

Does the Middle Conform or Compete? Quality Thresholds Predict the Locus of Innovation*

Anthony Vashevko

July 10, 2018

Abstract

Where does innovation come from? This research models producer incentives to innovate with a focus on the role of audiences in constructing quality thresholds within markets. Market audiences create mechanisms for identifying the highest quality producers in a market. I highlight a key distinction between fixed quality thresholds (such as accreditations) and quality thresholds that respond to producer quality (such as rankings or best-of-breed awards). Producers evaluate how the inherently risky nature of innovation interacts with these thresholds. The model predicts conditions under which innovation emerges from the best producers in a market, from producers near the threshold in a market, from both, or from nowhere. Such predictions generalize and simplify several existing organizational theories of innovation.

*I thank Amir Goldberg, William Barnett, Glenn Carroll, Jacob Model, and seminar participants at the Stanford Graduate School of Business for their advice and helpful comments throughout this work. I thank the editor and reviewers for giving this piece a final kick it sorely needed. Any remaining errors are my own.

Introduction

There are markets where only the best get better. There are markets where mid-tier firms innovate and end up surpassing their superiors. There are markets marked by change throughout, where no firm chooses or can choose to keep its product constant. There are markets with no change. There is a locus of innovation in any market, but the locus changes from market to market. Organizational theory has no shortage of explanations. Law firms feature a pattern of middle-status conformity, with mid-tier firms sticking to established practice, while the best and worst firms adopt novel, unusual practices (Phillips and Zuckerman, 2001). Law schools, in contrast, face constraint from rankings, with schools throughout the rankings changing their existing strategies to try to climb higher (Sauder and Espeland, 2009). Competing firms in general may enter Red Queen races, running to out-innovate their rivals just to stay in the same relative place (Barnett, 2008). At the same time, firms in general may innovate only so long as they fail to meet some aspiration level of performance (Greve, 2003b). There is no clear guide for when to apply any particular theory to any particular market.

While existing theory traces how specific market structures generate specific patterns of innovation, this research lays out a general framework for predicting where in a market innovation will emerge. The availability of innovation opportunities exerts a global effect on all producers in the market while quality thresholds exert a local effect on nearby producers; the two interact to determine which producers will innovate and which will stay put. In particular, quality thresholds represent a novel generalization of a number of concepts previously discussed in organizational theory. This article argues that quality thresholds come in two kinds, and that the nature of quality thresholds in a market interacts with the favorability of innovation risks in the market to produce a typology of innovation. This typology describes the conditions in which innovation in the market will come out of the highest quality producers, from those near the threshold, from everywhere in the market, or from nowhere.

The framework relies on two types of uncertainty inherent in markets. Markets bring together producers, who present offers, and audiences, who must select among them. Because producers can be complex and difficult to evaluate (Zuckerman, 1999; Hannan et al., 2007; Vergne and Wry, 2014), audiences create mechanisms to simplify their decisions. These mechanisms can take the

form of category schemes (Zuckerman, 1999), more formal rating systems (Kovács and Hannan, 2015; Rao et al., 2005), status distinctions (Podolny, 1993; Gould, 2002), or best-of-breed awards (Rossman and Schilke, 2014). Such mechanisms serve to simplify an audience member’s evaluation and selection of particular producers by distilling their hard-to-discern quality or qualities into a discrete evaluation (March and Simon, 1958; Greve, 2003b)—is this particular producer one of the good ones or one of the bad ones? In general, such mechanisms create thresholds separating lower-quality from higher-quality producers. Thresholds help audiences more easily decide how to divide their scarce resources and attention among the producers that deserve them.

Producers react to quality thresholds. Producers must balance the threat of competition from similar peers against the threat of falling outside audience consideration if they offer a particularly unusual product (Deephouse, 1999; Zuckerman, 1999). More concretely, if producers attempt to differentiate through an unusual product offer, they may find themselves excluded or penalized in ratings and rankings, dropped from award consideration, or left out of informal categorizations. They will find themselves below audience thresholds and starved for resources. When producers decide to innovate—to change their product—they must wrestle with this constraint imposed by their audience.

Quality thresholds, however, come in two kinds, and the nature of the quality threshold determines how a producer must react. The article highlights this distinction between two ideal types of thresholds: Fixed thresholds, such as accreditations, ratings, or categorical rules, demarcate boundaries that are stable over time. Responsive thresholds, such as rankings, or best-of-breed awards, demarcate boundaries that change when the quality or identity of the best producers change—they grade producers on a curve. If all producers in a market get better, a fixed threshold such as an accreditation will reward more of them; a responsive threshold such as a top-10 award will only ever reward the top ten.¹ Conversely, if a subpar rival improves their quality, a top producer may be ejected out of the top ten but would not lose an accreditation. Consideration of quality thresholds adds to the risk inherent in all attempts at organizational change and innovation (Cyert

¹This distinction is similar to the difference between grading students on a rubric and grading students on a curve. A literature in education discusses this as the difference between “criterion-referenced evaluation” and “norm-referenced evaluation” (Glaser, 1963; Natriello, 1987). Findings there appear primarily to focus on student motivation and outcomes, and are hard to map directly into questions of innovation. Some findings do, however, suggest that norm-referenced (i.e. responsive threshold) grading tends to improve effort and learning by students in competition for the highest grades, as compared to criterion-referenced (i.e. fixed threshold) grading. Threshold dynamics would predict similar motivational effects.

and March, 1963; Levinthal, 1997; Hannan et al., 2003). Under fixed thresholds, producers above the threshold will avoid innovation for fear of failure, while those below the threshold will innovate with the hope of getting lucky (c.f. Kahneman and Tversky, 1979; March and Simon, 1958; Greve, 2003b). Under responsive thresholds, however, the threat that innovation by low-quality producers will raise the threshold forces producers above the threshold to innovate defensively.

Finally, opportunities for innovation may vary from market to market, so that producers may expect that innovation risks are generally favorable and likely to raise a producer's quality, or generally unfavorable, and likely to lower it. When innovation risks are generally favorable, producers would want to pursue them. When innovation risks are generally unfavorable, producers would prefer to avoid them. Quality thresholds complicate this calculus, inducing some producers to pursue unfavorable innovations and others to avoid favorable opportunities.

The combination of these two features—the nature of quality thresholds and the favorability of innovation opportunities—generates a typology of innovation in markets, predicting which segments of a market will choose to innovate. Depending on the nature of thresholds and the favorability of opportunities in a market, innovation may emerge only from the highest-quality producers, only from producers near the threshold, from both, or from neither. The end result is a framework for describing the role that audience boundaries and producer reactions to them play in determining how innovation emerges in markets. This framework simplifies and generalizes existing theories of innovation, showing how a number of existing theories of markets can combine to produce a single coherent description of market behavior.

