production might have on audience and experts' evaluations of films.

The final model in Table 1 interacts year of production with our two measures of boundedness. For post-production boundedness to have a substantial effect on audience perceptions, a sufficiently long period – during which team members could potentially embark on new projects that would mature into film releases that audiences had the time to become familiar with – must have elapsed since the film was produced. We therefore create a dummy variable that represents whether a film was produced in 2001 or later, a five year window preceding the most recent rating included in our dataset. We expect to find that post-production boundedness affects audience perceptions only for non-recent films. Indeed, as expected, production boundedness increases in effect size for recent films, and post-production decreases (Model 4). The marginal effects of post-production boundedness on sophistication are visualized in Figure 5. Whereas a change in boundedness for recent films has no significant effect on their sophistication, the effect is very substantial for older films. The full range of post-production boundedness translates to more than one standard deviation change in perceived sophistication for films that at least five years have passed since their release. Films' whose members continued to work on other films with one another tend to be perceived as substantially more sophisticated than those whose members did not, net of the films' production circumstances, content, and whether or not they were lauded by the film establishment upon their release.

[—— Figure 5 about here ——]

The Commercial Implications of Sophistication

We began the analysis by arguing that aggregate measures of market success are insufficiently refined for getting a handle on how audiences infer identity. This does not mean, however, that identity is inconsequential for market success. In this final part of the analysis we demonstrate that perceived sophistication is related to how films fare commercially.

As Figure 6 demonstrates (left panel), films on either end of the sophistication scale engender far more disagreement than those in the middle. Disagreement is measured as the standard deviation of a film's ratings; the higher the standard deviation, the higher the overall disagreement between viewers about the film's quality. The estimates presented in Figure 6 were obtained using a simple bivariate model, where disagreement is estimated as a square function of sophistication. Movies on either end of the sophistication scale tend to exhibit almost two standard deviations higher disagreement than those in the middle of the scale. These are the films that have strong identities, and that elicit strong emotions – whether positive or negative – from moviegoers. The 2003 drama *Dogville*, with the highest rating standard deviation in the dataset, is exemplary of such disagreement. Using a minimalist stage-like set, this parablistic movie directed by critically acclaimed Danish director Lars von Trier, has the unquestionable hallmarks of an art film. But even critics could not agree on its quality. With 70 points on Rotten Tomatoes, *Dogville* was hailed as "singular and profound" by the *San Francisco Chronicle*, but criticized by the *Chicago Sun-Times*' Roger Ebert for exhibiting the "imagination of an artist and the pedantry of a crank." On the other hand, movies around the midrange of the sophistication scale tend to invoke far greater consensus.

[—— Figure 6 about here ——]

And as Figure 6 (right panel) illustrates, this consensus tends to be positive. The diagram plots films' return on investment (ROI) as a function of their sophistication, as estimated by model 1 in Table 2. ROI is a film's percent of net profit, and is calculated by dividing its gross income by its budget (log transformed). We include a variety of control variables that might affect a film's commercial success beyond its perceived quality, such as whether it was produced by a major studio, its runtime and whether or not it is a sequel. Sophistication exhibits an inverted U-shaped relationship with commercial profit (Model 1); its magnitude is highly substantial: films on the upper end of the sophistication scale are expected to yield roughly 80% less profit than those at the center of the scale. Slightly more variance is explained by Model 2, where sophistication is modeled as a categorical variable (corresponding to the latent categories inferred from the partitioning procedure), and crossover films are used as the omitted category. Crossover films are estimated to yield 2.13 times more profit than sophisticated films and 1.76 times more than common movies.

The differences between coefficient estimates in Tables 1 and 2 also shed light on how identity differs from commercial success. For example, whereas the coefficient for sequel is positive and significant when the dependent variable is ROI, it is negative when it is sophistication. This is likely because in a market that celebrates creativity and novelty, a sequel signals a lack in both qualities, but at the same time, because sequels usually involve individuals who have worked together in the past, they significantly reduce the coordination costs associated with a film production. Similarly, while critically acclaimed actors add to a film's perceived sophistication, because they are more expensive to acquire they reduce its overall return on investment. These different results suggest that different mechanisms affect how a film is perceived, and how it fares commercially.

[—— Table 2 about here ——]

Of course, it would be far fetched to argue that these results describe a causal relationship. Because sophistication is inferred from audiences' ratings, which are recorded after a film had been released and completed its theatrical distribution (in many cases, many years after its release), it is impossible to determine whether films are perceived as crossover as a result of their theatrical success, or whether their success is ipso facto the result of their inherent crossover appeal. We are therefore cautious not to draw any unwarranted conclusions. What these results do demonstrate, however, is that ultimately a film's commercial profit is strongly related to its ability to cross over audiences on both ends of the sophistication scale. Identity and market success are intertwined, but in a manner that undermines the intuitions informing recent scholarship on the detrimental effects of category spanning. Films that bridge the divide between sophisticated and low-status offerings are most likely to yield high return on investment.

