Producer Exploration Generates Categories without Audiences*

Anthony Vashevko

October 9, 2019

Abstract

Category theory finds that markets partition producers into categories and that producers who do not fit one specific category—or who span multiple categories—perform worse than their single-category peers. The major thread of category theory argues that categorizations stem from the bounded rationality of market audiences, who are forced to impose categorizations and ignore miscategorized producers in order to efficiently interact with the market. I present an alternative model in which producers in a market segregate into categories and experience an apparent miscategorization penalty without reliance on an audience process: In an uncertain world, producers imitate successful predecessors. An *ex-post* rationalization process identifies clusters of imitators as categories. Categories reflect, but do not cause, producer success. This model of exploration of an uncertain world not only accounts for the basic findings of category theory but further describes how categories shift and emerge over time. These dynamics align with recent attempts to describe the evolution of categories through audience-driven processes. Throughout the model, the limited knowledge of producers, not that of the audience, drives the apparent penalty to miscategorization. I establish these results in a formal model and simulation.

*The latest version of this paper is available at tony.vashevko.com/research/#explore

Introduction

Category theory combines a theoretical proposition with an empirical observation. The theory describes how individual cognition processes create a pressure to conform (Zuckerman 2017; Zhao et al. 2017): Markets are full of countless ambiguous objects for people to sort through. People simplify their search for objects by categorizing them. Objects that are hard to categorize are hard to understand, evaluate, and ultimately consume (Zuckerman 1999; Hannan, Pólos, and Carroll 2007). Evidence that individuals categorize is well-established in cognitive psychology (Rosch et al. 1976) and marketing (Shocker et al. 1991; Roberts and Lattin 1991). The theory aims to explain an empirical pattern: markets themselves regularly attach categorical labels to objects, and objects with multiple or ambiguous labels face penalties. Restaurants with a single cuisine (Korean or Mexican) earn higher reviews than restaurants that mix cuisines (Korean-Mexican, Kovács and Hannan 2015). Movie audiences prefer single-genre movies (romance or horror) to multi-genre movies (romance-horror, Hsu 2006). The path from theory to phenomenon seems direct: The failures of miscategorized objects in a market presumably stem from failures of individual categorization. Market objects that fall outside market categories fail because individuals find it hard to understand them. A stream of research measuring the penalty to miscategorized—poorly- and multi-labelled—firms appears to support the story (Hannan 2010; Durand and Paolella 2013; Vergne and Wry 2014).

Yet an undercurrent in this work complicates the simple audience categorization process. The miscategorization penalty appears, in certain cases, to reverse (Smith 2011; Leung 2014; Sgourev and Althuizen 2014; Paolella and Durand 2016). Markets may harbor multiple audiences with different tastes (Kovács and Liu 2016), even audiences with differing preferences for categorical ambiguity (Pontikes 2012; Goldberg, Hannan, and Kovács 2016)—venture capitalists and some cultural critics appear to reward the ill-categorized. Further still, theoretical arguments suggest that individuals may resort to ad hoc categorizations instead of reliably classifying a particular producer in a particular way (Durand and Paolella 2013). Such complexities make a producers decision to conform much more uncertain than it may first appear: Conform to what standard? If penalized, by whom?

I propose an alternative explanation for the market-level penalty to miscategorization and the appearance of categories. Beyond the confusion of categorical constraints, producers entering a market face a situation of general uncertainty (March, Sproull, and Tamuz 1991; Levinthal 1997). The characteristics that lead to market success are numerous and have complex interactions—an organization may be facing constraining audiences, but it may also be developing novel technologies, discovering the talents of its personnel, elaborating its production process, or even succumbing to chance events. At the same time, competitive markets force producers to differentiate from each other to seek competitive advantage (Wernerfelt 1984; Teece, Pisano, and Shuen 1997), and perfect imitation of a peer may be infeasible (Rivkin 2000). These tensions—the need to differentiate and the risk to doing so—put producers in a competitive bind. On one hand, firms feel drawn to imitate successful predecessors to reduce their own uncertainty. On the other, competitive pressure forces them into the unknown. They must find a balance.[1]

Tension between imitation and differentiation will produce a marketplace defined by clusters of success and isolated failures. All producer efforts to differentiate will begin alone. Subsequent entrants will imitate successful predecessors and shun failures,

[1]Contrast this with optimal differentiation research stemming from institutional theory, which argues that producers' optimal distinctiveness involves a balance between competitive and institutional conformity pressures (Zhao et al. 2017). In the present article, producers balance competition against a generalized fear that excessive differentiation will lead to something going wrong.

so that successful positions in the market will become more dense over time. If audiences tend to label clusters of objects (c.f. Hannan, Pólos, and Carroll 2007), it will appear that well-categorized objects outperform ill-categorized objects. The market will appear to feature a miscategorization penalty, even though audience cognition played no special role in producing the market.

Causality here flows in reverse: miscategorization does not cause failure; failure causes miscategorization. As this article further shows, the dynamics of producer entry tend to entice and create categories around unexpectedly successful positions. These dynamics parallel a number of arguments in the categories literature that describe categorical emergence and change through a cognitive perspective (e.g. Kennedy 2008; Zhao et al. 2018).

The core of this article is a formal model and simulation. The model explores how producers react to uncertainty in their environment and exploration by their peers. This model describes a world in which producers sequentially explore a world while observing the behavior and outcomes of their peers. It solves for producers' optimal exploration behavior given their bounded understanding of the world at any given point in time. Finally, it shows how this optimal exploration behavior generates differential clustering around prior successes that reproduces the basic pattern observed in the literature on market categories.

Theory and Model

To understand this article's relation to category theory, it is helpful to examine category theory through the relations among three constructs. Objects in a market have a position, receive a categorization, and experience an outcome. Position denotes the characteristics of an object, including the ways it might appeal to various needs. Similar objects have close positions. Categorization is the way audience members perceive the characteristics of an object and the way they communicate these characteristics to each other through shared labels. Audience categorization may misrepresent 'true' characteristics of the object, so that audience members fail to understand what uses an object could be put to. Outcomes measure object success in the market, translating an object's characteristics into audience appeal, evaluations, prices, or sales.

The dominant, audience-driven, perspective in category argues that objects' categorizations mediate the link from their positions to their outcomes: Audiences categorize objects based on their characteristics, but the categorization process itself determines whether objects are successful in the market (see Fig. 1a). The seminal works of Zuckerman (1999) and Hannan, Pólos, and Carroll (2007) lay out the argument: Zuckerman (1999) describes how audience categorizations operate in a sorting process that selects against miscategorized objects, establishing the second link of fig. 1a. Hannan, Pólos, and Carroll (2007) locate objects in an abstract feature space—their position—and describe how audiences assign categorical labels to clusters of similar objects in the space. This establishes how subsequent objects become categorized, the second link of fig. 1a.

Empirical work in category theory relies on the structure provided by these two core pieces. The feature space model of Hannan et al. allows for increasingly sophisticated measures of producer categorization (e.g. Pontikes 2012; Pontikes and Hannan 2014; Kovács and Hannan 2015). Zuckerman (1999) establishes the plausible causal link from these categorizations to observable outcomes. Empirical work then proceeds by showing an association between measures of categorization and outcomes.

I argue that the link between categorization and outcomes is spurious. Categorizations and outcomes are jointly, but independently, determined by how producers choose to position themselves in the market (fig. 1b). A common cause, not a direct

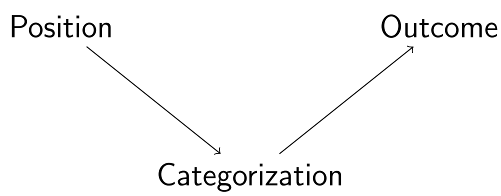

(a) Predominant logic of prior work

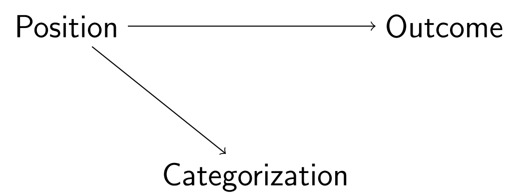

(b) Logic under producer exploration

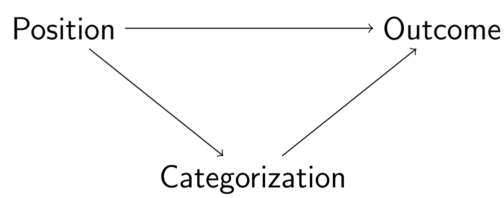

(c) Categorization as a partial mediator

Figure 1: Differing causal logics of past, present, and future work.

connection produces a correlation between the two. I argue that producer exploration generates a direct mechanism by which producers in tight clusters are more successful than those outside them. To the extent that an independent process of categorization exists to rationalize and make sense of the market, this process reflects but does not affect the results of producer exploration.

To the extent that categorization does determine producer outcomes, it must act as a partial mediator alongside the effects of producer exploration. As producers jockey for position in a market, audience categorizations may affect the outcomes they experience (see fig. 1c)—the role of empirical work is to measure the extent of this mediation. In more recent work, Zuckerman (2017) acknowledges that producers face such supply-side constraints and that excessive differentiation imposes strategic risks other than miscategorization; he maintains that categorization plays a primary causal role in producer outcomes. This argument does not, however, the numerous complications to audience categorizations that the literature has uncovered. Multiple audiences evaluate any given object (Kovács and Liu 2016), and they can differ in their taste for miscategorization (Pontikes 2012; Goldberg, Hannan, and Kovács 2016). Even individual audience members can categorize an object differently depending on the use to which they mean to put it (Durand and Paolella 2013) To the extent that it operates, audience categorization exerts no straightforward effect on producers' aggregate outcomes. The goal of my argument is to establish that strong category effects appear even in a world dominated by such inconsistent individual categorizations and market uncertainty.

Producer Exploration Generates Clusters of Success

Producers considering entry or expansion in a market face a difficult inferential problem. They can observe a set of peers or competitors that have made previous entries into the market. They can see the products their peers put into the market and some measure of how successful these forays were. But in general, producers have only a very limited understanding of what their peers are doing or how their behavior translates into observed levels of success (White 1981; March, Sproull, and Tamuz 1991; Levinthal 1997). Producers may have a poor understanding of how and why their own internal production processes work (Nelson and Winter 1982; Hannan, Pólos, and Carroll 2003; Bernstein 2012). Attempts to imitate dissimilar competitors may fail if producers fail to identify even a single key component of their competitor's strategy (Rivkin 2000). Competitors may rely on resources that are hard or impossible to duplicate (Wernerfelt 1984; Teece, Pisano, and Shuen 1997). Even if individual audience members rely on cognitive categorization to select partners, individual categorizations may fail to cohere into a meaningful whole.

Such uncertainty generates a critical tradeoff between imitation and exploration in market entry (c.f. March 1991): Close imitation of existing peers is likely to produce similar organizational outcomes; it comes at the cost of easy comparison and fierce competition against established peers. Differentiation, on the other hand, carries the risk of drastically misunderstanding a market or the production processes that lead to success in the market. An entirely novel business plan may end up being wildly successful, but it is just as likely to fail completely. Imitation offers a certain but more costly result. Exploration offers the chance of great success as well as great failure.

To the extent that organizations do attempt to imitate their peers, they will prefer to imitate successful peers rather than unsuccessful peers (c.f. Denrell and March 2001; Denrell and Le Mens 2007; Banerjee 1992; Strang and Macy 2001). On top of that, imitation of successful peers may reduce the burden of price competition: at similar levels of differentiation, a producer imitating a more successful peer should earn higher profits than a producer imitating an unsuccessful peer. Producers entering a market or changing their position then are likely to crowd around success and will crowd more around greater successes than around lesser successes. Over time, as producers enter a market, generally successful market positions will feature larger, denser clusters of producers than less successful positions.

Like category theory, this argument relies on a process of optimal differentiation (c.f. Zuckerman 1999; Zhao et al. 2017). But whereas audience-constrained producers fear differentiation because it carries the risk of taking them outside of categorical boundaries, producers here fear any of the ways that differentiation can go wrong. Businesses may fail and succeed not just because of socially shared categorizations, but for idiosyncratic or ephemeral reasons. A Russian restaurant may fail if it discovers that its prospective customers prefer more Americanized or more traditional food than expected. It may fail when it can't hire waiters to speak the language. Its chef's unique take on latkes might taste bad. Producers differentiate to avoid direct competition but fear differentiation because it robs them of control over how their product will land in the market. They differentiate until the risk of further differentiation becomes too high.