I lay out this theory using a formal model and simulation of producer behavior in the presence of audience thresholds. Producers strategically evaluate the value of pursuing risky innovations. They make these decisions with the expectation that their competitors may likewise pursue innovations, and that the success and failure of their competitors affects audience thresholds in the market. I look for equilibrium patterns of behavior in these markets. The article begins by discussing the basic constructs of the theory, then moves into the construction and results of the formal model, before concluding on the general consequences of the framework.

Quality Evaluations in Markets

Audiences meet producers in a market. Audiences comprise consumers and key stakeholders, such as government agencies, employees, or market intermediaries like critics and analysts. Audiences hold resources that producers want, and producers make offers to the audience in order to access these resources (c.f. Pfeffer and Salancik, 1978; White, 1981; Zuckerman, 1999). The set of offers audiences must sort through is complex and heterogeneous, in part because producers face strong incentives to differentiate their offers to avoid direct competition (Deepphouse, 1999). Yet audiences must find some way to sort through these offers.

I assume that in any market, audience members can aggregate the various characteristics of any particular offer into some measure of quality. Given available information, an audience member can recognize a producer as high or low quality. This perception of quality need not reflect some ‘true’ quality of the producer (Podolny, 1993; Lynn et al., 2009). I assume further that it pays to be high quality: a producer generally recognized as high quality can expect to receive greater resource than a low quality rival.

Because gathering information on offers and evaluating them is difficult, most audience members cannot perfectly evaluate the quality of any particular producer and instead rely on various institutions to aid their search. They may rely on public rating systems (e.g. Kovács and Hannan, 2015; Rao et al., 2005), rankings (e.g. Sauder and Espeland, 2009), awards (Rossman and Schilke, 2014; Kaniel and Parham, 2017), intermediaries (Zuckerman, 1999), status distinctions (Podolny, 1993, 2001; Gould, 2002), or implicitly understood categorizations (Zuckerman, 1999; Hsu, 2006b; Hannan et al., 2007). Such institutions create sharp distinctions in the perceived quality of otherwise similar producers and can generate discontinuities in the rewards they receive from the audience. Oscar awards, for instance, produce sharp differences between winners and also-rans (Rossman and Schilke, 2014); mutual fund tier lists drive investment into the highlighted funds (Kaniel and Parham, 2017); even otherwise continuous school rankings generate discontinuities for colleges that make it “above the fold” (Bowman and Bastedo, 2009). In effect, such institutions act as quality thresholds that producers must meet in order to participate fully in a market.

Even if individual audience members differ in their preferences or evaluations of particular producers (Kovács and Liu, 2016), such centralized institutions work to synchronize individual

decisions. As such, they not only simplify the decision-making process of audience members, but produce a shared reference point that producers in the market can target.²

Risky Innovation

Producers have incentives to differentiate from peers and improve their quality in order to win greater rewards from the audience. They can do this by choosing to innovate, by changing features of their offer or production process. I assume that such changes are invariably risky (c.f. Kacperczyk et al., 2015)—insofar as producers had any safe changes to make, they would have made them already. Any further changes take them in the direction of exploring a complex landscape of possible offers (Nelson and Winter, 1982; Levinthal, 1997; Hannan et al., 2003). Producers may have observed a number of alternatives in that space, along with the success of those alternatives. But whether they choose to imitate existing peers or explore novel offers, producers’ ability to predict the consequences of their actions is limited. When a producer innovates, they may find that their new offer is more appealing to some subset of the audience than they expected, or that it is less appealing. They may find that their new offer synergizes with their existing production process, or they may find that it upsets key stakeholders (e.g. by drawing boycotts, King and Pearce, 2010).

The uncertainty of innovation interacts with the discontinuities created by audience quality thresholds. The discontinuity induced by quality thresholds may cause producers just below the threshold to innovate, hoping to get lucky and cross the threshold; conversely, producers above the threshold will avoid innovation since a failed effort may push them below the threshold and cause a sharp drop in the resources they receive from the audience. At the same time, actors well above or well below the threshold face little risk that any specific innovation will push them over the edge. The threshold constrains those actors closest to it, while both high- and low-quality have greater leeway to consider innovations on their own merits.

Though these predictions bear a similarity to those of performance feedback theory (March and Simon, 1958; Greve, 2003b), quality thresholds and aspiration levels reach their conclusions by considering opposite aspects of a market. In both cases, a discontinuity in an organizational

²For exposition, the discussion below focuses on markets with a single, well-defined quality threshold. The appendix discusses the robustness of these assumptions, showing that predictions hold if the threshold is imprecise or producers are uncertain of its exact location, and that they hold to a more limited extent in markets with multiple thresholds. The assumption of a single threshold is meant to represent a typical rather than universal market situation.

utility function drives risk-taking. Aspiration levels, however, reflect the bounded rationality of the organization (Simon, 1947), while quality thresholds reflect the bounded rationality of the market audience, and the producer response to quality thresholds represents a rational response to this bounded rationality. This drives a second distinction: while aspiration levels are internal and idiosyncratic to an organization, all producers within a market face the same quality thresholds. In this capacity, quality thresholds may represent an avenue for the synchronization of aspiration levels across organizations. The common empirical strategy of defining average performance as an organization’s ‘social aspiration level’ may thus indirectly capture the effect of market thresholds: some research in fact suggests that an organization below such a social aspiration level is more likely to innovate than an organization below its historical performance level (Baum et al., 2005).

More broadly, a focus on organizations’ shared market situation rather than their idiosyncrasies may account for inconsistent empirical patterns in the literature on aspiration levels. Performance feedback theory, for example, predicts that organizations well above their aspiration level may or may not engage in slack search—the use of excess resources for innovation and exploration. Research finds slack search in some markets (Greve, 2003a; Baum et al., 2005) but not others (Greve, 1998; Miller and Chen, 1994; Iyer and Miller, 2008). Though some work ascribes such differences to differences in organizational attention (e.g. Boyle and Shapira, 2012), empirical differences across markets may stem from differences between markets and not from the nature of organizations within them. In the present framework, innovation above the threshold may be driven not by slack but by the abundance of innovation opportunities within the market. Finally, the novel distinction that this article draws between fixed and responsive thresholds predicts that market thresholds can induce innovation by organizations immediately above them; aspiration levels offer no such prediction.

Fixed and Responsive Thresholds

The critical distinction this article makes is between quality thresholds that are fixed and those that are responsive. While ratings and rankings, for instance, both generate quality thresholds in markets, producers react differently to the two kinds of thresholds.

Fixed thresholds demarcate fixed positions in the quality distribution. They reward any producer that meets some absolute minimum level of quality. Most explicitly, mechanisms such as accreditations or health codes (e.g. Lehman et al., 2014) specify rules to determine whether any

particular organization meets the threshold. More informally, rating systems such as Yelp or IMDB may reflect consumer application of tacit but well-understood categorical rules (Hsu, 2006b; Hannan et al., 2007)—audiences know what features a particular product has to have and evaluate producer offers against these feature sets. The defining characteristic of fixed thresholds is that they can reward any number of producers—if all producers raise their quality, more producers will receive threshold rewards. Consider, for example, how restaurant health codes or pharmaceutical drug regulations work to reduce consumer uncertainty in their respective markets. While restaurants or drugs may compete with one another, the certification of some new drug or restaurant as safe does not indicate that incumbents are dangerous.