Conclusions

Theories of market behavior often treat aggregate measures of success – such as a film's gross income – as information about consumers' revealed preferences. In contrast, this study examines consumers' actual stated preferences in great detail as a means to make more fine grained distinctions between the ways in which products – in this case, films – are perceived by audiences. Rather than making assumptions about how audiences interpret products, we infer these interpretations from individual rating behaviors. Relying on the lag between time of production and evaluation in our dataset, we were able to demonstrate that viewers' assessments of a film are influenced by team members' career paths years after the production team's disbandment.

Our findings add flesh to the oft recited conjecture that 'quality' is socially constructed (Lynn, Podolny & Tao 2009), in three different ways. First, we demonstrate that consumer perceptions are structured by a rigid hierarchy, spanning from common to sophisticated movies and embodying an implicit boundary between commercial films and 'art.' Whereas previous studies have focused on the manner by which institutionalized categories provide shared schematic templates that structure consumer expectations, this study sheds light on the structural foundations of market identities. We demonstrate that, given the absence of concrete information, audiences turn to signals emanating from production team members' location within a network of relationships as a a means to make inferences about a film's quality. Identity, in other words – at least in the film industry – is as much structural as it is categorical. Contra recent evidence on the deleterious effects of categorical ambiguity (Hannan 2010), our finding that sophistication has an inverted U-shaped relationship with commercial success suggests that a coherent market identity is not necessarily financially beneficial.

Second, we find that this hierarchy of sophistication traces to the social boundedness of cinematic production teams. Sociological literature on market competition tends to conceptualize status as cumulative advantage that inheres in ties to prominent others (Merton 1968, Podolny 1993, Stuart, Hoang & Hybels 1999). Our analysis shifts focus to the structural production of exclusivity through boundary work. In the eyes of observers, these emergent boundaries are interpreted as signals about quality. We do not argue that audience members are lay network analysts who are computing clustering coefficients on the fly. Rather, boundedness captures the extent to which a group of people tend to collaborate with one another repeatedly. When noticed by others, this groupness becomes reified as identity (White 1992).

Finally, our analytical strategy enables us to isolate the effects of post-production from production boundedness, thereby singling out how network structures function as prisms through which market actors ascertain each other's qualities (Podolny 2001). We demonstrate that the boundedness of films is related to their perceived sophistication because of the identities of the people involved in their production, and not exclusively due to the organizational circumstances of production and how those affect product properties. This finding is highly consequential to our understanding of how perceived quality is related to the interpersonal organization of creative production. The majority of the literature on networks and creative output conceptualizes social structure as the material infrastructure upon which information, knowledge and ideas travel. It finds that a fine balance between structural porousness and cohesion is necessary for facilitating the kind of conceptual conductivity that catalyzes new ideas (Burt 2004, Uzzi & Spiro 2005, Vedres & Stark 2010). Yet, networks also confer identity to those they are made up of. To the same extent that network closure generates a sense of belongingness among its members (Podolny & Baron 1997), we find that it bestows them with quality in the eyes of observers. Brokerage and closure affect not only what an actor does, but also how it is perceived (Kilduff & Krackhardt 1994).

Sociological research has focused extensively in recent years on how identity mitigates problems of uncertainty in markets. A central piece of this puzzle relates to where identity comes from. Though we have focused only on one industry, its project-based inter-organizational mobility, and its foundation on creative novelty, are also characteristic of a variety of other creative and entrepreneurial market domains. It remains to be seen whether structural identity, and the exclusionary identity-building dynamics it is structured on, is as salient in those domains as it is in Hollywood.

References

Ahuja, Gautam. 2000. "Collaboration Networks, Structural Holes, and Innovation: A Longitudinal Study." *Administrative Science Quarterly* 45(3):425–455.

Aral, Sinan & Marshall Van Alstyne. 2011. "The Diversity-Bandwidth Trade-off." *American Journal of Sociology* 117(1):90–171.

Baumann, Shyon. 2007. *Hollywood Highbrow: From Entertainment to Art.* Princeton, NJ: Princeton University Press.

Becker, Howard Saul. 1982. *Art Worlds.* Berkeley and Los Angeles, CA: University of California Press.

Bielby, William T. & Denise D. Bielby. 1994. "All Hits Are Flukes: Institutionalized Decision Making and the Rhetoric of Network Prime-Time Program Development." *American Journal of Sociology* 99(5):1287–1313.

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

Bourdieu, Pierre. 1979. *Distinction: A Social Critique of the Judgment of Taste*. Cambridge, MA: Harvard University Press. Translated by Richard Nice, 1984.