Categories and Labels

The above account offers no explanation for how categorizations and labels emerge in a market. Yet the ubiquity of labels and categories in markets would suggest that they serve some role. While cognitive theories of categories argue that social categorizations drive market outcomes, they may instead serve simply to facilitate communication about markets. As Hannan, Pólos, and Carroll (2007) argue, categories and labels may be most informative when they refer to well-defined, dense, and distinct clusters of similar objects. It is natural to assume that categories and labels emerge in ways that optimally fulfill this informative role.

Audience members need to communicate with each other, and to do so, they need to share a common language with common referents. Categories and labels serve as these common referents. Having a commonly understood set of restaurant categories makes the difference between inviting a friend for herring and sour cream, and inviting them for Russian food. It also helps the restaurant describe itself to me: I understand Samovar Russian Palace more easily than Samovar Herring and Sour Cream Hut. The absence of labels does not preclude meaningful communication: if I don't recognize the restaurant, I can look at their menu; if my friend is unfamiliar with Russian food, I can describe the dish. The presence of labels doesn't ensure quality: Samovar's chef may struggle to source ingredients; my friend may have no taste for preserved fish.

Efficient ways to communicate appear to emerge easily and spontaneously in groups of people (Weber and Camerer 2003). I argue that categories and labels primarily serve this communicative role in markets: labels emerge in order to efficiently describe producer positions in the market and they tend to emerge around dense and distinct clusters of producers. The market positions of products generates both market outcomes and labels that describe those positions, but the labels exert no independent effect on market outcomes.[2]

Model

This paper considers organizations operating in an uncertain environment in which optimal strategies are not readily apparent. Work in organizational theory has typically relied on NK landscapes to model such rugged performance environments (Levinthal 1997; Kauffman and Weinberger 1989). Several characteristics of NK landscapes render them unsuitable for the present context. First, this paper is critically concerned with the size and density of local producer clusters—NK landscapes are unable to support continuous distinctions among producers. Second, the structure of NK landscapes complicates the task of defining optimal search behavior, forcing researchers to rely on search heuristics.[3]

Instead, this paper models environmental uncertainty using a Brownian path. Such models have been recently adopted as an alternative framework for rugged fitness landscapes (Callander 2011; Ganz 2018; Callander and Matouschek 2014). I rely

[2] It must be pointed out that social classification is itself subject to interpretation as a strategic producer decision. Signalling is a well-established market dynamic (Spence 1973), and producers will adopt labels that aid their performance in the market (Podolny 2008; Etzion 2014; Pontikes and Kim 2017). This paper discusses the emergence of classification schemes in the absence of any market benefit to classification.

[3] Multi-arm bandit models suffer from the first issue, a lack of continuous positioning. PN landscapes, as recently proposed by Rahmandad (2019), suffer from the second, the difficulty of deriving optimal search behavior given bounded knowledge of the environment.

on an adaptation of the model. Producers operate in a market represented by the real number line: each point on the real number line represents some position in the market (c.f. Hotelling 1929; Salop 1979). The appeal of any position is given by the value of a one-dimensional Brownian path. Using the notation that the appeal of a producer at position x is given by W_x, the basic property of this appeal function is that the difference in appeal between two positions in the market is given by a normal distribution with variance proportional to the distance between positions:

$$W_x - W_y \sim N(0, (x - y) \cdot \sigma^2)$$

Positions that are close together will have relatively similar degrees of appeal to producers. A position that is far apart from another is likely to have either much higher or much lower appeal.

For a producer, the Brownian landscape is difficult to search over, and presents a degree of ruggedness similar to that of an NK landscape: local optimization is not guaranteed to identify global optima, and many local optima can exist in close proximity to each other. An optimal search strategy requires some plan to cope with the uncertainty of the environment. The key advantages of Brownian landscapes, however, are that organizations are able to locate arbitrarily close to one another, and that optimal search behavior is easier to derive. Appendix 2 compares Brownian and NK landscapes in greater detail, and suggests an underlying equivalence between the two models.[4]

In addition to the inherent appeal of a position, producers also face direct competition from neighboring producers. Here, competition is a greatly simplified representation of various competitive forces encouraging producer differentiation. This

[4]For an alternative use of Brownian walks in organizational research, see Levinthal (1991), Denrell (2004), or Le Mens, Hannan, and Pólos (2015).

can include price competition (D'Aspremont, Gabszewicz, and Thisse 1979; Salop 1979) among other institutional competitive pressures (c.f. Wernerfelt 1984). I assume that this competitive penalty depends only on the position of the nearest neighbors. Denoting the positions of all existing producers by in the market by X, the competitive penalty facing a producer at position x in the market is given by:

$$c(x, X) = \frac{1}{\min_{x<x_h \in X} |x - x_h|} + \frac{1}{\min_{x>x_l \in X} |x - x_l|}$$

The competitive penalty is nonlinear: as a producer becomes arbitrarily similar to her nearest competitor, the competitive penalty imposed by their proximity becomes arbitrarily high. Fig. 2 illustrates this penalty for a producer considering entry at a range of positions either to the right of a single competitor or between two competitors spaced one unit apart. The competitive penalty approaches infinity as a producer attempts to perfectly imitate their nearest neighbor.

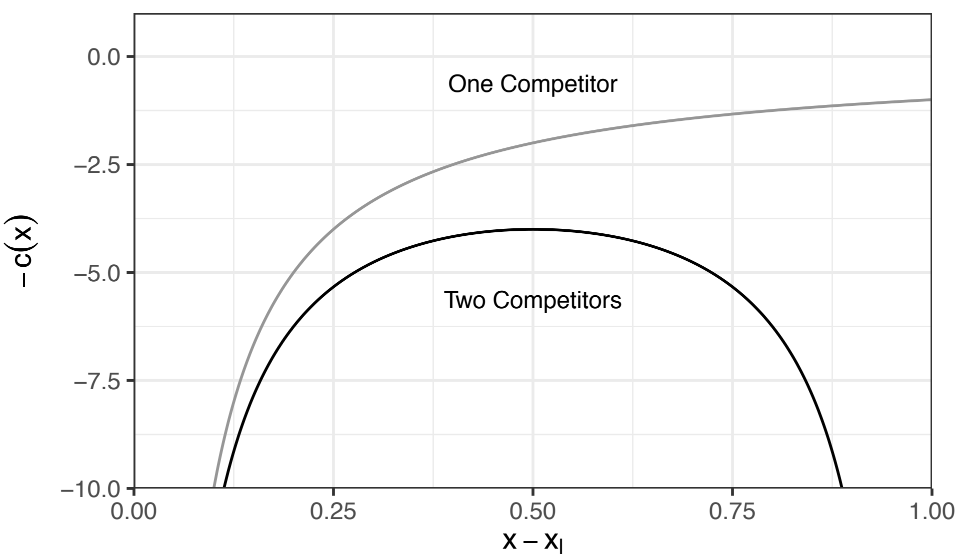

Figure 2: Competitive penalty with competitors at $x = 0$ and $x = 1$

Finally, producers are risk averse in their decision-making process, so that they evaluate the value of a given position using some risk-averse utility function u (Pratt 1964):

$$u(x, X) = u(W_x - c(x, X))$$

Stochasticity appears in this function in the form of W_x: for any position x that has not yet been attempted by some producer, W_x is a random variable whose value is determined by the Brownian motion described above.

Producer risk aversion is important to these results. Because the landscape is given by a Brownian walk, the distribution of risk is symmetric at any distance from existing points, and because the walk is drift-free, the distribution is symmetric about zero for producers at the edge of the landscape. A risk-neutral producer would be indifferent between taking on no risk by copying a competitor perfectly, and taking on arbitrary amounts of risk by differentiating. Competition, though, would induce risk-neutral producers to differentiate as much as possible from incumbents; a producer with competitors on only one side would attempt to differentiate by jumping infinitely far from incumbents.[5]

Producers enter the world one at a time. Some initial producer (or set of producers) begin at an arbitrary predetermined position in the space with some initial level of appeal. At every subsequent point in time, a new producer attempts to enter the market in a way that maximizes their expected utility. These new entrants do not know the shape of the appeal function: they can only observe the positions and realized

[5] N.B., in this model producers end up clustering more around successful positions than around unsuccessful positions, so that successful positions end up facing more future competition than unsuccessful positions. Forward-looking producers may expect greater competition upon success than upon failure. Such future expectation may have the effect of making risk-neutral producers act risk-averse in that imitative competition will reduce the net present value of successful positions more than it reduces the net present value of failed positions. In lieu of deriving these expectations, I simply assume risk aversion.

successes of incumbent producers. Given the positions of incumbent producers at time t, denoted $X^t = \{x_1, x_2, \ldots, x_{t-1}\}$, and the observed appeal at those positions, denoted $W^t = \{W_{x_1}, W_{x_2}, \ldots, W_{x_{t-1}}\}$, a producer enters at a position x that maximizes their expected utility:

$$\max_{x \in \mathbb{R}} E\left(u(x, X^t)\right)$$

Uncertainty in this expected value calculation stems entirely from a producer's uncertainty about the appeal of any given position. As an entrant considers positions arbitrarily far from any incumbent, the chance that the appeal of this position will be very low increases. As an entrant attempts to push out further from her competitors, she increases the chance that her position will be much worse than her most similar competitor. Fear of this risk pulls producers towards their competitors, while competitive pressure pushes them away, so that producers end up taking positions at some intermediate distance from their peers.[6] This optimal distance will depend on the relative prospects of positions within and outside of clusters, and it will vary from period to period.

Once an entrant enters at their preferred position x they discover the true value of their position. I assume, however, that relocation is prohibitively expensive, as producers may have committed to illiquid or irreplaceable resources. Producers are not able to move. This assumption may be consequential: If producers could move, preferentially migrate towards local peaks and depopulate local valleys; they

[6]This model is indebted to the Hotelling line model of a market. Hotelling (1929), however, predicts a 'principle of minimum differentiation': producers should end up imitating each other almost perfectly in such markets. As D'Aspremont, Gabszewicz, and Thisse (1979) show, this prediction cannot be sustained in equilibrium, and instead construct an example in which a 'principle of maximal differentiation' holds. Salop (1979) likewise constructs an equilibrium line market with maximal differentiation between producers. Neven (1987) finds differentiation in a sequential entry model. In general, it is not clear whether either principle represents the natural prediction (Biscaia and Mota 2013).

may also take greater initial risks in order to exploit the value of learning. Results may depend on whether audience cognition reflects all producer attempts or only current positions: if producers fully migrate out of local valleys, the difference between successful and unsuccessful producers would diminish, though it may not disappear entirely. Nevertheless, the fundamental mechanism driving the results of this model should remain: producers can better tolerate competition in successful positions than unsuccessful ones, allowing them to cluster more densely around high-appeal than around low-appeal positions.[7]

Simulation

Deriving the optimal behavior of individual producers in such a market is relatively straightforward. Deriving the properties of the market as a whole is more difficult. Instead, I simulate multiple instances of such markets in order to describe their general properties. Appendix 1 discusses the full details of the simulation. In brief, I simulate 1000 markets, and simulate the entry of 250 producers into each market. I seed each market with an initial set of producers, identical across all markets. Subsequent entrants make optimal entry decisions given their knowledge of existing producers' appeal. Their own entry then provides additional information to the next generation. The appeal function ends up distinct in each market: if two entrants into different markets happen to enter at the same position x, they will almost certainly end up with different levels of appeal.

I use a specific functional form for utility, $u(m) = am - \exp(-bm)$. It can be shown that this utility exhibits risk aversion, with risk aversion increasing in b. Further, u

[7] If producers are aware of their ability to move, they may use initial moves to learn about the environment in order to maximize the value of their future moves. This greatly complicates the derivation of optimal behavior, in that producers would now be solving a dynamic optimization problem, potentially in the presence of movements by competitors. The fundamental clustering pressure in the environment should, nevertheless, persist.

exhibits decreasing absolute risk aversion (Pratt 1964), which ensures the existence of the expected utilities that producers must calculate for the present results to hold (Callander and Matouschek 2014).