Unlike fixed thresholds, responsive thresholds shift in response to the changing quality of producers. Typically, they reward only some fixed number of the highest quality producers. Awards form the most prototypical example: only one movie wins the Academy Award for Best Picture in any given year, even if that year’s crop of movies were all particularly good or all particularly bad. The cutoff for actor rewards can, however, target any part of the quality distribution: the 40 positions of the French academy (presumably) reward France’s top 40 intellectuals (Merton, 1968); General Electric’s forced ranking system rewarded the 90 percent of employees with continuing employment (Stewart et al., 2010); mutual fund incentives reward managers that can beat the median (Brown et al., 1996). Tournaments and ranking systems in general consist of a series of responsive thresholds demarcating each rank, but they are also susceptible to sharp discontinuities at focal points in the ranking such as the front page, the top 10, or the top 100 (Bowman and Bastedo, 2009). Responsive thresholds bind producers above to those below the threshold: the success of inferior producers directly threatens the position of superior producers.

Fixed and responsive thresholds generate distinct patterns of innovation behavior. Fixed thresholds generate the well-known pattern of risk-taking below the threshold and risk-avoidance above the threshold: the former gamble on unlikely success, while the latter take all steps to avoid an unlikely failure. Producers below the threshold will innovate to get lucky, while producers above the threshold will stay the course to stay safe. Responsive thresholds generate a different pattern: under responsive thresholds, successful innovation by rivals below the thresholds threatens to raise an audience’s quality threshold and eject previously successful incumbents from the market. With sufficient innovation effort from below, this threat grows great enough that producers above the

threshold must innovate defensively in the hope that their own innovation efforts are more successful than those of the crowd below (c.f. Bothner et al., 2007).

Consider again the examples of law firms and law schools. Law firms face an audience that measures quality by the ability to deliver a set of specific services (Phillips and Zuckerman, 2001; Phillips et al., 2013)—law firms face a fixed quality threshold in the form of a categorical standard. Insofar as expansion into family law represents one possible threat to their ability to credibly deliver these services, they avoid expanding. Law schools, on the other hand, face an audience that measures quality using broad metrics aggregated into a single ranking, not through the adoption of specific practices that may meet the metrics (Sauder and Espeland, 2009). This induces schools to vary their practices in order to stay high in the rankings, even when such “innovations” involve attempts to cheat the system. While the audience threshold constrains the behavior of mid-status law firms, mid-quality schools may prefer to innovate insofar as innovation offers the opportunity of drastic upward jumps in the rankings.

Expected Value of Innovation

The last assumption I make is that producers in a market observe a common set of innovation opportunities. While the specific outcome of attempting an innovation will vary depending on the actor, all actors have a shared expectation of the probable outcomes of particular innovations—that is, actors have a common expectation of how much a given innovation will raise their quality. This common expectation of the probable outcomes of innovation defines the value of innovation in a market. When actors expect some set of innovations to improve quality on average, the expected value of innovation is favorable; when actors see no such set of innovations, the expected value of innovation is unfavorable.

Literature on complex landscapes provides a natural way to think of this kind of market-level risk (c.f. Levinthal, 1997; Rivkin, 2000). Producers have observed their peers’ strategies and the outcomes of those strategies. While they know the outcomes of incumbents, producers have at best informed but uncertain predictions of the consequences of novel strategies. More importantly, certain positions may appear to be more promising candidates for favorable outcomes than others. If this kind of information is relatively well-dispersed among actors, they will all perceive a similar set of potential innovation opportunities. The expected value of innovation in a market then indicates

the availability of specific opportunities with likely favorable outcomes.

Oil exploration offers an example of how the expected value of innovation can change in a market. Such firms face constant innovation decisions in the form of investments in potentially risky new projects, such as the development and exploration of fields (Mohn, 2008). Such projects can carry a number of risks such as over- or under-estimates of the oil they can extract, or possible political repercussions. More directly, their profitability is determined the price of oil—such projects have more room to fail as the prices increase. The global price of oil can determine whether the expected value of innovation is positive or negative, and hence whether firms take advantage of such opportunities or let them stand.

A Typology of Markets

The combination of the two key features of a market—the expected value of innovation and the responsiveness of audience thresholds—generates a typology of the locus of innovation or change in markets. The two features exert different kinds of effects on the market: the expected value of innovation exerts a global effect on the market, inducing all producers to either pursue or avoid innovation; a quality threshold exerts a local effect on sufficiently close producers, forcing them to tailor their innovation decision to the particular incentives induced by the threshold.

High- and low-quality producers are those sufficiently far from the quality threshold to ignore its incentives: they innovate in response to the global effect of innovation opportunities and their own taste for risk. High-quality producers will innovate if and only if available opportunities ensure that the expected value of innovation is positive. Low-quality producers will mirror high-quality producers to the extent that their own risk aversion allows it—they will pursue innovation only if it is favorable and they are sufficiently risk-seeking. In contrast to these unconstrained producers, middle-quality producers respond to the local effect of the threshold. They innovate in order to get above or stay above the threshold, whether or not the expected value of innovation is positive or negative.

The combination of these two effects generates four ideal patterns of market innovation. The four patterns determine whether innovation will emerge from high- and low- quality actors, from mid-quality actors alone, from all groups, or from none. Conversely, by observing a particular pattern of innovation, we may be able to infer the competitive dynamics of the setting.

	Fixed Threshold	Responsive Threshold
	“Mid-Status Conformity” <i>High (low) innovate, boundary stays</i>	“Red Queen” <i>All groups innovate</i>
Favorable In- novations	law firms, analysts Phillips and Zuckerman (2001) domestic partner benefits Briscoe and Safford (2008)	law schools Sauder and Espeland (2009) computer makers Barnett (2008)
	“Stasis” <i>No innovation</i>	“Mid-Status Innovation” <i>Boundary innovates, high (low) stay</i>
Unfavorable Innovations	farmers Simtowe et al. (2006) electronic health records Simon et al. (2007)	mutual funds Brown et al. (1996) Kaniel and Parham (2017) NASCAR Bothner et al. (2007)

Table 1: A Typology of Innovation in Markets

Tbl. 1 gives an overview of the typology. When innovation opportunities in the setting are favorable (i.e. the expected value of innovation is positive) but the audience imposes a fixed quality threshold, the market shows a pattern of innovation consistent with middle-status conformity (Phillips and Zuckerman, 2001): actors in the middle of the hierarchy, threatened by proximity to the threshold, avoid innovation, while those at the top seek it out. This quadrant makes an additional prediction for low-quality firms: if low-quality firms are relatively risk-seeking (i.e. if they can tolerate failure), they will mirror high-quality firms and innovate. The original settings supporting middle-status conformity may be understood through this lens: We expect high-quality law firms and high-quality investment analysts to attempt unusual strategies when they perceive market demand for such strategies, whereas uncertainty about the market’s response to such innovation keeps middle-quality firms and analysts from making the same attempt. In another context, Briscoe and Safford (2008) describe a situation in which a categorical constraint against corporate LGBT acceptance conflicted with growing market and activist demands for domestic partner benefits: adoption was led by high-quality companies known to resist activism, whose position was least threatened by this risky initiative. Edelman (1990) and Rao et al. (2005) describe similar innovative leadership by the largest or highest-quality producers in categorically constrained markets featuring

novel opportunities.