Brubaker, Rogers & Frederick Cooper. 2000. "Beyond "identity"." *Theory and Society* 29(1):1–47.

Burt, Ronald S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110(2):349–399.

Burt, Ronald S. 2005. *Brokerage and Closure: An Introduction to Social Capital*. Oxford: Oxford University Press.

Crossley, Nick. 2009. "The Man Whose Web Expanded: Network Dynamics in Manchester's Post/Punk Music Scene 1976–1980." *Poetics* 307(1):24–49.

DiMaggio, Paul J. 1987. "Classification in Art." *American Sociological Review* 52:440–455.

DiMaggio, Paul J. 2011. Cultural Networks. In *The SAGE Handbook of Social Network Analysis*, ed. John Scott & Peter J. Carrington. Thousand Oaks, CA: SAGE Publications pp. 286–300.

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

Faulkner, Robert R. & Andy B. Anderson. 1987. "Short-Term Projects and Emergent Careers: Evidence from Hollywood." *American Journal of Sociology* 92(4):879–909.

Gieryn, Thomas F. 1983. "Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists." *American Sociological Review* 48(6):781–795.

Gioia, Dennis A., Majken Schultz & Kevin G. Corley. 2000. "Organizational Identity, Image, and Adaptive Instability." *Academy of Management Review* 25(1):63–81.

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

Gould, Roger V. 2002. "The Origins of Status Hierarchies: A Formal Theory and Empirical Test." *American Journal of Sociology* 107(5):1143–1178.

Hannan, Michael T. 2010. "Partiality of Memberships in Categories and Audiences." *Annual Review of Sociology* 36:159–181.

Hannan, Michael T. & John Freeman. 1984. "Structural Inertia and Organizational Change." *American Sociological Review* 49:149–164.

Hannan, Michael T., Lásló Pólos & Glenn R. Carroll. 2007. *Logics of Organization Theory: Audiences, Codes, and Ecologies*. Princeton, NJ: Princeton University Press.

Hansen, Morten T. 1999. "The Search-Transfer Problem: The Role of Weak Ties in Sharing Knowledge across Organization Subunits." *Administrative Science Quarterly* 44(1):82–111.

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

Hsu, Greta & Michael T. Hannan. 2005. "Identities, Genres, and Organizational Forms." *Organization Science* 16(5):474–490.

Hsu, Greta, Michael T. Hannan & Özgecan Koçak. 2009. "Multiple Category Memberships in Markets: An Integrative Theory and Two Empirical Tests." *American Sociological Review* 74:150–169.

Kilduff, Martin & David Krackhardt. 1994. "Bringing the Individual Back in: A Structural Analysis of the Internal Market for Reputation in Organizations." *Academy of Management Journal* 37(1):87–108.

King, Geoff. 2009. *Indiewood, USA: Where Hollywood Meets Independent Cinema.* New york, NY: I.B. Tauris & Co.

Klinger, Barbara. 1994. *Melodrama and Meaning: History, Culture, and the Films of Douglas Sirk.* Bloomington, IN: Indiana University Press.

Krackhardt, David. 1999. "The Ties that Torture: Simmelian Tie Analysis in Organizations." *Research in the Sociology of Organizations* 16:183–210.

Lamont, Michèle & Virág Molnár. 2002. "The Study of Boundaries in the Social Sciences." *Annual Review of Sociology* 28:167–195.

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

Lévi-Strauss, Claude. 1963. *Structural Anthropology.* New York: Basic Books.

Lynn, Freda B., Joel M. Podolny & Lin Tao. 2009. "A Sociological (De)Construction of the Relationship between Status and Quality." *American Journal of Sociology* 115(3).

Merton, Robert. 1968. On sociological theories of the middle range. In *Social theory and social structure.* 1968 enlarged edition ed. New York: The Free Press pp. 39–53.

Mohr, John W. 1998. "Measuring Meaning Structures." *Annual Review of Sociology* 24:345–370.

Murphy, Gregory. 2004. *The Big Book Of Concepts.* Cambridge, MA: MIT Press.

Onnela, Jukka-Pekka, Jari Saramäki, János Kertész & Kimmo Kaski. 2005. "Intensity and coherence of motifs in weighted complex networks." *Physical Review E* 71:065103.
URL: *http://link.aps.org/doi/10.1103/PhysRevE.71.065103*

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

Podolny, Joel M. 1993. "A Status-Based Model of Market Competition." *American Journal of Sociology* 98:829–872.

Podolny, Joel M. 1994. "Market Uncertainty and the Social Character of Economic Exchange." *Administrative Science Quarterly* 39(3):458–483.

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

Podolny, Joel M. & James N. Baron. 1997. "Resources and Relationships: Social Networks and Mobility in the Workplace." *American Sociological Review* 62(5):673–693.