Producers must evaluate their utility either on 'open' intervals, in which they have only one competitor because they are searching on the very edge of a market, or on 'bridge' intervals, in which they have competitors on both sides. On 'open' intervals, producers maximize their expected utility by selecting Δ, the distance to their nearest competitor, at x. On 'bridge' intervals, producers maximize their expected utility by selecting Δ, the distance to their leftmost competitor, at x_l. For convenience, the distance from their rightmost competitor, at x_r is given by $\bar{\Delta} = x_r - x_l - \Delta$. The expected utility takes the form:

$$E(u(\Delta)) = aM(\Delta) - \exp\left(-bM(\Delta) + \frac{1}{2}b^2V(\Delta)\right)$$

where $M(\Delta)$ is the expected appeal at Δ, and $V(\Delta)$ is the contribution of variance to the expected utility. On open intervals, these equal:

$$M(\Delta) = W_x + c(\Delta)$$
$$V(\Delta) = \Delta\sigma^2$$

On bridge intervals, these equal:

$$M(\Delta) = W_{x_l} + \frac{W_{x_r} - W_{x_l}}{x_r - x_l} \cdot \Delta + c(\Delta) + c(\bar{\Delta})$$
$$V(\Delta) = \frac{\Delta \cdot \bar{\Delta}}{x_r - x_l}$$

Entrants evaluate their expected utility at all available intervals and enter market in the interval and with a Δ that maximizes their expected utility.

Modeling the Categorization Process

The model described above only encapsulates the positions and outcomes of producers in the market. To model the behavior of audience cognition in these markets, I attempt to identify dense clusters of producers. Each cluster would likely represent a distinct category of producers, with producers falling outside such clusters representing producers that are difficult to categorize or that fall into multiple categories.

I identify such clusters by fitting a finite Gaussian mixture model to the positions of producers in each market (Dempster, Laird, and Rubin 1977; Xu and Wunsch II 2005). A finite Gaussian mixture model is a general tool for classifying points into similar clusters. Characteristics of this model make it particularly appealing for modeling category membership.

This model assumes that the positions of a set of points are given by some set of normal distributions, $\{N(\mu_1, \sigma_1^2), \ldots, N(\mu_k, \sigma_k^2)\}$, centered at different means and with potentially different variances. Roughly, the model assumes that each point is generated by first randomly selecting one of the k distributions, proportionally to its probability p_i. The position of the point is then drawn from the chosen distribution. Such mixtures tend to place the centers of clusters at particularly dense parts of the set. I select the number of clusters k using the Bayesian Information Criterion (Schwarz 1978; Steele and Raftery 2009).

There are two major advantages to modeling clusters with a probabilistic mixture model, as here. Unlike hierarchical or partitioning (e.g. k-means) cluster models, a probabilistic model estimates the likelihood of each point being in each cluster.

This allows for a measure of each point's grade of membership in each cluster—high grade-of-membership positions are those close to the center of the cluster and which the cluster is likely to generate. In addition, it allows for the identification of points that straddle clusters, points whose grade of membership across clusters is roughly equal.

Fig. 3 gives an example of one simulated market, the clusters identified in it by the Gaussian mixture model, and the associated measures of category membership described below. Panel **A** shows the positions of producers and the appeals W of their positions; the bottom of the panel shows the local density of producers across the market. Panel **B** shows the probability density functions of the two clusters identified in the market: the two clusters align with the peaks seen in Panel **A**.

I define the grade of membership GOM of a point x in cluster i as the logarithm of a point's predicted likelihood of being in a cluster, $p_i(x)$, normalized by the peak predicted likelihood of the entire cluster:

$$\mathrm{GOM}_i(x) = -\left(\max_z \left[\log p_i(z)\right] - \log p_i(x)\right)$$

where z is the point with the highest grade of membership in the category. Comparison against the peak-likelihood point controls for differences in cluster width: points in more diffuse clusters have lower likelihoods on average. Panel **C** of fig. 3 shows the grade-of-membership functions associated with each cluster identified in the example market: even though Cluster 2 is more diffuse than Cluster 1 (panel **B**), the most central producers in each clusters have identical grades of membership.

I also construct measures of the extent to which a point lacks categories or spans multiple categories. For every point, I consider the two clusters in which the point has its highest grades of membership. Taking $p_1(x)$ and $p_2(x)$ to be the highest and second

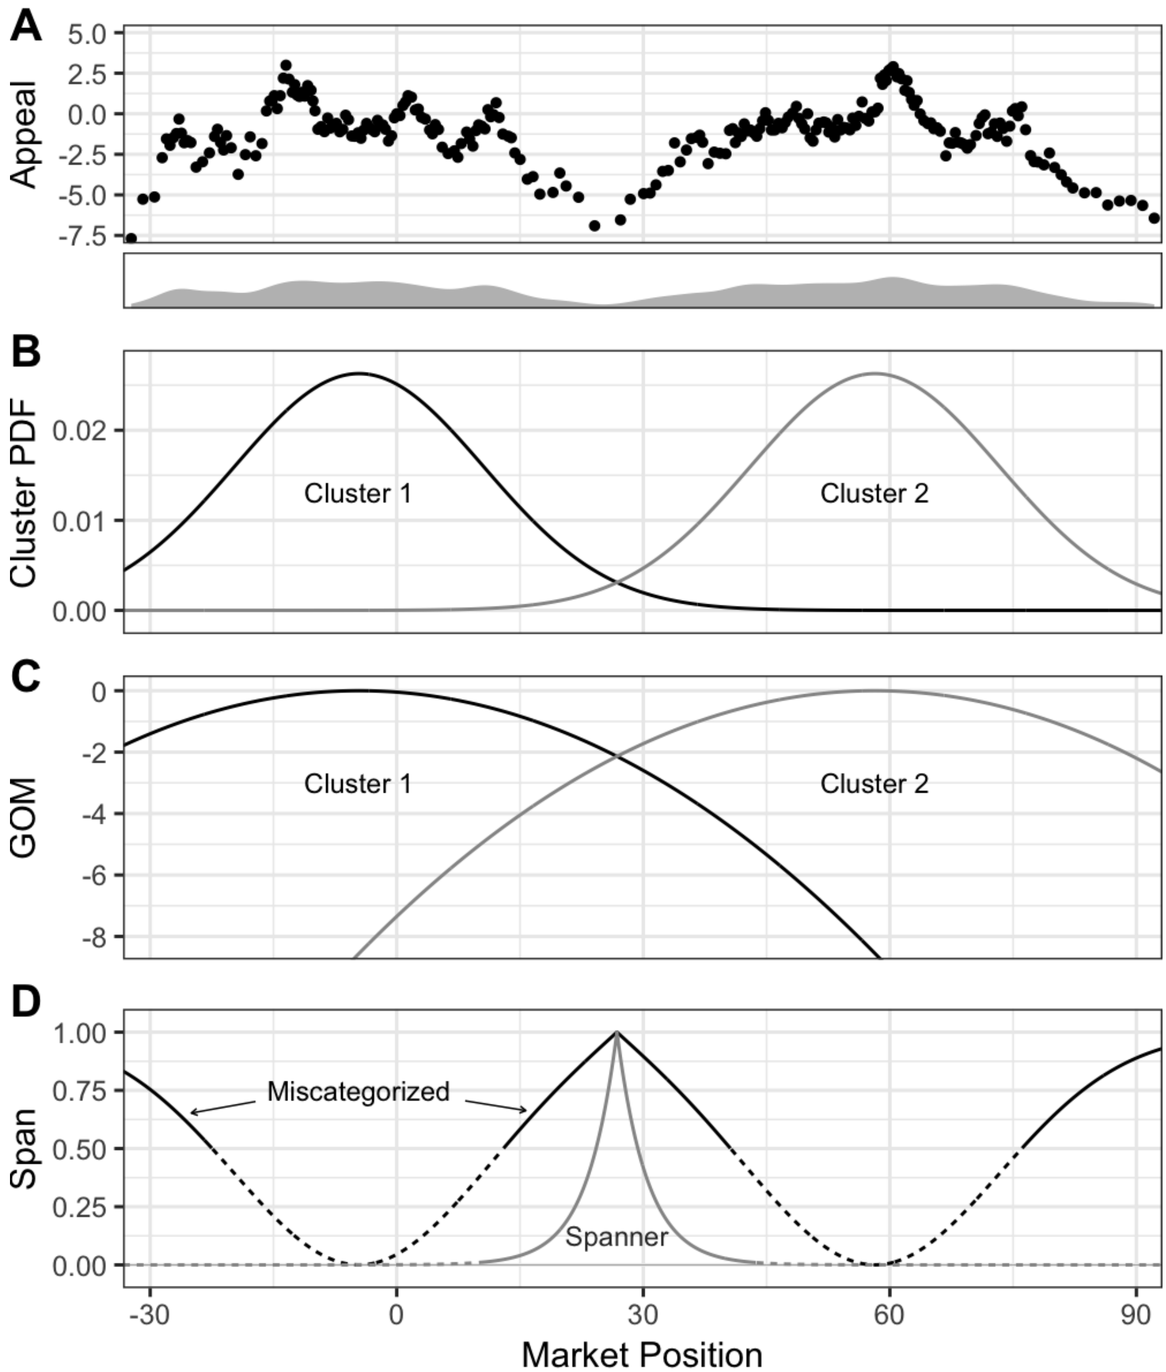

Figure 3: Example of a simulated market, estimated clusters, and measures of category membership

highest predicted likelihoods for the point, I first define a measure of miscategorization and a measure of spanning, as depicted in Panel **D** of fig. 3:

$$\text{Miscat}(x) = 1 - (p_1(x) - p_2(x))$$
$$\text{Span}(x) = \frac{p_2(x)}{p_1(x)}$$

Under these definition, points are miscategorized either if they are unlikely to belong in any cluster or if they are about equally likely to belong to multiple clusters. On the other hand, points are spanners if they are similarly likely to belong in multiple clusters. Next, I reduce these to binary measures, so that a point x is considered miscategorized if $Miscat(x) > 0.5$, and it is considered a spanner if $Span(x) > 0.01$. The particular thresholds were chosen by inspection, but analysis is robust to variation around these specific values. Finally, I construct a third binary measure of non-categorization if a point is miscategorized and is not a spanner. This allows me to examine the effects of spanning and non-categorization separately.[8]

Results

Existence of a Miscategorization Penalty

The aim of this article is to show that a process of optimal market entry under uncertainty can generate the market-level categorization phenomenon described by

[8]The regressions I consider below include *Spanner* and *Non-Categorized* as predictors. Using *Spanner* and *Miscategorized* as predictors instead causes the coefficient on spanner to reflect the effect of spanning net of the underlying effect of miscategorization. All else equal, spanning points tend to occur in denser, higher value regions in which two clusters exist close together, which causes *Spanner* to predict a positive effect of spanning relative to the miscategorization alone. Treating spanning and non-categorization as mutually exclusive allows for closer comparison to prior research.

category theory. As described above, the prediction made by category theory is that well-categorized producers are generally more successful and appealing than ill-categorized producers. Well-categorized producers tend to fall into dense clusters of similar producers, whereas ill-categorized producers tend to fall outside of or between such dense clusters.

There are three ways of examining this proposition in the markets simulated here. First, if the categorization penalty holds, then category-spanning positions will tend to be less appealing to the audience than non-spanning positions. More fundamentally, positions with lower grade of membership in any category should be less appealing to the audience than positions with higher grade of membership in any category. That is, positions which generally do not fit in any category should generally underperform positions that fit in some category. Most fundamentally, the miscategorization argument can be applied directly to positions in feature space. Insofar as categories signify dense clusters of positions, the miscategorization penalty should be stronger for more isolated positions, and weaker for positions in dense clusters. The miscategorization penalty would predict that a producer's distance from his predicts lower appeal. The propositions below summarize these three arguments:

Proposition 1 (Category spanning)

Positions spanning multiple categories will have lower appeal than positions that do not span multiple categories.

Proposition 2 (Grade of membership)

Positions with a lower grade of membership in their nearest category will have lower appeal than positions with a greater grade of membership in their nearest category.

Proposition 3 (Distance)

Positions with greater distance to their neighbors will have lower appeal than positions with lower distance to their neighbors.

I test these relationships with linear regressions estimating the effect of each measure of category membership on the appeal of the position. I estimate the effects within each period of each market. Specifically, within each period t and each market i, I estimate the regressions

$$W_{x,i,t} = \alpha + \beta_{1,i,t} NoCategory_{x,i,t} + \beta_{2,i,t} Spanner_{x,i,t} + \epsilon_{x,i,t}$$
$$W_{x,i,t} = \alpha + \beta_{i,t} GOM_{x,i,t} + \epsilon_{x,i,t}$$
$$W_{x,i,t} = \alpha + \beta_{i,t} Distance_{x,i,t} + \epsilon_{x,i,t}$$

where x indexes all producers within a market during the period, and each regression considers one of the measures of category membership: distance to nearest neighbor; grade of membership; or distance to nearest neighbor. Thus, in each market and each period, I estimate three separate β coefficients representing how miscategorization predicts producer appeal.