With favorable innovation but a responsive quality threshold, we expect a Red Queen pattern, with actors experiencing an overall increase in quality: actors at the top innovate to take advantage of opportunities, while actors in the middle innovate to stay ahead of the pack. Low-quality actors innovate to the extent that they can tolerate failure. Barnett (2008) describes this pattern most explicitly across several generations of computer makers, showing how incumbent manufacturers had to constantly compete against the efforts of newcomers developing increasingly smaller computers. Such a pattern may also map into the situation of law schools (Sauder and Espeland, 2009): a ranking system induces instability in the position of schools, especially those lower in the rankings at risk of falling below key thresholds in the ranking. This instability, along with the availability of favorable opportunities like investment in career services, induces persistent innovation in school strategies throughout the quality distribution. First-tier schools channel their resources into promising educational innovations at the same time as lower-tier schools resort to more questionable attempts to game the system.

If innovation becomes unfavorable, a responsive threshold will instead produce a pattern of “middle-status innovation”: actors in the middle innovate in order to stay ahead of the pack, but innovation is likely to make their product less desirable to audience members; actors at the top brace themselves and stay put. In highly competitive financial markets defined by a general lack of exploitable opportunities, this pattern may describe the tendency of mutual fund managers to fudge their portfolios to appear to outperform close peers (Brown et al., 1996; Kaniel and Parham, 2017). This pattern may also describe the behavior of NASCAR drivers, who pursue life-threatening risks in order to stay ahead in tournament rankings (Bothner et al., 2007). This pattern may, however, be dynamically unstable— any actor whose successful innovation puts them well above the threshold will stop there, so that the market bifurcates into a static top and a falling middle that is steadily driving itself out of business.

Finally, if an audience imposes a fixed threshold and innovation is unfavorable, the market enters stasis: an actor that has met the audience threshold has no further incentive to innovate; low quality actors have no innovation opportunities to pursue. Only actors that find themselves just below a threshold may pursue innovation if it offers a sufficiently high chance of clearing the threshold; such actors will continue to innovate until they find themselves outside of this region, either by

having met the threshold, or, as is more likely, by having fallen too low to have hope of further success. Since such markets feature little innovation, they draw little formal study. Nevertheless, this pattern may account for incidences of non-adoption of questionable practices (Rogers, 2003): Simtowe et al. (2006) describe the reluctance of Malawian farmers to adopt hybrid maize given that failure of this novel crop could threaten their survival; Simon et al. (2007) describe the low adoption rate of electronic health records in the face of legal and financial risks of adoption. In both situations, high quality actors avoid adopting questionable practices at the same time as constrained actors avoid adopting risky practices for fear of economic failure.

The dynamic evolution of these market patterns merits further discussion. The model discussed here treats innovation as a function of actor position in the quality distribution: the predictions of this model may be understood as predicting which regions in a quality distribution will innovate and which will stay put. Since innovation changes actors' positions in the quality distribution, innovative regions will tend to expel actors and lose population over time, while non-innovative regions will tend to capture innovating actors and grow over time. In fixed threshold environments, the boundaries of each region stay put over time; in responsive threshold environments, the boundaries of the threshold region may shift, capturing previously excluded actors or excluding actors that were previously affected by the threshold. For example, actors just above the threshold innovate defensively to the extent that actors below the threshold exist to threaten them. If the population of actors below the threshold falls, the number of mid-quality actors willing to innovate defensively will likewise fall.

These dynamics suggest additional predictions for the four market patterns. In general, the region just below the threshold will lose population over time, while the region just above the threshold will gain it. In responsive threshold markets, however, this change will be slower, as the threshold will chase after actors that have successfully innovated up. As such, innovation by actors near the boundary, both above and below, will persist longer in responsive threshold markets than in fixed threshold markets. An additional prediction applies to "mid-status innovation" markets in particular: as the region below the threshold depopulates, innovation in such markets will slow or stall, even near the threshold. Over time, such markets will bifurcate into those actors above threshold and those well below it and will come to resemble the static market pattern.

In addition, a market with favorable innovation opportunities may see those opportunities exhausted over time, so that mid-status conformity markets devolve into static markets, and Red

Queen markets devolve into mid-status innovation.

Because of these dynamics, static markets may be the general pattern in the world: Exhaustion of innovation opportunities will push markets into the static and mid-status innovation quadrants, and mid-status innovation markets will steadily devolve into static markets. Insofar as innovation research selects markets based on the dependent variable—innovation—markets that lack innovation never enter formal analysis. Most markets may, however, be dominated by a general stability. Utility companies may provide one example: they generally operate in closely regulated markets without price competition, and face fixed thresholds in the form of local laws and regulations. Absent clear opportunities to improve (as with recent pushes to renewable energy), such companies have remained largely static.

While this typology describes four distinct market patterns, these patterns are not permanent, stable characteristics of their markets. A market with favorable innovation opportunities may see those opportunities exhausted over time. Conversely, an unexpected discovery or invention may introduce novel opportunities to a market, producing middle-status conformity in static markets and Red Queen races from mid-status innovation. Likewise, innovation itself may reveal audience dissatisfaction with existing rules of the market, transforming fixed thresholds into responsive ones or vice versa. Such market disruption may involve turning a mid-status conformity market into a Red Queen market, or generate mid-status innovation among previously static producers (c.f. Christensen and Bower, 1996).

It is worth emphasizing one final point: this typology describes an idealized effect of market thresholds. Thresholds exert a local effect, affecting actors sufficiently close to the threshold. In practice, markets may have multiple thresholds, of varying strength and type. Predictions of innovation in such markets depend critically on the specific locations and strengths of the various thresholds, but the effect of thresholds is always local. When multiple thresholds are essentially coincident, we may expect the more consequential threshold to dominate; when multiple thresholds apply to different parts of the quality distribution, there will be multiple bands of innovation behavior, each band consistent with its respective threshold. Price competition itself may be understood as a responsive threshold operating in most markets, albeit with varying strength. Audience segments with distinct criteria represent another avenue for multiple thresholds (Hsu, 2006a; Kovács and Liu, 2016).