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

Rao, Hayagreeva, Phillipe Monin & Rodolphe Durand. 2005. "Border Crossing: Bricolage and the Erosion of Categorical Boundaries in French Gastronomy." *American Sociological Review* 70(9):868–891.

Roberts, Peter W., Mukti Khaire & Christopher I. Rider. 2011. "Isolating the Symbolic Implications of Employee Mobility: Price Increases after Hiring Winemakers from Prominent Wineries." *American Economic Review* 101(3):147–151.

Rosch, Eleanor H. 1978. Principles of Categorization. In *Cognition and Categorization*, ed. Eleanor H. Rosch & Barbara B. Lloyd. Hillsdale: Lawrence Erlbaum pp. 27–48.

Rossman, Gabriel, Nicole Esparza & Phillip Bonacich. 2010. "Id Like to Thank the Academy, Team Spillovers, and Network Centrality." *American Sociological Review* 75(1):31–51.

Ruef, Martin & Kelly Patterson. 2009. "Credit and Classification: The Impact of Industry Boundaries in 19th Century America." *Administrative Science Quarterly* 54:486–520.

Sauder, Michael, Freda Lynn & Joel M. Podolny. 2012. "Status: Insights from Organizational Sociology." *Annual Review of Sociology* 38:267–283.

Stuart, Toby E., Ha Hoang & Ralph C. Hybels. 1999. "Interorganizational Endorsements and the Performance of Entrepreneurial Ventures." *Administrative Science Quarterly* 44(2):315–349.

Tortoriello, Marco & David Krackhardt. 2010. "Activating Cross-Boundary Knowledge: The Role of Simmelian Ties in the Generation of Innovations." *Academy of Management Journal* 53(1):167–181.

Traag, V. A. & Jeroen Bruggeman. 2009. "Community Detection in Networks with Positive and Negative Links." *Physical Review E* 80:036115.
URL: *http://link.aps.org/doi/10.1103/PhysRevE.80.036115*

Turner, Graeme. 1999. *Film as Social Practice.* 3 ed. London: Routledge.

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

Vedres, Balázs & David Stark. 2010. "Structural Folds: Generative Disruption in Overlapping Groups." *American Journal of Sociology* 115(4):1150–1190.

Whetten, David A. 2006. "Albert and Whetten Revisited: Strengthening the Concept of Organizational Identity." *Journal of Management Inquiry* 15(3):219–234.

White, Harrison C. 1992. *Identity and Control: A Structural Theory of Social Action.* Princeton, NJ: Princeton University Press.

Yogev, Tamar. 2010. "The Social Construction of Quality: Status Dynamics in the Market for Contemporary Art." *Socio-Economic Review* 8(3):511–536.

Zaheer, Akbar & Giuseppe Soda. 2009. "Network Evolution: The Origins of Structural Holes." *Administrative Science Quarterly* 54(1):1–31.

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

A Appendix: Data, Matching and Structuring

The Netflix Prize was an open competition launched by Netflix Inc., an American online DVD rental service and on-demand media streaming provider, on October 2$^{\text{nd}}$, 2006. The objective of the competition was to improve the company's recommendation algorithm, based exclusively on users' previous rating activity, by at least 10 percent. The $1M prize was eventually awarded almost three years later, on September 21$^{\text{st}}$, 2009.

The data used in this study were downloaded from the Netflix Prize website, at `www.netflixprize.com`, in January 2009. They comprise 100,480,507 ratings, provided by 480,189 unique users to 17,770 unique titles. Users are identified solely by a unique numerical identifier. Titles are similarly identified by a numerical identifier in the dataset. Each datapoint contained in the dataset is a quadruplet containing the user id, title id, date of rating and rating. Ratings range from 1 to 5, in full integer increments. A separate dataset provided by Netflix contains additional title identifiers: a textual string corresponding to the title, and a number corresponding to year of production. No additional identifying information is provided in the original data.

In order to match titles with their corresponding entries in Netflix's database, the official Netflix API was used (for more details, see `developer.netflix.com`). This allowed retrieving additional information about each title, including cast members and director(s). Overall, 81 titles, accounting for 0.03% of all titles, could not be matched thorough Netflix's API, and were removed from the dataset. Duplicate titles, where the same original title appears in multiple versions in the raw data, were also removed. Such duplicates include special editions (e.g. 20$^{\text{th}}$ anniversary edition, collector's edition, bonus material), versions (e.g. director's cut) and formats (e.g. widescreen).