In each case, larger β represents a larger effect of category membership on producer appeal. Propositions 1-3 predict negative β for distance to nearest neighbor and category spanner status, and positive β for grade of membership: category-spanningness, decreasing grade of membership, and increasing distance all reduce producer appeal.

Fig. 4 plot these coefficients across all markets over time. Fig. 4a shows the difference between category spanning and non-spanning positions; fig. 4b shows the effect of increasing grade-of-membership; Fig. 4c shows the effect of an increasing distance to neighbor on producer appeal. Pale lines track how the relationship evolves within each individual market over time, as more entrants come into the market. The thick middle line show the average value of the coefficient across all markets within a period, with a 95% confidence interval represented by a white band around the line. Finally two thin gray lines show the relationship at the 10th and 90th percentile of markets in each period.

Each of the figures supports the predicted relationships. Category spanners generally have lower appeal—the estimated effect of spanning is generally negative within markets and it is negative when aggregating across all markets. Positions with high grade of membership in some category generally have higher appeal. Most fundamentally, isolated positions generally have lower appeal. These relationships appear to strengthen over time as markets approach something like their equilibrium state.

Within the context of these models, these effects may be better termed as a correlation between clustering and appeal than as a miscategorization penalty: the correlation reflects producer tendency to imitate success instead of the limits of audience cognition. Markets in which producers imitate success naturally generate a positive correlation between density and appeal. When audiences assign high category membership to dense clusters, ill-categorized producers tend to also have low appeal, reproducing an apparent miscategorization or category spanning penalty.

This correlation between clustering and appeal, however, is not absolute. As a close look at the thin lines of fig. 4 suggests, many individual markets experience periods in which the correlation between clustering and appeal disappears. Looking at markets at least 50 periods away from initial conditions, 54 percent of markets and 6 percent of individual market periods experience an inversion of the penalty as measured by grade-of-membership; 80 percent of markets and 22 percent of individual market periods experience an inversion as measured by category-spanner status. A

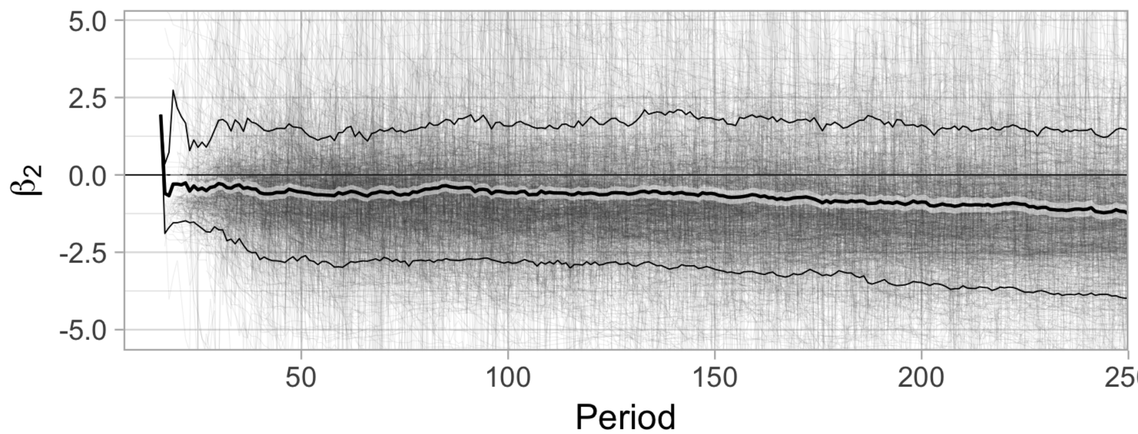

(a) Effect of category spanning. $W = \alpha + \beta_1 \cdot NoCategory + \beta_2 \cdot Spanner$

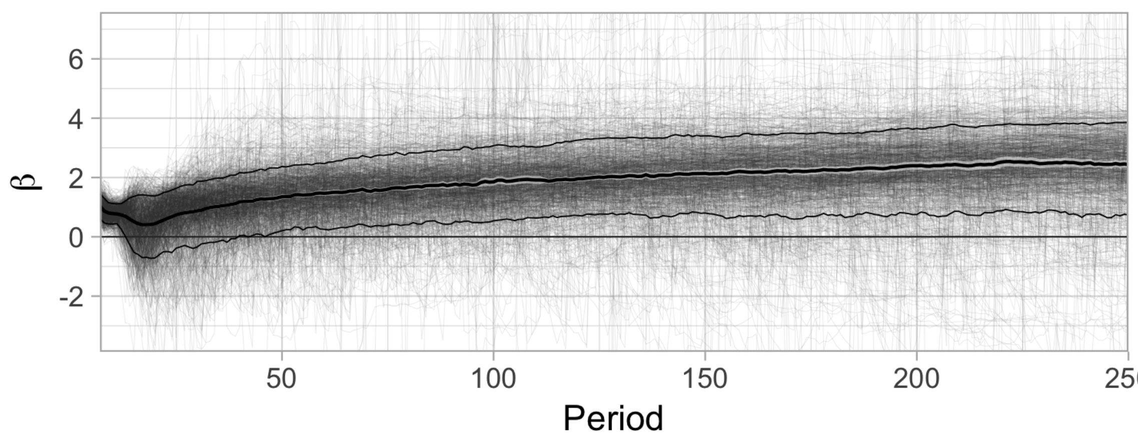

(b) Effect of increasing grade of membership. $W = \alpha + \beta \cdot GOM$

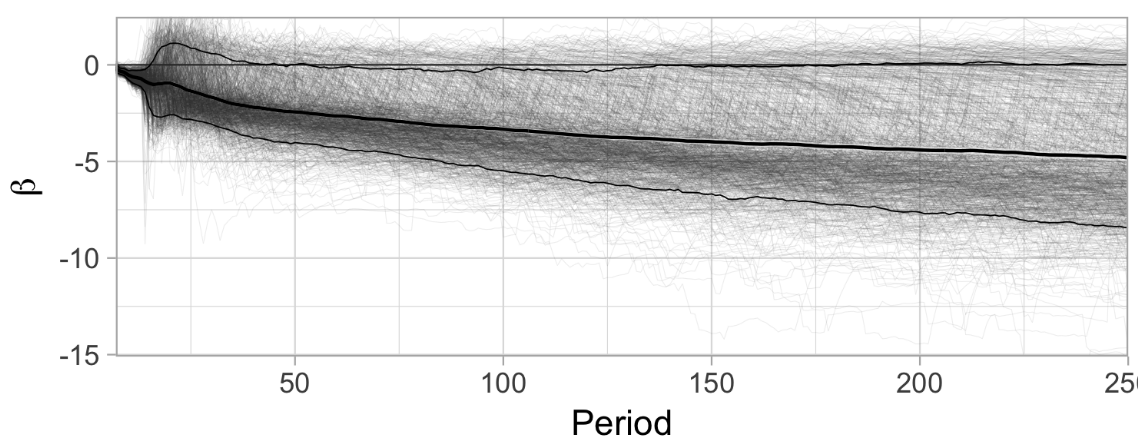

(c) Effect of increasing distance to neighbor. $W = \alpha + \beta \cdot distance$

Figure 4: Effect of position characteristics on position appeal, 1000 simulated markets, 250 periods.

fair amount of the time, the most successful producers in a market will be those outside of major clusters or categories. Such inversions occur entirely because the appeal function is unknown: as producers discover highly appealing regions outside of existing clusters, they find themselves simultaneously outside of existing categories and in positions of high appeal. If many producers discover success, enough producers may wind up outside existing clusters that the overall positive correlation between clustering and appeal disappears or reverses.

It bears repeating that the appeal function, and the audience members and production processes which generate it, do not depend on the categorical structure of the market. Categories reflect producer positions but do not influence where producers choose to enter the market. The correlations reported here arise entirely from producer willingness to imitate success, no matter the sources of that success. It may be tempting to interpret areas of high appeal as representing latent categories. Producer appeal may stem from audience evaluation and may, in the extreme, even result entirely from an audience's cognitive inability to evaluate certain kinds of producers. It remains the case that producers have no awareness of these categories above and beyond their knowledge of the appeal function. Audience members, too, have either a limited understanding of such latent categories or a very limited ability to convey their understanding to producers. Neither producers nor the audience can predict the appeal of a novel object, above and beyond its similarity to prior examples.

The categories and labels that exist in practice have a limited ability to predict these latent features. More importantly, while the clustering-appeal correlation is a typical feature of markets, it is not an absolute but a probabilistic relationship: miscategorized producers experience lower appeal only when existing categories closely align with peaks in the appeal function. The process of producer imitation tends to reinforce this alignment. When the alignment breaks down, the clustering-appeal relationship, and any apparent miscategorization penalty, attenuates or disappears. The analysis below explores this dynamic instability of the categorical system further.

Change and Emergence of Categories over Time

The cluster model estimated here not only estimates the position of clusters but the number of clusters as well. This allows for an investigation of category dynamics in markets.

The literature on categories has a mixed record of dealing with category emergence. Much of the formal theory treats categories as *a priori* constructs in the market (e.g. Zuckerman 1999; Hannan, Pólos, and Carroll 2007). A separate strand of the literature has discussed category and form emergence as a highly intentional process. Some research discusses the role of social movements in creating new categories (Lee, Hiatt, and Lounsbury 2016; Weber, Heinze, and DeSoucey 2008; Carroll and Swaminathan 2000). Some discusses the role of influential vanguards (Koçak, Hannan, and Hsu 2014; Rao, Monin, and Durand 2005; Ruef 2000). Scholars are generally pursuing mechanisms of category emergence (Glynn and Navis 2013). What these approaches share is an emphasis on the intentionality of category emergence and a stress on category existence as a prerequisite for market action. If the existing system of categories is a poor fit for the market, some group of actors in society must typically create the category before producers can take advantage of it.

In the model presented here, categories follow producers instead of leading them. The category system shifts to reflect the reality of producer behavior as it changes over time. This allows for a specific prediction for the timing of category change. The number of categories in a market may increase as new categories emerge, or may decrease if two existing categories combine into one. As discussed above, though the clustering-appeal correlation is a common feature of markets, it is not ubiquitous, and markets often undergo periods in which this correlation weakens or reverses. A weakening (or reversing) clustering-appeal correlation indicates that the existing system of categories poorly describes the set of positions that producers have found to be appealing to the audience. Such situations are ripe for categorical reconfiguration and emergence.

Proposition 4 (Category Change)

The number of categories in a market is more likely to change when the clustering-appeal correlation is weakening.

If the proposition holds, the number of categories in a market is more likely to change when the clustering-appeal correlation has experienced a sustained negative trend. I evaluate the relationship in each period of each market as described in the previous section, measuring the effect of decreasing grade-of-membership and the effect of spanning multiple categories. I measure trend by the overall change in the estimated β coefficient for each effect across the prior 25 periods, with appropriate sign changes to reflect the discussion here. A negative trend implies that the correlation between clustering and appeal is weakening (i.e. the miscategorization penalty is weakening); a positive trend, that the correlation is becoming stronger. I then predict the probability that the number of categories will change this period as a function of trend.

Tbl. 1 presents the estimates of these logistic regressions. The table shows that a recent positive trend (increasing clustering-appeal correlation) strongly predicts stability in the number of categories in a given period. If instead producers discover success at the borders of existing categories, weakening the penalty, the category system will change to reflect their discovery.

The theoretical implications of this proposition are straightforward: changes in

	res1	res2	res4a.1	res4a.2	res4a
(Intercept)	-2.840^{***}	-2.855^{***}	-2.968^{***}	-2.970^{***}	-2.969^{***}
	(0.010)	(0.009)	(0.012)	(0.012)	(0.012)
Penalty Trend (Distance)	-1.720^{***}				
	(0.125)				
Penalty Trend (GOM)		-1.464^{***}			
		(0.121)			
Penalty Trend (No Cat.)			-0.515^{***}		-0.466^{***}
			(0.071)		(0.073)
Penalty Trend (Spanner)				-0.440^{***}	-0.307^{**}
				(0.096)	(0.098)
Num. obs.	219000	219000	151837	151837	151837

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 1: Trends in the clustering penalty predict category change

the set of categories will naturally track the exploratory and imitative behavior of producers, even if producers are ignorant of the category system. There are interesting empirical implications in this proposition as well insofar as the miscategorization penalty is observable in markets. Markets with a weaker miscategorization penalty likely offer opportunities for further innovation, or already feature a set of innovating producers. This proposition also offers an alternative way to understand the market-maker/market-taker distinction drawn by Pontikes (2012). Pontikes argues that different audiences may have different tastes for categorical ambiguity, showing that venture capitalists are more favorable to category spanners than are general audiences. The mechanism described here instead suggests that venture capitalists may be operating in systematically different markets than the general audience. More specifically, VCs may have identified a set of markets in which innovation and categorical reconfiguration are within reach, and existing clusters do not strongly predict successful or appealing market positions. VC preferences do not lead them to ignore miscategorization, nor do their abilities enable them to overcome the cognitive limits of categorical thinking. Instead, their preferences and skills lead them to identify and invest in markets in which categorical boundaries are misaligned. Similar dynamics may cause the preference for atypical hedge funds identified by Smith (2011).