The next section formalizes these arguments in a model and illustrates them in simulation. The dynamics outlined above represent rational reactions not just to quality thresholds themselves but to other producers' own reactions to those thresholds. The model describes how an equilibrium pattern of behavior emerges from these reactions, and how this pattern depends on the particular configuration of producers in a market.

Quality Thresholds and Innovation Risk

The rules of the model are general. Multiple actors compete in a market. Actors face an audience controlling some desired resource. The audience yields this resource up in response to the actor's quality: each actor has a well-defined publicly known position on the real number line, and audiences reward higher-quality actors. A concrete instance of this general framework may be firms competing for money in a market—higher quality-firms receive greater profits than lower-quality firms.

A market contains some number N of actors. Each actor has some initial quality position q_i^0 in the market. Audiences observe the quality of all actors ($q^t = (q_0^t, \dots, q_N^t)$) and reward each with some level of resources $m(q)$ proportionate to their quality. Higher quality actors receive greater rewards.³ This reward function encapsulates any dynamics of a market, including the existence of quality thresholds—quality thresholds represent discontinuities in rewards received from the audience. In the discussion below, audience rewards will always be a 1:1 function of an actor's quality along with a fixed bonus if an actor meets the audience quality threshold.

Actors are risk averse. Actors value market rewards according to a common utility function $u(m)$. The utility function follows a typical form, with declining marginal utility ($u' > 0$) and risk aversion ($u'' < 0$). In addition, I require that actors exhibit declining absolute risk aversion (DARA)—each actor is more willing to take a given risk as she becomes wealthier, i.e. has a higher certain level of the desired resource. This assumption captures the intuition that as actors come to feel more secure in their position, they become more willing to accept losses that offer the chance to make greater gains (Pratt, 1964).⁴

³More formally, at a fixed level of quality for other actors, q_{-i}^t , m is non-decreasing in the quality of the focal actor q_i^t . Holding everyone else's position constant, raising an actor's quality will net her greater (or at least not lesser) rewards.

⁴I assume this for plausibility. The broad pattern of results holds for both constant absolute risk aversion (CARA) and a linear utility function. Under linear utility, high- and low-quality actors are more likely to innovate, low-quality actors will always take neutral (zero-mean) risks, and high-quality actors will always avoid them. Under CARA utility,

While actors' initial quality is outside of their control actors may try to change their quality by attempting innovations. Innovation is risky, and its outcome follows a random distribution, signifying a shock to their quality. If the actor does not innovate, her initial quality persists into the future, $q_i^1 = q_i^0$. If the actor innovates, her quality will instead be $q_i^1 = q_i^0 + X$ where X follows some continuous probability distribution with cumulative distribution function F_X .⁵ All actors decide whether or not to innovate simultaneously—they must decide to innovate without knowing whether or not their rivals will do so as well. This may stand in not only for explicitly simultaneous decision-making, but also for contexts in which the outcome of innovation comes long after the decision to innovate, so that actors may be in the process of innovating without their rivals knowing about it.

The two market features described above—the nature of quality thresholds and the favorability of innovation—map onto features of this model. Quality thresholds appear as discontinuities in the reward function m . Fixed thresholds are specific points in the quality distribution, while responsive thresholds are discontinuities whose location is a function of the qualities of all actors. The favorability of innovation is encoded in the characteristics of the innovation risk distribution X . Specifically, innovation is favorable when the expected value of innovation is positive, $E(X) > 0$, and unfavorable when the expected value of innovation is negative, $E(X) < 0$. The sign of this expectation can be read as the expected return to innovation: if an actor invests a dollar in R&D, will she expect at least a dollar back? Even if actors expect to lose money, the risk may be worthwhile if it offers them the chance to pass some quality threshold. Conversely, even if innovation is likely profitable, actors just above the threshold may prefer to avoid the risk of failure.

These model characteristics capture the general problem. The actor starts at some point in the quality distribution and faces competitors likewise distributed. Depending on her position, the positions of the competitors, and the market dynamics, the actor decides whether or not to innovate. She makes this choice if and only if the expected utility from innovation exceeds the expected utility from standing still:

low- and high-quality actors essentially mirror each other, innovating only when innovation has positive expected value. Risk aversion primarily affects the behavior of actors away from the threshold.

⁵Formally, the requirements on the distributions are that it has a continuous, strictly increasing CDF over its support. The logic of the model appears to only require that being a little farther is a little less likely—otherwise, the model may allow for diverging but not terribly informative edge cases.

$$E[u(m(q_i^0 + X, q_{-i}^1))] \geq E[u(m(q_i^0, q_{-i}^1))] \quad (1)$$

Because the reward function m potentially depends on the behavior of the other actors, the utility of standing still is a random variable—it depends on the realized positions of the competitors. Each actor makes this innovation decision simultaneously, and I look for equilibrium solutions wherein each actor best responds to the optimal behavior of her rivals.

In the discussion below, I will assume that innovation risks follow a normal distribution with some mean μ and variance σ^2 . Variance of the distribution plays a limited role in the model, because the scale of actor quality is arbitrary. In addition, asymmetry of innovation risks (e.g. most of the time innovation reduces quality a bit, but some of the time, it increases quality a lot) plays only a very limited role in these results.

Fixed Thresholds

Fixed quality thresholds have a straightforward effect on an actor's decision to innovate. Quality thresholds act as discontinuities in the reward function—they indicate points at which a small increase in quality causes a large increase in audience rewards. Fixed quality thresholds are discontinuities whose location on the quality line is fixed. Such thresholds may indicate membership in some prominent group or possession of some key certification. I model this using a reward function of the form:

$$m(q_i, q_{-i}) = \begin{cases} q_i, & q_i \leq c \\ q_i + C, & q_i > c \end{cases}$$

If an actor crosses the threshold at c , she receives a bonus reward of $C > 0$. Such a threshold may correspond, for example, to a situation in which the rating of a restaurant changes from three to four stars, or if the restaurant picks up a Michelin star. If the restaurant can use the event to command a price premium or drive media attention, it should be able to see higher profits. The position of this quality threshold does not depend on the qualities of other actors in the setting—it is fixed.

In this situation, the expected utility criterion eq. (1) reduces to a comparison between the

expected utility of the innovation risk and the certainty of the current position. For notational convenience, I define a function U to denote the difference in utility due to crossing the threshold:

$$U(q) = u(q + C) - u(q)$$

With this notation, the expected utility criterion for innovation simplifies to

$$\begin{aligned} E[u(q_i^0 + X)] + E[U(q_i^0 + X)|X > c - q_i^0] \cdot (1 - F_X(c - q_i^0)) &\geq \\ &\geq u(q_i^0) + (U(q_i^0)) \cdot \mathbf{1}\{q_i^0 > c\} \end{aligned}$$

where X follows a distribution with cdf F_X .