The resulting dataset was matched with the Internet Movie Database (IMDB), available online at `www.imdb.com/interfaces`. IMDB is a comprehensive database including information on films, television shows and video games. Because film titles can appear differently on Netflix and IMDB (especially, though not exclusively, foreign language films), matching between the two databases is not a straightforward task. Moreover, many film titles are not unique (e.g. Franco Zafirelli's 1968 Romeo & Juliet and Baz Luhrmann's 1996 film of the same title), and the databases vary, and at times are internally inconsistent, in their use of special characters (e.g. apostrophe), numerical characters and concatenation of subtitles. To overcome these problems, we matched titles on text, director name, and year of production. IMDB provides alternative titles, where available, and text was matched with all alternative titles. We used regular expressions (implemented with Python programming language), as well as text matching algorithms (e.g. Levenshtein distance), to standardize special characters, remove redundant text and ultimately match between non-identical strings. Where matching was not reliably attainable using these algorithms, titles were matched with human supervision (overall, 1,094 titles were matched with some human supervision).

The IMDB 'kind' classification distinguishes between different types of titles: feature films (classified as 'movie'), video movies, video games, tv movies, tv series and mini-series, and individual tv series episodes. Because our focus is the film industry, we retained only titles classified as movies by IMDB. The remaining dataset comprises 9,845 films, ranked by 479,578 unique users. As we discuss above, we retained only the 2,876 titles that correspond to 95% of all ratings provided by users who rated films.

We matched titles in the dataset with entries in the movie review aggregation website *Rotten Tomatoes*, using the website's API (see `developer.rottentomatoes.com` for details). The same text matching algorithms used to match between Netflix and IMDB titles were used to match between IMDB and Rotten Tomatoes titles. Eleven films could not be matched on Rotten Tomatoes, either because they do not appear in the database, or because they do not have an aggregation score.

B Appendix: Robustness to the Recommendation Engine

Because our analysis assumes that Netflix ratings reflect the true dimensions along which users perceive movie quality, it was important to exclude any possible confounding effects of Netflix's own *Cinematch* recommendation engine. By 2008, recommendations drove between 50 and 60 percent of the movies that users watch (see "If You Liked This, You're Sure to Love That" by Clive Thompson, p. MM74, New York Times, November 23, 2008). One of our concerns was that people rated movies that Netflix suggested differently than movies that they found through other sources, so that our measure of sophistication proxied quality of recommendation instead of quality of film. Using users' rating behavior to identify movies that people watched outside the influence of Netflix, we are able to exclude this possibility and show that film-viewers judged movies on the same dimensions of quality no matter how they stumbled onto the films.

In general, the recommendation system should not trouble our analysis. Sophistication as measured here is the dimension predicting the most covariance in user ratings. Netflix recommendations may well shift the movies that people watch, but in this capacity, they only drive the selection of movies, not the ultimate rating that users give. If the audience's dimensions of quality are fixed, their ratings do not depend on how they came to the movie, even if the movie was selected to appeal to the the audience. Two movies that fall equally high on underlying sophistication will draw similar ratings from similar users, and principle component analysis will identify this similarity. The worst case scenario is that Netflix's recommendation engine *has* identified sophistication and uses it to suggest movies: if users only ever rate suggested movies a 4 or a 5 instead of the full five-point scale, recommendations would raise the uncertainty of the observed rating covariance matrix as compared to the true distribution of ratings. This would be equivalent to measurement error in movie sophistication and would only serve to bias our coefficients toward zero, underestimating the true effect.

A more serious problem occurs if the audience's dimensions of quality are unstable and recommendations can shift the criteria that drive movie ratings. In this case, there would be two rating covariance matrices - the matrix of 'true' rating correlations, C_{True}, that reflects how the audience would judge movies in the absence of external recommendations; and the matrix of Netflix-biased rating correlations, $C_{Netflix}$, that reflects the alternative rating criteria users use on the website. If we knew both matrices, we could test for difference between them: If the two matrices differed, we would have evidence that Netflix ratings fail to capture true dimensions of quality. If the two matrices were identical, we could be confident that the ratings we observe from Netflix users accurately reflect the true parameters of sophistication.

Of course, these matrices are unknown. It is possible, though, to estimate them because of how Netflix users behave. Before Netflix can issue reliable recommendations for a given user, it has to tease out their existing tastes. At the same time, users can only rent movies a few at a time, and we should expect users to rate rented movies soon after watching them. As such, ratings will come in two flavors: on some days users will rate the few movies they have most recently rented. On other days, users will rate tens of movies at once, filling out ratings for those films they had watched before. By partitioning all ratings into those that occured on peak rating days, when a user rated many movies at once, and off-peak rating days, when a user rated only a few, we can estimate the true and Netflix-influenced rating correlation matrices C_{True} and $C_{Netflix}$ by C_{Peak} and $C_{Offpeak}$, respectively.