Stability of Categorical Boundaries

Producer exploration affects not only the system of categories as a whole but also the categorical fit of individual producers. The categorical system described here reacts to the explorations of producers. As producers explore the environment, some end up luckier than expected, and others less lucky. Subsequent entrants end up facing a substantially similar set of market opportunities as their immediate predecessors but for the lessons learned by those predecessors. If their predecessors were unlucky, subsequent entrants will pursue the next-best opportunities available in the market. Lucky predecessors, however, inform the market of a new set of available opportunities. When an entrant is luckier than expected, she will inspire a flow of close imitators. Since the category system moves to reflect dense clusters of producers, unexpectedly lucky producers will pull nearby categories towards themselves or will generate entirely new categories.

Proposition 5 (Category Seeding)

Unexpectedly successful entrants will become more categorically central following entry.

A formal model can offer a deep view into the decision-making of actors in a market. In particular, the model is aware of each producer's expected appeal at the position they chose to enter at. Once they enter, producers observe the position's true appeal. I construct a measure of the difference between realized and expected appeal:

	GOM	10 Per.	25 Per.	Miscat.	10 Per.	25 Per.
(Intercept)	-0.421^{***}	-0.772^{***}	-0.677^{***}	-0.840^{***}	-0.036^{***}	-0.232^{***}
	(0.000)	(0.001)	(0.000)	(0.000)	(0.002)	(0.002)
ΔW	0.091^{***}	0.078^{***}	0.115^{***}	-0.214^{***}	-0.132^{***}	-0.226^{***}
	(0.000)	(0.001)	(0.001)	(0.001)	(0.003)	(0.002)
Time		0.031^{***}	0.012^{***}		-0.069^{***}	-0.030^{***}
		(0.000)	(0.000)		(0.000)	(0.000)
Time$\times\Delta$W		0.010^{***}	0.001^{***}		-0.028^{***}	-0.004^{***}
		(0.000)	(0.000)		(0.000)	(0.000)
Num. obs.	31354000	2674000	6154000	31354000	2674000	6154000

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 2: Effect of unexpected appeal on category centrality

$$\Delta W_x = W_x - E(W_x)$$

I estimate the effect of ΔW on a producer's grade of membership in the nearest category as well as the likelihood of the producer being a category spanner over the 10 or 25 periods subsequent to producer entry. I truncate the tracking period because the long-term effect of success is more ambiguous than the short-term effect: over long periods of time, successful entrants may spawn yet more successful imitators, so that a category that at first moves towards the successful entrant later ends up moving away from them. In the real world, consider how MySpace become the prototypical social network before its more successful competitor Facebook took that spot.

Tbl. 2 presents estimates of this proposition. Positive ΔW indicates unexpected success, while negative ΔW indicates unexpected failure. The *Time* variable measures the number of periods since producer entry. The first three models show the effect of unexpected success on grade of membership, and the second three models show the effect on the likelihood of being a category spanner. The interaction term shows how these effects evolve over time. The models show that unexpectedly high appeal raises grade of membership in the immediate future and lowers the likelihood of category spanning.

Like the dynamics of category emergence, this result has several theoretical and empirical implications. It strongly suggests that categorical boundaries and categorical prototypes may not be stable over time. As producers test the boundaries of existing categories and end up successful, existing categories will migrate towards their success and new categories will emerge around them. This reflects a process of passive sensemaking, as audience members learn which positions in a market convey peak performance (Antonio et al. 1999; Zhao et al. 2018). Moreover, the model suggests a potential, albeit difficult, empirical test of the proposition. Although prospective entrants' expectations are difficult to observe, it may be possible to identify natural experiments within potential entrants, or to estimate the expected appeal of novel product configurations, before observing their realized appeal. Successful entrants should inspire imitation and categorical change. Failed entrants should inspire caution and stability.

Discussion

The model described here should be understood as an alternative explanation to category theory that is both able to reproduce its results and that makes several distinguishing predictions. In principle, it should be possible to distinguish the producer exploration theory described here from an audience cognition theory of a market categorization penalty. At the very least, the model described here raises the empirical bar that an audience cognition theory of categories must meet.

To summarize the argument, the model examines producer entry decisions in a complex, uncertain environment. Producers enter the market with a limited set of information: they know the outcomes of prior entrants along with their market position or the strategy they pursued. From this, new entrants can venture a guess on how successful a new strategy might be, based on how similar it is to prior attempts. Producers fear both failure and the threat of competition. They choose to enter the market with some differentiation from incumbents. In doing so, the stream of entrants slowly explores the environment, learning which positions are successful and which not.

Key to the result is that producers imitating highly successful peers are more willing to tolerate close competition: losing profit to a competitor is easier when there is much profit than when there is little. As such, entrants tend to cluster more densely around successful, appealing positions, and less densely around unsuccessful positions.

Categories emerge to rationalize the outcome of this process. Categorization processes recognize existing clusters of producers and so correlate with dense positions. Because dense clusters form around appealing positions, categories also come to indicate successful positions, so that ill-categorized producers will tend to underperform well-categorized producers in the market. Markets generate a correlation between producer clustering and producer appeal, which may be appear as a miscategorization penalty or a category spanning discount. These penalties are only side effects of the interaction between producer exploration and an automatic categorization process. Categories play no causal role in the creation of a spanning penalty—they are merely an instrument for measuring it.

So far this argument reconstructs the predictions of category theory under an alternative mechanism. It makes at least two additional predictions that distinguish it. First, because categories track the success of producers, their meaning tends to migrate towards successful producers, especially as their success attracts imitators.

The only description of categorical boundary drift under an audience cognition theory of categories makes a different set of predictions (Pontikes and Hannan 2013). Second, the producer exploration mechanism identifies conditions under which the mis-categorization penalty weakens. This helps explain previous findings about differential audience preference for categorical ambiguity (e.g. Pontikes 2012), by arguing that such audiences engage in market selection, favoring markets where the penalty happens to be weaker. It also makes a specific prediction on the timing of category emergence—it occurs when category spanners start to perform relatively well.

The alternative formulation of categories described here can also help explain several puzzles in existing work on category theory. First, empirical identification of the category spanning penalty relies on observation of successful single-category producers as well as unsuccessful multi-category producers. Category theory gives no explanation for why such multi-category failures exist, or more precisely, for why they chose to enter the market with their specific poorly designed configuration. The mechanism here explains the motivation of such producers: they took a risk and failed.

Second, the producer-focused described here bears at least a nominal resemblance to the production-side category penalty described by Hsu (2006). The argument there goes that producers have limited resources to spread across multiple efforts. Because markets reward specialists, producers that specialize in a single category outperform generalists that spread their efforts across many domains. The puzzle here comes from the definition of a specialist. Hsu gives no means by which researchers could identify a "pure-type" market domain. That is, it is not clear whether romantic comedy films should be considered a specialist genre or a blend of romance and comedy: their success argues for the former and their name for the latter. The model described above provides a general way to avoid this problem: there are no inherent specializations in a market. Producers generalize when they attempt to bridge islands of success.

Such bridging tends to fail precisely because islands of success sit atop the *most* successful possible positions in a market. If producers instead attempted to generalize by bridging islands of failure, we would expect generalists to outperform specialists.

The last point of discussion concerns a potential counterargument to the proposed model. One of the critical components of this model is the stochastic appeal function describing the terrain that producers explore. This terrain abstracts away from the details of why one producer's strategy succeeds and another's fails. A critical reader could object that this model simply takes the logic of category theory and buries it in a mysterious function with a different name. Cognitive limits on categorization may be a major or even a dominant contributor to the appeal of a given position.

This abstraction should instead be understood as a critical empirical challenge to category theory. An established literature in organization theory, with traces through recent work, argues that nobody really understands how organizations work, including the people that run them (March and Olsen 1976; March, Sproull, and Tamuz 1991; Levinthal 1997; Hannan, Pólos, and Carroll 2003). Many idiosyncractic factors feed into the success or failure of a particular strategy, so that a random model of organizational outcomes may be the appropriate approach. In particular, any latent effect of categories must exist alongside all of the other factors feeding organizational success. Given that a theory of producer exploration of a random terrain generates the same empirical predictions as the existing category theory, category theory must look for ways to distinguish itself above and beyond the empirical finding of a mis-categorization penalty (see Fig. 1c).

More fundamentally, if audience-side conformity pressure plays a major role in determining product appeal in the model presented here, this would generate a category theory entirely alien to the one understood in existing work. Intuitively understood categories may well determine why audience members appreciate some products and avoid others. Audience members, however, appear to have very little ability to communicate their internal preferences to the producers that must fulfill those preferences. Whether this stems from some barrier to communication between producers and the audience, or whether it stems from an inability by the audience itself to know what it wants, it is hard to see what predictions such a theory of categories could make.

Conclusions

The model presented here describes a world in which producers take risks in order to explore an uncertain world. This world reproduces many of the key empirical predictions of category theory. This suggests at the very least that category theory must look for additional empirical patterns in order to support its theoretical mechanisms. It appears as though the null hypothesis of category theory has been that categorical penalties should not exist in markets. This research resets the null: category penalties and apparent pressures to conform will randomly appear throughout the world.

While this may appear to be a hostile conclusion to a theory of categories relying on audience categorization, the model described suggests avenues to advance a category theory further. Such efforts may in fact be more fruitful than an attempt to adjudicate which model of the world is more correct. First, one of the key processes in this model is that producers maintain and react to an information set about the opportunities available in the world. In the model presented here, this information set is accurate. A category theory may instead advance arguments on how categories bias and shape producer exploration efforts, and prevent spanning from occurring in the first place. It is important to note that such advances will still find it difficult to distinguish categorically-constrained information from unbiased information sets. It will be difficult to discern the social construction of opportunity from an objective perception of opportunity (Zuckerman 2012).

Perhaps the more fruitful path forward is a reconfiguration of category theory from a cognitive to a political theory of the market (c.f. Fligstein 1996). While the mechanism presented here argues that coalitions and social movements are not required for category formation, in practice, social movements seem to be involved in the creation or maintenance of categorical boundaries (e.g. Carroll and Swaminathan 2000; Rao, Monin, and Durand 2005; Weber, Heinze, and DeSoucey 2008). Such movements may be necessary precisely in order to prevent or encourage producer willingness to push categorical boundaries. Actors may use categories as organizing symbols to advance their material ends. Categories in such a world would serve not as an indicator of cognitive failure but as a coordinating device for recognizing and reacting to market threats. Expansion of the theoretical structure of a category theory will build a stronger theory.

References

Antonio, José Rosa, Joseph F. Porac, Jelena Runser-Spanjol, and Michael S. Saxon. 1999. "Sociocognitive Dynamics in a Product Market." *Journal of Marketing* 63: 64–77.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107 (3): 797–817.

Bernstein, E. S. 2012. "The Transparency Paradox: A Role for Privacy in Organizational Learning and Operational Control." *Administrative Science Quarterly* 57 (2): 181–216. doi:10.1177/0001839212453028.

Biscaia, Ricardo, and Isabel Mota. 2013. "Models of Spatial Competition: A Critical Review." *Papers in Regional Science* 92 (4): 851–71. doi:10.1111/j.1435-5957.2012.00441.x.

Callander, Steven. 2011. "Searching and Learning by Trial and Error." *American*

Economic Review 101 (October): 2277–2308. doi:http://www.aeaweb.org/aer/.

Callander, Steven, and Niko Matouschek. 2014. “Managing on Rugged Landscapes.”

Carroll, Glenn R., and Anand Swaminathan. 2000. “Why the Microbrewery Movement? Organizational Dynamics of Resource Partitioning in the US Brewing Industry.” *American Journal of Sociology* 106 (3): 715–62.