On the left-hand side, this criterion separates into the baseline of expected utility in the absence of a threshold reward along with the expected excess utility associated with meeting the threshold. The right-hand side likewise separates into the baseline level of utility along with some excess if the actor is over the threshold already. On the left, uncertainty due to taking the risk creates a continuous increase in expected utility as an actor's quality increases; on the right, the threshold induces a sharp discontinuity at the cutoff. As such, as an increase in an actor's quality causes her to cross the cutoff, she will experience a sharp negative shock to the expected utility of innovation—crossing the threshold causes a sudden jump in excess utility $U(q_i^0)$ on the right hand side.

This sudden change in rewards causes actors just above the cutoff to value innovation much less than they would if they were just below the cutoff. At the same time, if the distribution of the innovation risk is well-behaved, innovating actors experience only a gradual change in their expected utility. If the excess utility of crossing the threshold is higher than the loss due to risk aversion, actors just below the cutoff will certainly choose to innovate. On the other side, if the excess utility due to crossing is larger than the expected increase in quality due to innovation, actors just above the threshold will certainly avoid innovation—if the risk pays off, they will not gain as much as they would lose if the risk failed and pushed them below the threshold.

Fig. 1 shows the desirability of innovation in the presence of a fixed threshold at a quality of 0.5, as the favorability of innovation varies. Those below the threshold will innovate even if innovation

is unfavorable (i.e. the risk distribution has negative mean)—the chance of crossing the threshold is sufficiently high that it makes up for the more likely but slight drop in quality. Those above the threshold will avoid innovation even if it is favorable (i.e. the risk has positive mean)—the risk of dropping below the threshold is sufficiently high that it outweighs the more probable increase in quality.




Figure 1: Expected utility of innovation with a fixed threshold

The figure also indicates that the behavior of actors well above and well below the threshold is primarily determined by the favorability of innovation. If innovation is favorable, actors find innovation desirable both just below the cutoff and well above it. If innovation is unfavorable, only actors below the threshold will innovate. As such, favorable innovation generates a tripartite pattern innovation: actors well above and below the threshold innovate, while those just above the threshold in the middle of the quality distribution avoid innovation.

In line with the typology presented earlier, it is possible to map existing theories of organizational behavior and innovation into this framework. Fixed quality thresholds with favorable innovation opportunities generate a pattern of behavior akin to the predictions of middle-status conformity. The presence of an audience threshold constrains the behavior of actors close to that threshold. Those just above the threshold must avoid innovation for fear of falling below it. Those well above the threshold, on the other hand, experience little constraint from the audience and are free to

pursue innovations. The behavior of low-quality actors diverges: Actors just below the threshold innovate, but they do so in order to try to fully participate in the market, not because their behavior is unconstrained by the audience; actors well below the threshold avoid innovation because of their extreme risk aversion. Over time, however, innovating actors will either succeed, jumping above the threshold, or fail, falling too far below it. We should expect such settings—with a fixed threshold and favorable innovation opportunities—to stabilize into a set of high-quality actors pulling away from the pack, a stable cluster of middle-tier actors, and well below those, a stable cluster of low-quality actors.

Fig. 1 also indicates the possibility of a pattern of actor stasis altogether distinct from the predictions of middle-status conformity. If innovation opportunities in a setting are unfavorable (i.e. quality-reducing in expectation), only actors just below the threshold will pursue them. As above, such low-quality innovators are unlikely to persist over time, so that the setting steadily evolves to a pattern in which no actors innovate. The market will feature high-quality actors well above the threshold, a glut of middle-quality actors just above the threshold, and a cluster of low-quality actors some distance below them. No part of the market will innovate.

The dividing line between the pattern of middle-status conformity and a pattern of stasis is thin. Insofar as innovation opportunities change over time, innovation may shift from favorable to unfavorable. This transition will be apparent primarily in the behavior of high-quality actors, who may suddenly begin to take unusual actions and explore novel strategies. Conversely, insofar as innovation exhausts available opportunities, the behavior of high-quality actors will stabilize. The long-term stability of a pattern of middle-status conformity or of stasis cannot be assured—circumstances in a setting may cause transitions between the two patterns at any time.

Responsive Thresholds

Responsive thresholds generate a different pattern of behavior, shifting as the quality of actors shifts. A market may feature substantial rewards in response to entering some key group or meeting some certification, but the criterion for such certification may vary with the general quality of actors in the setting. The threshold for being a top-10 college, for example, depends on the behavior the current top 10.

I model this using a reward function of the form:

$$m(q_i, q_{-i}) = \begin{cases} q_i, & q_i \leq c(q) \\ q_i + C, & q_i > c(q) \end{cases}$$

If an actor crosses the threshold at c , she receives a bonus reward of $C > 0$. The threshold c , however, may be a function of q , the positions of all the actors—it may, for instance, be some percentile or the average of the quality distribution.⁶ In the results below, I take c to be the median of the quality distribution. Once again, I define the convenience function $U(q) = u(q + C) - u(q)$ to denote the difference in utility due to crossing the threshold. In this situation, the utility criterion features uncertainty on both sides of the inequality:

$$\begin{aligned} E_{X_i}[u(q_i^0 + X_i)] + E_{X_i}[U(q_i^0 + X_i) \cdot P(q_i^0 + X_i > c(q^1))] &\geq \\ &\geq u(q_i^0) + U(q_i^0) \cdot P(q_i^0 > c(q^1)) \end{aligned}$$

where X follows a distribution with cdf F_X .

On the left hand side, the criterion is substantially similar to that described in the previous section: An actor's expected value to innovation can be separated into her expected utility associated with a change in quality resulting from the innovation alone, along with an excess reward if she crosses a threshold. The location of the threshold, however, can now shift depending on the choices and outcomes of the other actors in the setting. On the right hand side, the actor receives the utility from her present position along with an excess reward if the behavior of the other players causes the threshold to remain or drop below the focal actor's current position.

Because the threshold is responsive, innovation by rivals can cause it to shift upwards and leave previously safe actors high and dry. Actors above the threshold cannot ensure a particular stream of rewards simply by staying put—the risk of the threshold moving may induce actors above the threshold to defensively innovate in order to hedge against this possibility.

Illustrating optimal behavior under responsive thresholds is more complicated than illustrating behavior under fixed thresholds. In particular, it involves identifying equilibrium decisions—actors'

⁶The only formal criterion on $c(q)$ should be that it falls somewhere between the minimum and maximum of actors qualities in either the initial or final period. This ensures that at least *some* actor can meet the threshold and that some actor is below it. Without this requirement, the threshold cannot affect incentives to take risk.

optimal decisions given the decisions of their peers. I illustrate actor behavior in such markets in three steps: First, I derive an analytic result in a limited two-actor game to illustrate the formal intuition. Second, I simulate the innovation process in settings with larger numbers of actors to show an emergent pattern of defensive innovation. Third, I show the results of this innovation process over time, showing how markets with fixed and responsive thresholds evolve.