We defined *off-peak* rating days as days when a user rated ten or fewer movies, and *peak* rating days as those when a user rated between 20 and 150 movies. These capture 34.1 and 41.7 percent of all ratings respectively. We excluded days with over 150 ratings to ensure that a minority of unusual users didn't dominate the data - a small number of users rated 1000 or more movies in a single day, with one user rating 5,446. We compared correlation matrices for the most popular 200 movies only to ensure that there were enough ratings in every cell of the matrix. Figure B.1 shows the a strong positive relationship ($\rho = 0.9655$) between correlations on peak and offpeak days - if two movies had a high correlation among peak, or 'true,' ratings, they're likely to have a high correlation among offpeak, or recommender-influenced, ratings.

[—— Figure B.1 about here ——]

To confirm this positive relationship, we ran a Mantel test for the difference between the two matrices, which showed a significant relationship with probability $p = 0$ under the null hypothesis that the two matrices arose from distinct processes. The nonparametric Mantel test permutes the rows and columns of the two matrices to check whether the observed correlation could have arisen by chance as a function of the distance structure encoded in each matrix. As such, we can be confident that we find a significant similarity between the two correlation matrices. This suggests that we can be comfortable in interpreting Netflix ratings as an accurate indicator of users' external standards of movie quality, and need not worry about bias from the recommendation engine.

Table 1: OLS of Perceived Sophistication

	(1)	(2)	(3)	(4)
Boundedness:				
- Production	0.184***	0.166***	0.135***	0.176***
- Post-production		0.099***	0.092***	0.230***
- Recent*production				0.107*
- Recent*post-production				-0.187***
- Recent				0.111
Recognition:				
- % Established nominations			1.597***	1.812***
- Independent nominations			0.724***	0.646***
Content:				
- Ln(Niche Width)	-0.268*	-0.273*	-0.215	-0.223
- Runtime (hours)	0.193***	0.199***	0.038	0.058
- Sequel	-0.264***	-0.263***	-0.243***	-0.258***
Human Capital:				
- Mean individual experience	-2.142***	-2.142***	-1.840***	-1.767***
- Awarded individuals	0.099***	0.103***	0.083***	0.077***
Resources:				
- Ln(Team Size)	-0.193***	-0.181***	-0.176***	-0.191***
- Major studio	-0.346***	-0.338***	-0.321***	-0.300***
Intercept	0.422**	0.275	0.432**	0.315*
N	2876	2876	2872	2872
adj. R^2	0.422	0.425	0.480	0.459
Genre dummies	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	No

Sophistication scale is standardized to have a mean of 0 and standard deviation of 1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: OLS of Return on Investment

	(1)	(2)
Sophistication		
- Sophistication scale	0.005	
- Sophistication scale^2	-0.132***	
- Common		-0.568***
- Sophisticated		-0.757***
Content		
- Ln(Niche Width)	0.196	0.109
- Runtime (hours)	-0.293*	-0.287*
- Sequel	0.229*	0.237*
Human Capital		
- Awarded indiviuals	-0.050*	-0.044*
Resources		
- Ln(Team Size)	0.163*	0.100
- Major studio	0.174**	0.138*
Intercept	0.643	1.291***
N	1740	1740
R^2	0.233	0.252
Genre dummies	Yes	Yes
Year fixed effects	Yes	Yes

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

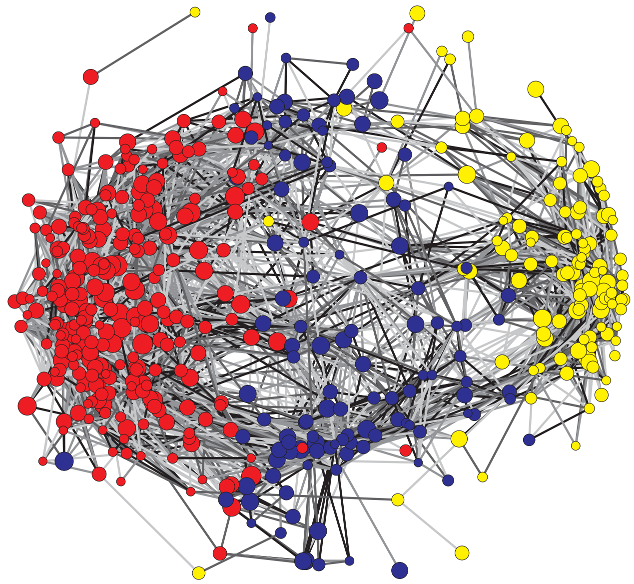

Figure 1: Visualization of the network of the 500 most rated films. Each node corresponds to a film, and edges are weighted by the correlation coefficient between the two films they connect. Edge shading corresponds to this weight. Node sizes correspond to the total number of users ranking a film (log transformed). Node color coding corresponds to the partitioning of the network into three clusters.