D’Aspremont, C., J. Jaskold Gabszewicz, and J.-F. Thisse. 1979. “On Hotelling’ S "Stability in Competition".” *Econometrica* 47 (5): 1145–50.

Dempster, A.P., N.M. Laird, and D.B. Rubin. 1977. “Maximum Likelihood from Incomplete Data via the EM Algorithm.” *Journal of the Royal Statistical Society. Series B (Methodology)* 39 (1): 1–38.

Denrell, Jerker. 2004. “Random Walks and Sustained Competitive Advantage.” *Management Science* 50 (7): 922–34. doi:10.1287/mnsc.1030.0143.

Denrell, Jerker, and Gaël Le Mens. 2007. “Interdependent Sampling and Social Influence.” *Psychological Review* 114 (2): 398–422. doi:10.1037/0033-295X.114.2.398.

Denrell, Jerker, and James G. March. 2001. “Adaptation as Information Restriction: The Hot Stove Effect.” *Organization Science* 12 (5): 523–38.

Durand, Rodolphe, and Lionel Paolella. 2013. “Category Stretching: Reorienting Research on Categories in Strategy, Entrepreneurship, and Organization Theory.” *Journal of Management Studies* 50 (6): 1100–1123. doi:10.1111/j.1467-6486.2011.01039.x.

Etzion, Dror. 2014. “Diffusion as Classification.” *Organization Science* 25 (2): 420–37.

Fligstein, Neil. 1996. “Markets as Politics: A Political-Cultural Approach to Market Institutions.” *American Sociological Review* 61 (4): 656. doi:10.2307/2096398.

Ganz, Scott C. 2018. “Hyperopic Experimentation: Organizations Learning About Managers Learning About Strategies.” Preprint. SocArXiv. doi:10.31235/osf.io/aycnf.

Glynn, Mary Ann, and Chad Navis. 2013. “Categories, Identities, and Cultural Classification: Moving Beyond a Model of Categorical Constraint.” *Journal of Management Studies* 50 (6): 1124–37. doi:10.1111/joms.12023.

Goldberg, Amir, Michael T. Hannan, and Balázs Kovács. 2016. “What Does It Mean to Span Cultural Boundaries? Variety and Atypicality in Cultural Consumption.” *American Sociological Review* 81 (2): 215–41. doi:10.1177/0003122416632787.

Hannan, Michael T. 2010. “Partiality of Memberships in Categories and Audiences.” *Annual Review of Sociology* 36 (1): 159–81. doi:10.1146/annurev-soc-021610-092336.

Hannan, Michael T., László Pólos, and Glenn R. Carroll. 2003. “The Fog of Change: Opacity and Asperity in Organizations.” *Administrative Science Quarterly* 48 (3): 399–432.

———. 2007. *Logics of Organization Theory: Audiences, Codes, and Ecologies*. Princeton: Princeton University Press.

Hotelling, Harold. 1929. "Stability in Competition." *The Economic Journal* 39 (153): 41–57.

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

Kauffman, Stuart A., and Edward D. Weinberger. 1989. "The NK Model of Rugged Fitness Landscapes and Its Application to Maturation of the Immune Response." *Journal of Theoretical Biology* 141 (2): 211–45. doi:10.1016/S0022-5193(89)80019-0.

Kennedy, Mark Thomas. 2008. "Getting Counted: Markets, Media, and Reality." *American Sociological Review* 73 (2): 270–95. doi:10.1177/000312240807300205.

Koçak, Özgecan, Michael T. Hannan, and Greta Hsu. 2014. "Emergence of Market Orders: Audience Interaction and Vanguard Influence." *Organization Studies* 35 (5): 765–90. doi:10.1177/0170840613511751.

Kovács, Balázs, and Michael T. Hannan. 2015. "The Space of Categories and the Consequences of Category Spanning." *Sociological Science* 2: 252–86. doi:10.15195/v2.a13.

Kovács, Balázs, and Chengwei Liu. 2016. "Audience Structure and Status Multiplicity." *Social Networks* 44: 36–49. doi:10.1016/j.socnet.2015.07.008.

Lee, Brandon H., Shon R. Hiatt, and Michael Lounsbury. 2016. "Market Mediators and the Tradeoffs of LegitimacySeeking Behaviors in a Nascent Category." *Organization Science*, no. forthcoming: 1–57.

Le Mens, Gaël, Michael T. Hannan, and László Pólos. 2015. "Organizational Obsolescence, Drifting Tastes, and Age Dependence in Organizational Life Chances." *Organization Science* 26 (2): 550–70. doi:10.1287/orsc.2014.0910.

Leung, Ming D. 2014. "Dilettante or Renaissance Person? How the Order of Job Experiences Affects Hiring in an External Labor Market." *American Sociological Review* 79 (1): 136–58. doi:10.1177/0003122413518638.

Levinthal, Daniel A. 1991. "Random Walks and Organizational Mortality." *Administrative Science Quarterly* 36 (3): 397–420.

———. 1997. "Adaptation on Rugged Landscapes." *Management Science* 43 (7): 934–50. doi:10.1287/mnsc.43.7.934.

March, James G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2 (1): 71–87.

March, James G., and Johan P. Olsen. 1976. "Organizational Choice Under Ambiguity." In *Ambiguity and Choice in Organizations*. Universitetsforlaget.

March, James G., Lee S. Sproull, and Michal Tamuz. 1991. "Learning from Samples of One or Fewer." *Organization Science* 2 (1): 1–13.

Nelson, Richard R., and Sidney G. Winter. 1982. "Organizational Capabilities and Behavior." In *An Evolutionary of Economic Change Theory*.

Neven, Damien J. 1987. "Endogenous Sequential Entry in a Spatial Model." *International Journal of Industrial Organization* 5 (4): 419–34. doi:10.1016/S0167-7187(87)80004-8.

Paolella, Lionel, and Rodolphe Durand. 2016. "Category Spanning, Evaluation, and Performance: Revised Theory and Test on the Corporate Law Market." *Academy of Management Journal* 59 (1): 330–51.

Podolny, Joel M. 2008. *Status Signals: A Sociological Study of Market Competition*. Princeton: Princeton University Press.

Pontikes, Elizabeth G. 2012. "Two Sides of the Same Coin: How Ambiguous Classification Affects Multiple Audiences' Evaluations." *Administrative Science Quarterly* 57 (1): 81–118. doi:10.1177/0001839212446689.

Pontikes, Elizabeth G., and Michael T. Hannan. 2013. "The Geometry of Social Classification."

Pontikes, Elizabeth G., and Ruben Kim. 2017. "Strategic Categorization." In *From Categories to Categorization: Studies in Sociology, Organizations and Strategy at the Crossroads Research (Research in the Sociology of Organizations)*, edited by Rodolphe Durand, Nina Granqvist, and Anna Tyllström, 51:71–111. Emerald Publishing Limited. doi:10.1108/S0733-558X20170000051003.

Pontikes, Elizabeth, and Michael Hannan. 2014. "An Ecology of Social Categories." *Sociological Science* 1: 311–43. doi:10.15195/v1.a20.

Pratt, John W. 1964. "Risk Aversion in the Small and in the Large." *Econometrica* 32 (1): 122–36.

Rahmandad, Hazhir. 2019. "Interdependence, Complementarity, and Ruggedness of Performance Landscapes." *Strategy Science* 4 (3): 234–49. doi:10.1287/stsc.2019.0090.

Rao, Hayagreeva, Philippe Monin, and Rodolphe Durand. 2005. "Border Crossing: Bricolage and the Erosion of Categorical Boundaries in French Gastronomy." *American Sociological Review* 70 (6): 968–91.

Rivkin, Jan W. 2000. "Imitation of Complex Strategies." *Management Science* 46 (6): 824–44.

Roberts, John H., and James M. Lattin. 1991. "Development and Testing of a Model of Consideration Set Composition." *Journal of Marketing Research* 28 (4): 429–40.

Rosch, Eleanor, Carolyn B. Mervis, Wayne D. Gray, David M. Johnson, and Penny Boyes-Braem. 1976. "Basic Objects in Natural Categories." *Cognitive Psychology* 8

(3): 382–439. doi:10.1016/0010-0285(76)90013-X.

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

Salop, Steven C. 1979. “Monopolistic Competition with Outside Goods.” *The Bell Journal of Economics* 10 (1): 141–56. doi:10.2307/3003323.

Schwarz, Gideon. 1978. “Estimating the Dimension of a Model.” *The Annals of Statistics* 6 (2): 461–64. doi:10.1214/aos/1176344136.

Sgourev, Stoyan V., and Niek Althuizen. 2014. “"Notable" or "Not Able": When Are Acts of Inconsistency Rewarded?” *American Sociological Review* 79 (2): 282–302. doi:10.1177/0003122414524575.

Shocker, Allan D., Moshe Ben-Akiva, Bruno Boccara, and Prakash Nedungadi. 1991. “Consideration Set Influences on Consumer Decision-Making and Choice: Issues, Models, and Suggestions.” *Marketing Letters* 2 (3): 181–97. doi:10.1007/BF00554125.

Smith, Edward Bishop. 2011. “Identities as Lenses: How Organizational Identity Affects Audiences’ Evaluation of Organizational Performance.” *Administrative Science Quarterly* 56 (1): 61–94. doi:10.2189/asqu.2011.56.1.061.

Spence, Michael. 1973. “Job Market Signaling.” *The Quarterly Journal of Economics* 87 (3): 355. doi:10.2307/1882010.

Steele, Russell J., and Adrian E. Raftery. 2009. “Performance of Bayesian Model Selection Criteria for Gaussian Mixture Models.”

Strang, David, and Michael W. Macy. 2001. “In Search of Excellence: Fads, Success Stories, and Communication Bias.” *American Journal of Sociology* 107 (1): 147–82. doi:10.5465/APBPP.1999.27625893.

Teece, David J., Gary Pisano, and Amy Shuen. 1997. “Dynamic Capabilities and Strategic Management.” *Strategic Management Journal* 18 (March): 509–33.

Vergne, J. P., and Tyler Wry. 2014. “Categorizing Categorization Research: Review, Integration, and Future Directions.” *Journal of Management Studies* 51 (1): 56–94. doi:10.1111/joms.12044.

Weber, Klaus, Kathryn L. Heinze, and Michaela DeSoucey. 2008. “Forage for Thought: Mobilizing Codes in the Movement for Grass-Fed Meat and Dairy Products.” *Administrative Science Quarterly* 53: 529–67. doi:10.2189/asqu.53.3.529.

Weber, Roberto A., and Colin F. Camerer. 2003. “Cultural Conflict and Merger Failure: An Experimental Approach.” *Management Science* 49 (4): 400–415. doi:10.1287/mnsc.49.4.400.14430.

Wernerfelt, Birger. 1984. “A Resource-Based View of the Firm.” *Strategic Management Journal* 5 (2): 171–80. doi:10.1002/smj.4250050207.

White, Harrison C. 1981. "Where Do Markets Come from?" *American Journal of Sociology* 87 (3): 517–47.

Xu, Rui, and Donald Wunsch II. 2005. "Survey of Clustering Algorithms for MANET." *IEEE Transactions on Neural Networks* 16 (3): 645–78. doi:10.1109/TNN.2005.845141.

Zhao, Eric Yanfei, Greg Fisher, Michael Lounsbury, and Danny Miller. 2017. "Optimal Distinctiveness: Broadening the Interface Between Institutional Theory and Strategic Management." *Strategic Management Journal* 38 (1): 93–113. doi:10.15713/ins.mmj.3.

Zhao, Eric Yanfei, Masakazu Ishihara, P. Devereaux Jennings, and Michael Lounsbury. 2018. "Optimal Distinctiveness in the Console Video Game Industry: An Exemplar-Based Model of Proto-Category Evolution Optimal Distinctiveness in the Console Video Game Industry: An Exemplar-Based Model of Proto-Category Evolution. Organization Science," no. May. doi:10.1287/orsc.2017.1194.

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

———. 2012. "Construction, Concentration, and (Dis)Continuities in Social Valuations." *Annual Review of Sociology* 38 (1): 223–45. doi:10.1146/annurev-soc-070210-075241.

———. 2017. "The Categorical Imperative Revisited: Implications of Categorization as a Theoretical Tool." In *From Categories to Categorization: Studies in Sociology, Organizations and Strategy at the Crossroads Research (Research in the Sociology of Organizations)*, edited by Rodolphe Durand, Nina Granqvist, and Anna Tyllström, 51:31–68. Emerald Publishing Limited. doi:10.1108/S0733-558X201751.