Two Actors

A two-actor game illustrates the logic of equilibrium decision-making that characterizes behavior under a responsive threshold.⁷ The subsequent section extends the same logic to settings with more actors.

The rules of competition are simple: I consider the quality threshold $c(q) = \min(q)$ —the better actor gets the threshold reward and the worse one does not. I consider the specific utility function $u(m) = a \cdot m - \exp(-b \cdot m)$ —this function satisfies the risk aversion criteria described earlier. Finally, I consider an innovation whose outcome follows a normal distribution with mean μ and variance σ^2 . Deriving the equilibrium depends on two questions: how does an actor best respond when the opponent is choosing to innovate, and how does the actor best respond when the opponent chooses not to?

Figs. 2 and 3 present plots of this criterion, and the resulting equilibria, for two situations: mean negative risk (fig. 2), and mean positive risk (fig. 3). On all figures, the horizontal axis represents the current position of the focal player (ego) and the y-axis represents the current position of the opponent (alter). Points below the 45-degree line indicate positions in which ego is higher quality than alter. On all figures, blue (positive) regions indicate positions at which ego prefers to innovate, and red regions (negative) indicate positions where ego prefers to stay in place.

The equilibrium calculation is simplest with unfavorable innovation risks (fig. 2). First, consider point A—with ego quality at 0.4 and alter quality at 0.75, ego has much lower quality than alter.

Figs. 2a and 2b indicate that the value of innovation here is negative: Ego prefers not to innovate

⁷The games I discuss are highly simplified two period models. There is more extensive modeling of such ‘races’ in economic literature on innovation—in particular, Cabral (2003) and Hörner (2004) present multi-period models of two-player races. Their results are substantively similar to the ones discussed here, suggesting that intuitions developed by these simpler models are accurate for more dynamic settings. Those models are, however, limited to two-player races. In addition, the races they model do not discuss thresholds explicitly but rely on alternate ways to induce concave utility for players.





Figure 2: Innovation behavior with unfavorable innovation opportunities ($\mu = -0.01$).

because the probability of being lucky enough to beat the opponent is too low. The symmetric point A' shows alter's behavior in this situation. Figs. 2a and 2b once again indicate that the expected value of innovation is negative—alter also refuses to innovate, given that doing so will probably reduce his quality. As such, the equilibrium outcome here is that both ego and alter refuse to innovate: fig. 2c indicates that the points A and A' have a “no innovation” equilibrium.

If alter is closer in quality, so that ego faces the situation at B , there is a different set of equilibria: As figs. 2a and 2b indicate, ego now prefers to innovate whether alter is doing so or not. Though innovation is likely to reduce ego's quality, the chance of ego being lucky enough (and alter being unlucky enough) for ego to beat alter is high enough to take the risk. The symmetric point B' shows that alter would prefer not to innovate regardless of what ego is doing. As such, the only equilibria is for the lower quality actor to innovate and the higher quality actor to avoid doing so, as fig. 2c indicates.

Finally, if ego and alter are very close in quality, say at point C , we observe the same equilibrium as above, but under a slightly different logic. As fig. 2a shows, if alter stays put, ego prefers to innovate, because the chance of beating alter at such a small distance is very high. But as fig. 2b shows, if alter *is* innovating, ego prefers not to: because innovation is quality-reducing, alter is likely to fall below ego with no action on ego's part. Ego prefers to stay put and watch alter trip up. But as the symmetric point C' shows, alter would prefer not to innovate in any situation: at best, his quality will probably drop, and at worst, it'll drop below ego, causing a massive loss in rewards. Fig. 2c shows that in equilibrium, ego innovates and alter does not.

The general takeaway in an unfavorable innovation environment is that either nobody innovates, or if they are close enough in quality, the lower quality actor may attempt to.

This situation does not persist in a favorable innovation environment (Fig. 3). The logic behind these figures is substantially similar, with three differences worth discussing. First, consider the situation at A' , where the two actors are close in quality. In this case, alter prefers to innovate whether or not ego is doing so: if ego is not, innovation generates the possibility of surpassing them; if ego is innovating, their expected quality is higher in the next period, but not high enough that ego cannot beat them with some luck. This creates a threat for ego. Ego (at A) will not innovate if alter stays put, for fear of dropping below the cutoff; but if alter is already innovating, ego is in danger of being surpassed, and prefers to innovate as well. As such, there is an equilibrium in which





Figure 3: Innovation behavior with favorable innovation opportunities ($\mu = 0.01$).

alter certainly innovates, and ego also innovates to protect herself from alter.

Situation *B* reflects the logic of a fixed threshold. Both ego and alter are too far from the threshold for it to affect their decisions. Because the innovation is favorable (quality-increasing in expectation), the high-quality ego is willing to attempt it, whereas the low-quality alter will refuse to do so because of their high risk aversion. This produces equilibria in which very high quality actors innovate and very low quality actors refuse to.

Point *C* represents a more complicated interaction between alter and ego. Though alter and ego are far apart, they are close enough that simultaneous innovation may cause alter to surpass ego. This generates a cycle: (1) If alter stays put (fig. 3a, *C*), ego will innovate; (2) but if ego innovates, alter will innovate as well, hoping to she will get particularly lucky and ego will be particularly unlucky (fig. 3b, *C'*). (3) Fearing the same outcome, if alter innovates, ego will stay put (fig. 3b, *C*); (4) but then alter becomes less likely to beat ego, so alter would also stay put (fig. 3a, *C'*). Because there is a cycle and not an equilibrium in pure strategies, the equilibrium solution will be to have both actors decide at random whether to take the action or not. As fig. 3c indicates, this generates an equilibrium in mixed strategies at *C*.

These models reveal three basic patterns of innovation. Actors sufficiently high above the threshold will innovate when innovation is favorable. Actors just below the threshold will almost always innovate. Most importantly, actors on both sides of the threshold will innovate only when the expected value of innovation is positive. Only then does a low-quality actors present enough threat to a higher-quality actor that the high-quality actor must react defensively.

This requirement for defensive innovation breaks down when more than two actors compete. Intuitively, an actor above the threshold competes not against a single rival, but against the luckiest of her rivals. With enough competition, ego must rely on one of her rivals being successful enough to raise the quality threshold. In that situation, an actor above the threshold must innovate defensively both in situations of favorable and unfavorable innovation.

Multiple Actors

Extending the analysis above to a general n -player game is complicated by several factors. First, the distribution of many responsive thresholds, such as median or quantile functions, becomes intractable (Bapat and Beg, 1989). Second, visualizing the decision-making of n distinct actors

simultaneously is difficult. Instead, I attempt to demonstrate the same point by simulating actors' decision process.