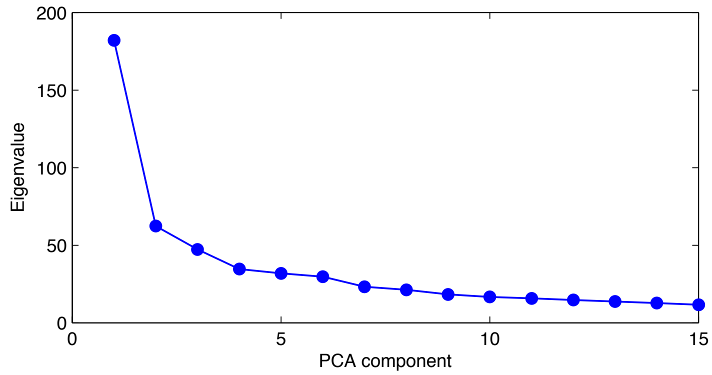

Figure 2: Variance explained by PCA components.

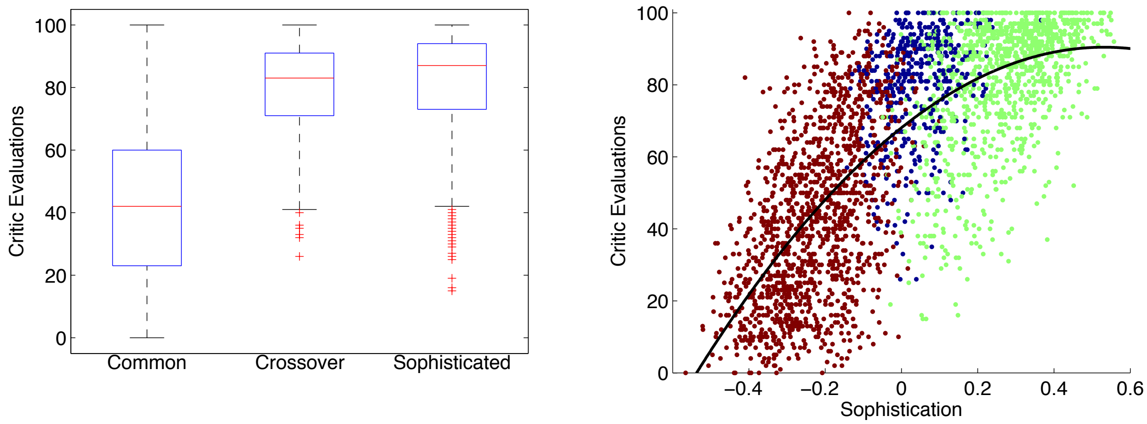

Figure 3: Critic evaluations as a function of the decomposition of network A. On the left, distributions are plotted by the network's partitioning into three groups. On the right, evaluations are plotted as function of the first PCA component. Films are color coded by their assignment to different categories. The black line corresponds to a fitted cubic model.

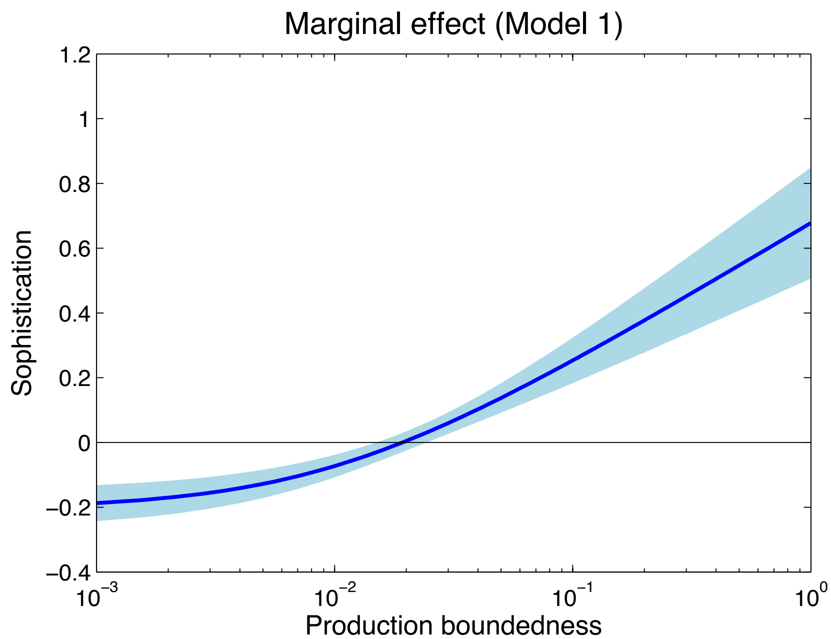

Figure 4: Marginal effects of production boundedness on perceived sophistication. Sophistication is standardized to have a mean of 0 and standard deviation of 1. Shades outline 95% confidence interval.