Appendix 1: Simulation Details

This appendix discusses the details of the simulation and the derivation of its key equations.

The simulation models the course of 100 markets over 250 periods. In each market, producers can take positions along the real number line. The key characteristic of each market is the fitness function assigning producer fitness to specific positions along the number line. This fitness function is given by a Brownian walk W with drift parameter $\mu = 0$ and variance parameter $\sigma^2 = 1$. Zero-drift ensures that there's no overall direction of improvement to the market, while the choice of variance is arbitrary. Each simulated market is independent, with its specific course determined by its realization of the Brownian walk.

I seed each market with an initial set of six producers representing two initial clusters. One cluster of three producers is centered at position $x = 0$ with neighbors at -4 and 4, and a second cluster of producers is centered at $x = 50$. The central producers have $W = 0$ and their neighbors have $W = -1$, representing half of a standard deviation fall in the random walk. Alternative simulations starting at a single point $x = 0$ with $W(0) = 0$ show qualitatively similar results: the major difference is that the simulations presented here always begin with two initial clusters in the Gaussian mixture model.

In each subsequent period, one producer enters the market at some position. This producer can observe the positions of all previous entrants as well as the value of W at those positions. The new entrant chooses her entry position by conditioning on this prior knowledge; once she enters, she discovers the value of W at her chosen position, both for herself and for any subsequent entrants. An entrant enters at the position that maximizes her expected utility, as described below.

An entrant can enter either between two existing positions or at the extremes of the market (i.e. left of the leftmost producer or right of the rightmost producer). If the entrant enters at the edge of the market, we can denote the closest producer by x_0, with fitness given by $W(x_0)$, and denote the entrant's chosen position by δ if the entrant chooses to enter at $x_0 + \delta$ (or at $x_0 - \delta$ on the left of the market). If the producer enters between two existing positions, we can denote the leftmost and rightmost neighbors by x_l and x_r, with their fitnesses given by $W(x_l)$ and $W(x_r)$, and we can denote the entrant's chosen position by δ where the entrant chooses to enter at $x_l + \delta$. In period t of a market, t producers have entered the market, so that the new entrant must consider entry in $t + 1$ different intervals: 2 intervals at the left and right of the market, and the $t - 1$ intervals between each existing pair of producers in the market. In each of these $t + 1$ intervals, we can thus consider an optimal entry position denoted by δ.

Each producer derives utility from the amount of income m they receive at a given position according to the utility function $u(m)$. In the Brownian walk fitness landscape, the value of a given position follows a normal distribution with mean and variance given by distance from known positions. Within each interval, we define the mean function $M(\delta)$ and variance function $V(\delta)$. The expected utility at δ is given by:

$$U(\delta) = E\left[u(M(\delta) + \sqrt{V(\delta)}Z)\right]$$

The expectation is taken over Z, a standard normal variable. Chipman (1973) describes conditions on the utility function u that ensure that the expected utility exists. The utility functions I use here (and describe below) satisfy the conditions.

The entrant picks an optimal position within each interval by picking δ to maximize expected utility. For twice differentiable u, and differentiable M, V, these positions can be identified by setting the first derivative to zero, giving the following criterion condition:

$$
\begin{aligned}
0 = U'(\delta) &= E\left[u'\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right] M'(\delta) + \\
&\quad + \frac{1}{2}\frac{V'(\delta)}{\sqrt{V(\delta)}} E\left[u'\left(M(\delta) + \sqrt{V(\delta)}Z\right)\cdot Z\right] \\
&= E\left[u'\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right] M'(\delta) + \\
&\quad + \frac{1}{2}V'(\delta) E\left[u''\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right] \\
&= E\left[u'\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right] \cdot \\
&\quad \cdot \left[M'(\delta) + \frac{1}{2}V'(\delta)\frac{E\left[u''\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right]}{E\left[u'\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right]}\right]
\end{aligned}
$$

If $u' > 0$ everywhere, i.e. marginal utility is declining in income, the relevant criterion reduces to the second component:

$$
0 = M'(\delta) + \frac{1}{2}V'(\delta)\frac{E\left[u''\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right]}{E\left[u'\left(M(\delta) + \sqrt{V(\delta)}Z\right)\right]} \tag{1}
$$

With a specific utility function u and precise specifications for M and V, (1) can be simplified further.

Here, the actor's utility in income is given by

$$
u(m) = am - \exp(-bm)
$$

With this u, the expected utility function U and its derivatives reduce to:

$$
\begin{aligned}
E\left[u\left(M(\delta)+\sqrt{V(\delta)}Z\right)\right] &= u\left(M(\delta)-\frac{1}{2}bV(\delta)\right)+\frac{1}{2}abV(\delta) \\
&= aM(\delta)-\exp\left(-bM(\delta)+\frac{1}{2}b^2V(\delta)\right) \\
E\left[u'\left(M(\delta)+\sqrt{V(\delta)}Z\right)\right] &= u'\left(M(\delta)-\frac{1}{2}bV(\delta)\right) \\
&= a+b\exp\left(-bM(\delta)+\frac{1}{2}b^2V(\delta)\right) \\
E\left[u''\left(M(\delta)+\sqrt{V(\delta)}Z\right)\right] &= u''\left(M(\delta)-\frac{1}{2}bV(\delta)\right) \\
&= -b^2\exp\left(-bM(\delta)+\frac{1}{2}b^2V(\delta)\right)
\end{aligned}
$$

The mean and variance functions M and V derive from the characteristics of the Brownian walk W and a competition function c indicating the loss of income due to competition from nearby producers. Here I use

$$
\begin{aligned}
c(\delta) &= -\frac{1}{\delta} \\
c'(\delta) &= \frac{1}{\delta^2}
\end{aligned}
$$

On an open interval, in which the entrant has only one immediate neighbor at x_0, M and V take simple forms:

$$
\begin{aligned}
M(\delta) &= E\left[W(x_0+\delta)\right] + c(\delta) \\
&= W(x_0) + c(\delta) \\
M'(\delta) &= c'(\delta) \\
V(\delta) &= \delta\sigma^2 \\
V'(\delta) &= \sigma^2
\end{aligned}
$$

On a bridge interval, the entrant has a left neighbor at x_l and a right neighbor at x_r. With δ the distance from x_l, we denote the distance from x_r by $\bar{\delta} = x_r - x_l - \delta$. M and V take the following forms:

$$
\begin{aligned}
M(\delta) &= E\left[W(x_l+\delta)\right] + c(\delta) + c(\bar{\delta}) \\
&= W(x_l) + \frac{W(x_r) - W(x_l)}{x_r - x_l}\delta + c(\delta) + c(\bar{\delta}) \\
M'(\delta) &= \frac{W(x_r) - W(x_l)}{x_r - x_l} + c'(\delta) - c'(\bar{\delta}) \\
V(\delta) &= \frac{\delta\bar{\delta}}{x_r - x_l}\sigma^2 \\
V'(\delta) &= \frac{\bar{\delta} - \delta}{x_r - x_l}\sigma^2
\end{aligned}
$$

Substituting these values into (1) allows us to solve for the entrant's optimal entry point on each possible interval. The ratio of expected utilities $\frac{Eu''}{Eu'}$ is always negative here, and in practice the location of this root reflects the declining marginal effect of competition (M') matching the growing marginal effect of variance (V').

Finally, I identify these optima numerically with R's *uniroot* command (R Core Team 2016). The presence of the exponential function in u tends to cause numerical instability for extremely large or extremely small δ. In practice, (1) can be easier to solve after logarithmic transformation, looking for roots of the criterion

$$\begin{aligned} 0 &= -\log M'(\delta) + \log\left(-\frac{1}{2}V'(\delta)\frac{Eu''}{Eu'}\right) \\ &= \log V'(\delta) - \log(2) + \log\left(-\frac{Eu''}{Eu'}\right) - \log M'(\delta) \end{aligned}$$

with care taken to ensure that optimization is done over positive values of M' and V'.

Once the entrant has identified optimal entry points on each subinterval, she enters at the point with the highest overall expected utility.

References

Chipman, John S. 1973. "The Ordering of Portfolios in Terms of Mean and Variance." *Review of Economic Studies* 40 (2): 167–90.

R Core Team. 2016. "R: A Language and Environment for Statistical Computing." Vienna, Austria: R Foundation for Statistical Computing.

Appendix 2: Modeling Uncertain Organizational Environments

This paper models producer search behavior in a complex, uncertain environment. It models the environment as drift-free Brownian walk. Much of the recent literature on organizational search in complex environments has instead relied on the framework of NK landscapes (Kauffman and Weinberger 1989), which directly model complex interdependencies of the environment (Levinthal 1997; Rivkin 2000; Levinthal and Posen 2007; Siggelkow and Rivkin 2009; Levinthal and Workiewicz 2018).

This appendix discusses the choice of Brownian walk landscapes over NK landscapes. It reviews several characteristics of NK models and discusses how they complicate the modeling of agent behavior relative to Brownian walk models (or similar analytic models). This appendix also outlines an apparently novel equivalence result between NK landscapes and Brownian walks: It examines movements on a subset of NK landscapes as discrete random walks, showing that in the infinitesimal limit, they converge to Brownian motion. As such, Brownian landscapes allow for an examination of rational agent behavior on NK landscapes.

NK Models

The intent of NK models is to generate a "rugged" fitness landscape with local optima in which simple hill-climbing optimization heuristics would not perform very well. Organizational research using NK landscapes typically drops agents into such a terrain and demonstrates how some classes of search rules underperform others. Because the limitations of the model stem from its fundamental assumptions, it is helpful to review those here.

The most basic version of the NK model creates worlds based on the two parameters, N and K. The world has N binary dimensions: any position in the world can be described as some combination of N characteristics. In an $N = 3$ toy market, for example, a toy might be identified by whether or not it has a head, whether or not it has wheels, and whether or not it is red. A position can be concisely identified as a string or vector N items long containing only zeros and ones. A world has 2^N possible positions in it.[1]

Each position also has an associated fitness value, so that producers at position 000 might be better off than producers at position 010. The value of these positions is determined by the K parameter and a specific value assignment rule. The fitness of a position x is given by the sum of the fitnesses of all substrings of length K contained in x:[2]

$$V(x) = \sum_{i=1}^{N} f(x_{i,K})$$

$x_{i,K}$ represents the substring of x starting at the ith element of the vector and containing K total elements. For convenience, this substring operation loops back to the beginning of x if it runs out of values to draw from: if $i + K > N$, then $x_{i,K}$ contains the last $N - i + 1$ followed by the first $K - (N - 1 + 1)$ terms of x. For $x = 001$, and $K = 2$, we would have $x_{1,2} = 00$, $x_{2,2} = 01$, and $x_{3,2} = 10$. The function f assigns each possible substring some value and is a fixed characteristic of the world. Within a given NK world, each occurrence of the substring 00 in a position contributes the same value to the overall fitness of the position. The values of specific substrings

[1]For $N = 3$, these eight would be 000, 001, 010, 011, 100, 101, 110, 111.

[2]Here I use a specific construction of the NK landscape to aid subsequent derivation. In general, NK models can vary in whether the K parameter implies substrings of length K or of length $K + 1$. They can also vary in whether the fitness landscapes are given by the sum or the average of substring contributions.

are randomly drawn according to some distribution. In practice, this is often a draw from the unit interval uniform distribution, so that each substring contributes some random number between 0 and 1 to the overall fitness.

Values of K greater than one make the model interesting. If $K = 1$ then the world has two possible substrings, 0 and 1, and each has some value associated with it, $f(0)$ and $f(1)$. If an agent in this world learns that $f(1) > f(0)$ he should immediately jump to the all-ones position. With $K = 2$ such decisions become much more complicated. There are now four possible substrings (00, 01, 10, 11). An agent starting at 0000 might want to move to 0010. This changes two parts of the value function—it changes one 00 substring ($0\mathit{00}0$) to a 01 substring and a second 00 substring ($00\mathit{00}$) into a 10 substring. This improves the agent's fitness if the combination of the two changes is positive, $f(10) - f(00) + f(01) - f(00) > 0$. An attempt to replicate this success by moving to 0110 may backfire, however, because such a jump introduces the new 11 substring, and eliminates the old 00 substring. The agent will end up worse off if $f(11) - f(00) < 0$. Values of K greater than one represent environmental interdependencies that generate local optima and generally complicate the process of environmental search for actors in the environment.

Rational Actors in NK Models

Most if not all organizational studies involving NK models describe the behavior of agents following researcher-determined search rules. While these search rules draw on valid insights from organizational theory, they do not represent the behavior of a rational actor operating in such a terrain. That is, they do not represent the behavior of an actor that has looked at the rules by which this terrain is generated and deduced an optimal search process. As such, many of the search rules examined in work on NK organizational landscapes are vulnerable to the criticism that they may not reflect how even a boundedly rational actor would respond to such an environment. This failure stems from the great difficulty in deriving an optimal search rule on an NK landscape: the NK model does its best to put obstacles and hard decisions in the path of a researcher hoping to analyze optimal search rules.

The major question with a rational actor NK model is how much information about the world should the actor have. Should an actor know what N and K are? Should an actor know that the substring value function f is fixed within the world? Should an actor know their own position vector? In particular, should an actor know if they are a 0 or a 1 on dimension i, or whether dimensions i and $i+1$ are more related to each other than i and $i+10$?

Affirmative answers to any of these questions would appear to generate incentives for perverse actor learning behavior aimed at identifying and gaming the rules of the world. For small values of K, an actor can conduct or observe a relatively small number of market experiments in order to learn the value function f. An actor that knows the value of K can estimate how many such experiments to run. An actor that knows that f is fixed can attempt to learn all of its values through experimentation in order to deduce the global optimum.[3] Actors that do not know the details of their specific vectors but understand the rules of the world may nevertheless be able to deduce their position by some program of local experiments.

Reasonable modeling decisions will point to a restriction of agents' information sets to prevent such exploitation of the model. An actor whose rationality is bounded only in their incomplete knowledge of the terrain must decide on some number of dimensions of their strategy to change. Two features of the NK model will remain to

[3] With $K = 2$, for example, three experiments appear to suffice to determine the global optimum: a comparison of an all-zeros position, a single-one position, and an all-ones position should determine whether the global optimum is at all-zeros, all ones, or at an alternating series of 0's and 1's. Similar experiment should be possible for larger K.

characterize their behavior: first, they will experience an effectively random shift in their performance as a result of this change; and second, for $K < N$, a larger number of changes lead to a higher variance performance shift. While optimal behavior can be derived for such agents, it induces the additional modeling tedium of highly discretized decision-making: not only is the distribution of fitness changes discretized, but agents must also make an integral number of changes to their position vectors, when they might prefer a fractional change. For researchers, discrete optimization is much more analytically tedious than continuous optimization.

Brownian Walks

Brownian walks offer an alternative way to model a complex environment that features the two desirable characteristics of NK landscapes: deviations produce uncertain outcomes and larger deviations produce greater uncertainty, so that performance at nearby positions is correlated. A Brownian walk, W, has a drift parameter μ and a variance parameter σ^2 and is defined by the following relationship between every two points x and y, $y > x$:

$$W_y - W_x \sim N(\mu(y-x), \sigma^2(y-x))$$

The difference in fitness between any two points follows a normal distribution with drift and variance proportional to the distance between the two points. While some organizational research has used Brownian with drift to model certain organizational decisions (Callander and Matouschek 2014; Ganz 2018), a drift-free Brownian walk ($\mu = 0$) is a more reasonable choice for modelling a multi-organizational environment: in a drift-free landscape, all positions are *ex ante* identical in expectation.

Brownian walks make deriving rational search rules substantially simpler than do NK models. Brownian landscapes allow for continuous positioning, and the outcome at the end of a jump follows a normal distribution. Under common utility functions, expected utility allows for straightforward concave optimization. This allows a researcher to model behavior that is rational up to uncertainty about the external environment, so that organizational outcomes are driven purely by uncertainty, rather than by faulty decision-making.

Similar dissatisfaction with NK landscapes has led to alternate searches for more analytic landscapes (e.g. Fourier landscapes, Winter, Cattani, and Dorsch 2007).

An Equivalence Between NK and Brownian Worlds

The difference between NK and Brownian landscapes seems to make for a partisan modeling decision. The two classes of model appear to be modeling extremely different situations: NK models describe a world of extreme interdependency, while Brownian walk models describe a world of smooth variation. This section attempts to bridge the gap between the two classes of model by considering the behavior of NK models with infinite N, and showing that under certain assumptions, their behavior reproduces the key characteristics of a Brownian walk.

We begin by considering some properties of NK models under finite N. When an agent changes one point of her position, she changes K substrings: the substring beginning at the changed point, as well as the $K-1$ substrings immediately prior to the changed point. Assuming the fitness of each contribution of each substring is drawn from some distribution F, each of the K changes induces a change in fitness that follows F', the difference between two variables drawn from F. Under the bounded rationality assumptions described above, we treat each of these K changes as independent of each other and of the overall position.[4]

[4]More precisely, each point change will follow some derivative distribution F', with F' being a

If an agent makes two changes, analysis becomes more complicated. If the two changed points are very far apart, two pointwise changes will generate $2K$ substring changes. If the two changed points are near each other, however, their affected substrings will overlap. Since multiple changes to a single substring randomize its value as effectively as a single overlap, nearby point changes will affect anywhere from $K + 1$ substrings (in the case of changes to consecutive points) to $2K$ in the case of no overlap. We can continue this further: if the agent makes P point changes, she will change anywhere from $K + P$ to PK substrings. If she induces C substring changes, her fitness function will change by

$$S_C = \sum_{i=1}^{C} X_i, \quad X_i \sim F'$$

We can now consider expanding the scale of the landscape. We consider increasing N, as K remains fixed. In addition, we consider the agent changing some fraction p of all N positions, such that $p = P/N$. We can also consider the fraction of substrings changed, $c = C/N$. Under the bounded rationality assumptions, we may assume that the P changes are randomly distributed within the position vector, even if the agent is intentional about the changes she's making. As such, for any given set of changes, c is a random variable taking some value between $(K + P)/N$ and PK/N, depending on the degree of overlap. It is possible to derive a precise distribution for c using inclusion-exclusion counting: we add up the fraction of substrings changed without considering overlaps, then subtract a correction for all double-counted overlaps, then add back all triple-counted overlaps, continuing on with all appropriate corrections. This count will approximately equal:

discrete distribution taking on 2^{2K+1} distinct values, one for each possible combination of all values of K consecutive substrings.

$$
\begin{aligned}
c &= \binom{P}{1} \cdot \frac{K}{N} - \binom{P}{2} \cdot \frac{O(K^2)}{N^2} + \dots \\
c &= (pN) \cdot \frac{K}{N} - (pN)(pN-1) \cdot \frac{O(K^2)}{N^2} + \dots \\
c &= pK - p^2 \cdot O(K^2) + \frac{O(K^2)}{N^2} + \dots \\
c &\approx pK - p^2 \cdot O(K^2) + \dots
\end{aligned}
$$

Here the order function $O(x)$ indicates an upper bound given by $A \cdot x$ for constant A. The last approximation holds as $N \to \infty$, as K remains constant.

Finally, we can consider the total change in the fitness function induced by a change to a fraction p of points in the position vector. We assume that the variance of F' falls as N grows, such that $\mathrm{Var}(F') \cdot \sqrt{N} = \sigma^2$. This allows an application of a central limit theorem to argue that the total change in the fitness function is given by a normal distribution with mean 0 and variance $c\sigma^2$. By the counting principle above, for $p \ll 1$, higher order terms of p disappear, allowing for a first order approximation of c: $c \approx pK$. Thus, the variance of the fitness function equals $pK\sigma^2$: it is increasing in the number of changes as well as the degree of interdependence in the environment.

More precisely, the result follows almost immediately from application of Donsker's theorem (a generalization of the CLT to random walks): interpreting the finite sequence of changes S_C as a random walk, the rescaled walk

$$
W^{(n)}(c) = \frac{S_{\lfloor nc \rfloor}}{\sqrt{n}}
$$

converges to a Brownian walk on $[0, c]$ with variance parameter σ^2.

The Brownian model described in the paper relies on two features of the Brownian walk: search on open intervals of the Brownian walk, and interpolating search between known points of the Brownian walk (i.e. bridge intervals). The construction above maps onto open interval search almost directly: a search of length Δ corresponds to some search of some small length c in a NK landscape. Constructing an analogy to bridge search requires further intuition, but provides substantial insight into modeling decisions in the organizational search literature.

Brownian Bridges in an NK World

Suppose we operate in an NK landscape and know the positions and fitness values of two positions, x and x'. Because we can perfectly observe both positions, we can consider the n-length vector of differences between these positions, $D = x' - x$, where $D_i = x'_i - x_i$, the difference at the ith point of the position. More concisely, we can consider only the vector of non-zero differences between the two positions, $d = (d_1, \ldots, d_m)$, $d_i = D_{l_i}$, where l_i identifies the ith non-zero element of D.

We can consider an actor starting at x and implementing some or all of the changes in d. If she implements none of the changes, she remains at x, and remains at fitness level $V(x)$. If she implements all of the changes, she ends up at x' and its fitness level $V(x')$. She can, however, make only some subset of the changes, in which case we must make further assumptions about the world to describe the consequences of her behavior.

Per the bounded rationality assumptions described above, the elements of x and d have no particular order, and may be internally permuted with no effect on the actor. As such, we can consider what happens as the actor steps through these changes one by one, first making the change d_1, then the change d_2, etc. If $K > 1$, each individual change affects an entire substring of the position, and so by the arguments above, each change induces a random deviation to the fitness level, either positive or negative, relative to the overall trend in moving from $V(x)$ to $V(x')$. If $K = 1$, we cannot guarantee that each change affects a substring, and hence we must assume more directly that each change induces a random deviation from the overall trend. This should hold so long as the jump from x to x' was not perfectly optimized, or was taken in haste, so that each of the individual changes in d does not represent a strict increase in V.

The process of stepping through d then represents a random walk over a sequence of IID random variables following the distribution F' described above. By Donsker's theorem, this random walk converges to a Brownian walk as the number of positions N and the number of differences m increase, and insofar as the endpoint of the walk is known, this process converges to a Brownian bridge.

Discussion

By the equivalence described here, a change to a fraction p of N points in a large N NK model can thus be reinterpreted as a jump of length pK along a Brownian walk with variance parameter σ^2. Conversely, jumps on a Brownian landscape correspond to jumps in an NK world. To the best of my knowledge, this is a novel result in the organizational search literature, and it opens the door to a number of reinterpretations, reconstructions, and extensions of previous work.

For the purpose of the present paper, this equivalence allows the reader to import intuitions about NK landscapes into the Brownian motion setting. The characteristics of Brownian motion contribute two key advantages over modeling organizational behavior in an NK landscape: First, actors can engage in rational search behavior, which isolates the effect of environmental uncertainty on organizational outcomes. Second, actors are able to exploit a continuum of positions, allowing for a more precise examination of the effects of local density than NK landscapes allow.

References

Callander, Steven, and Niko Matouschek. 2014. "Managing on Rugged Landscapes."

Ganz, Scott. 2018. "Hyperopic Experimentation: Organizations Learning About Managers Learning About Strategies." Preprint. SocArXiv. doi:10.31235/osf.io/aycnf.

Kauffman, Stuart A., and Edward D. Weinberger. 1989. "The NK Model of Rugged Fitness Landscapes and Its Application to Maturation of the Immune Response." *Journal of Theoretical Biology* 141 (2): 211–45. doi:10.1016/S0022-5193(89)80019-0.

Levinthal, Daniel A. 1997. "Adaptation on Rugged Landscapes." *Management Science* 43 (7): 934–50. doi:10.1287/mnsc.43.7.934.

Levinthal, Daniel A., and Hart E. Posen. 2007. "Myopia of Selection: Does Organizational Adaptation Limit the Efficacy of Population Selection?" *Administrative Science Quarterly* 52 (4): 586–620. doi:10.2189/asqu.52.4.586.

Levinthal, Daniel A., and Maciej Workiewicz. 2018. "When Two Bosses Are Better Than One: Nearly Decomposable Systems and Organizational Adaptation." *Organization Science* 29 (2): 207–24. doi:10.1287/orsc.2017.1177.

Rivkin, Jan W. 2000. "Imitation of Complex Strategies." *Management Science* 46 (6): 824–44.

Siggelkow, Nicolaj, and Jan W. Rivkin. 2009. "Hiding the Evidence of Valid Theories: How Coupled Search Processes Obscure Performance Differences Among Organizations." *Administrative Science Quarterly* 54: 602–34.

Winter, Sidney G, Gino Cattani, and Alex Dorsch. 2007. "The Value of Moderate Obsession: Insights from a New Model of Organizational Search." *Organization Science* 18 (3): 403–19. doi:10.1287/orsc.1070.0273.