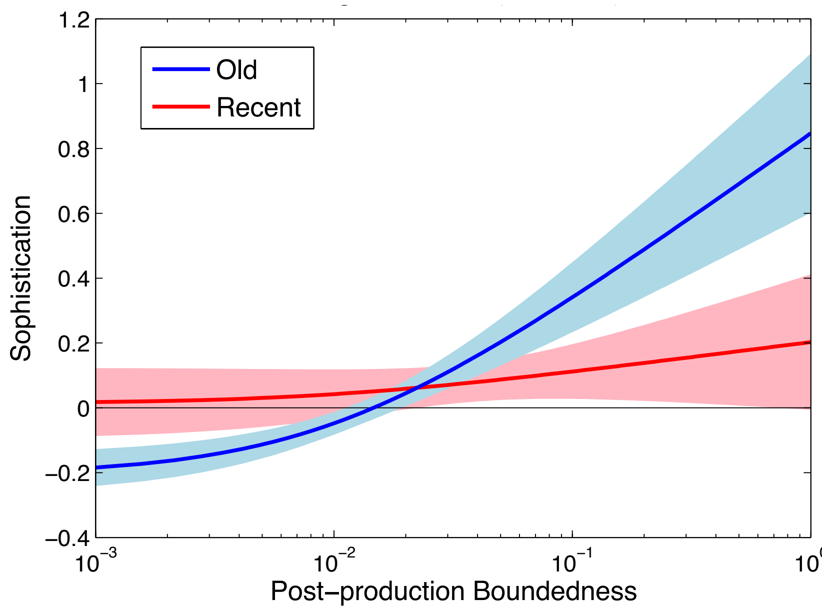

Figure 5: Marginal effects of post-production boundedness, interacted with time of production, on perceived sophistication. Sophistication is standardized to have a mean of 0 and standard deviation of 1. Shades outline 95% confidence intervals.

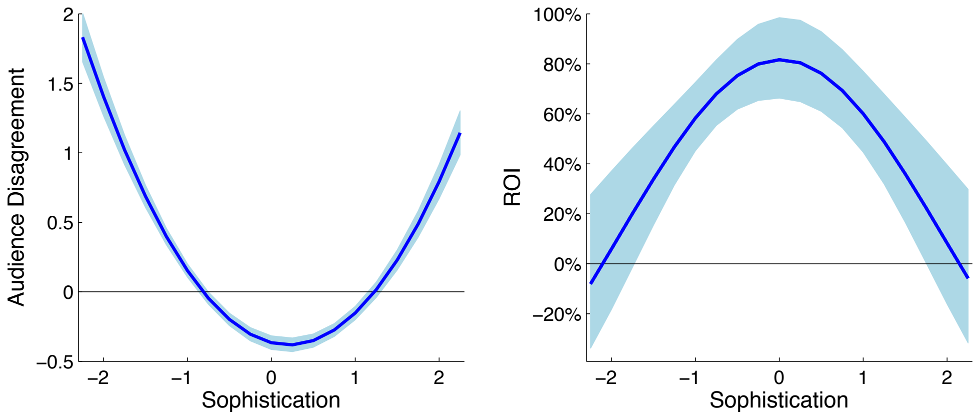

Figure 6: On the left, estimated audience disagreement (rating standard deviation) as a function of sophistication. On the right, estimated return on investment (logged) as a function of sophistication. Sophistication and disagreement are standardized to have a mean of 0 and standard deviation of 1. Shades outline 95% confidence intervals.

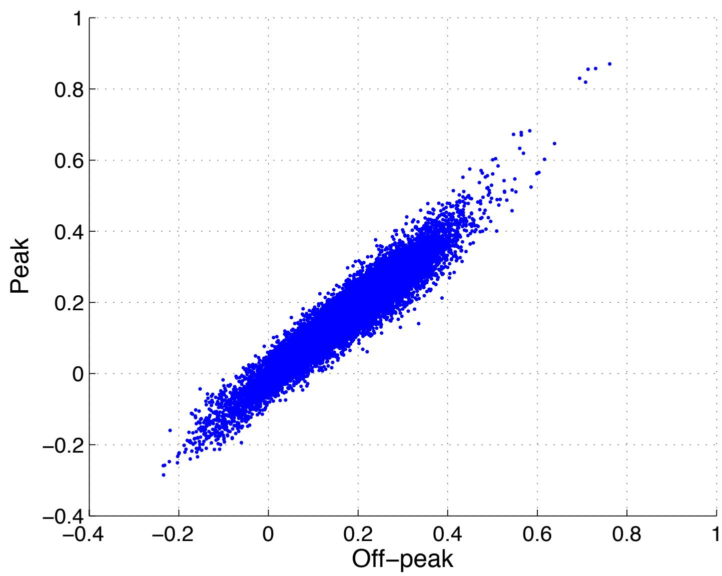

Figure B.1: Correlation between ratings on peak and offpeak days.