I generalize the game described above to a multiple-actor setting, illustrating both the situation of fixed and responsive thresholds. Once again, all actors have the option of making some innovation that shifts their quality position according to a normal distribution with some mean and standard deviation.⁸ The outcome is independent across actors. Actors receive a reward either for beating the median (a responsive threshold) or a fixed threshold at 0.5. The actors' initial position is drawn from a uniform distribution on $[0, 1]$. I present results for environments of ten actors—results are robust to the number of actors, so long as the number is greater than two.

Though I cannot calculate exact equilibrium solutions in this setting, I calculate actors' optimal behavior using an iterated simulation approach. Given their initial position, each actor begins by deciding not to innovate. Then each actor makes a calculation: Holding fixed what all others have decided to do, each actor simulates two possible worlds—one in which she innovates and one in which she does not, calculating the difference in expected value between them. In the next round of iteration, other actors then assume that the focal actor adopted the choice with the higher expected value. All actors make these calculations simultaneously,⁹ generating a potentially different set of innovation decisions for the next round of iteration. The actors then iterate again, taking the new set of equilibrium actions as fixed.

This procedure is guaranteed to terminate either in a fixed vector of equilibrium actions or in a cycle of best responses.¹⁰ Termination in a cycle indicates that there exists an equilibrium in mixed strategies. Once iteration has terminated either way, I record the final equilibrium action of each actor, including whether the player is mixing.

Fig. 4 present aggregated results of these simulations for the four quadrants of the innovation

⁸These simulations take risk to be distributed with standard deviation $\sigma = 0.1$. The mean is $\mu = 0.01$ in the favorable risk situations, and $\mu = -0.01$ in the unfavorable risk situations.

⁹Unfortunately, it does not appear possible to show that modeling actor decisions as simultaneous is innocent. In general, iterating on best responses does not guarantee identification of a unique equilibrium—it is possible that if actors were to make their decisions in some other order, they would converge to some other equilibrium state or cycle. Any such alternative equilibria should, however, be qualitatively similar to those observed here.

¹⁰Consider the vector of all actors' innovation choices. The set of all possible innovation vectors is finite. Given any vector of innovation choices, any given actor's best response to it will always be the same—no matter how many times the simulation gets to a particular innovation vector, the best reply in the next iteration will not change. This procedure of starting at one vector and jumping to another can visit only a finite number of unique vectors, because the whole set of these vectors is finite. At some point, it must loop back to some previously seen vector—this loop may involve just a single vector, in which case the simulation has reached equilibrium.

typology. Fig. 4a shows individual decisions of actors close to the threshold. Each point in the figure represents some actor in some run of the simulation. The horizontal axis shows the actor’s distance from the threshold, with actors on the right above the threshold, and those to the left below it. The vertical axis shows the expected value of innovation for that actor—actors whose expected value is above zero innovate. As the figure shows, fixed thresholds induce a sharp divide between innovators and abstainers at the threshold. But with a responsive threshold, some actors above the threshold choose to innovate, whether or not they expect innovation *per se* to raise their quality.


Fig. 4b zooms out from the threshold and shows the aggregated results of these individual decisions—it shows what proportion of actors a given distance away from the threshold across all runs chooses to innovate. The figure confirms that high-quality actors, those well above the quality threshold, innovate only when they expect innovation to be quality-improving. At the same time, the figure confirms that responsive quality thresholds can induce the middle-quality actors just above the cutoff to innovate as well.

The multiple actor version of the game shows directly that a responsive reward threshold can induce actors to innovate on either side of the threshold. This occurs even in settings where innovation is unfavorable, that is to say, even in settings where actors above the threshold expect it to hurt them. With responsive thresholds, middle-quality actors cannot reliably remain above the quality threshold. They innovate in order to hedge against the innovative potential of those below them.


Dynamic Results

To more fully establish the results of the model, I simulate the multiple actor game described in the previous section across multiple periods. This dynamic model suggests that, far from describing minor differences in actor behavior, fixed and responsive thresholds generate qualitatively different market outcomes.

The logic of the multi-period simulation reflects the procedure described above. Actors start initially distributed in the unit interval. In each period, actors calculate their optimal behavior given the market they’re in. Actors that chose to innovate then change position according to the innovation risk distribution, and responsive thresholds shift to reflect the new distribution. This



(a) Expected utility of innovation in multi-actor environments



(b) Proportion of actors innovating in multi-actor environments

Figure 4: Innovation dynamics in multi-actor environments

process then repeats. I simulate each market across 50 periods.¹¹


I present the results of this simulation in fig. 5. Fig. 5a presents the overall balance of innovation between markets with fixed and responsive thresholds over time. Solid lines show the base-10 logarithm of total innovation volume in each type of market. The dashed line shows the difference between the two. When it is positive, markets with responsive thresholds are more innovative than markets with responsive thresholds. Shading represents the 95% confidence interval for the difference.

Fig. 5a suggests that markets with responsive thresholds are more innovative than markets with fixed thresholds, especially over time. Because the difference is a logarithm, the difference between the two market types is substantial: responsive threshold markets feature about three times as much innovation as fixed thresholds markets. Fig. 5a does, however, suggest that this difference disappears over time in markets with unfavorable innovation: innovation in both fixed and responsive thresholds markets disappears after 50 periods, though responsive threshold markets remain generally more innovative over this interval.


Figs. 5b and 5c explore which parts of responsive threshold markets produce this increased rate of innovation. I separate actors into four segments: High-quality actors are those higher than two standard deviations of the risk distribution (0.2 units) above the threshold; low-quality actors are those lower than two standard deviations below it; high-middle and low-middle actors are actors that are within two standard deviations and fall either above or below the threshold, respectively. This two standard deviation bandwidth roughly captures those actors affected by threshold dynamics; bandwidths of one or three standard deviations show essentially identical results.

Fig. 5b shows the proportion of total innovation that each segment is responsible for over time. The dashed line shows what fraction of its peak innovative output the market produces in each period. The upper left quadrant, for instance, shows that in fixed threshold markets with highly favorable innovation opportunities ($\mu = 0.01$), overall innovative output decreases over time, but that high quality producers are responsible for an increasing proportion of it; low-quality producers generate a substantial stable fraction of overall innovative output, and low-middle producers (i.e. those just


¹¹I simulate markets with a fixed threshold at 0.5 or a responsive threshold at the median. In the simulations presented here, I assume an exponential utility function ($u(m) = m - \exp(-m)$), and a normal risk distribution with $\sigma = 0.1$. I vary the expected value of innovation (μ) across 0.01, 0.003, 0, -0.003, and -0.01 (10%, 3%, 0%, -3%, and -10% of σ , respectively).



(a) Balance of innovation decisions between fixed and responsive thresholds



(b) Proportion of innovation by quality segment



(c) Proportion of population by quality segment

Figure 5: Innovation dynamics over time by market type

below the threshold) generate a small and falling proportion of total innovation in such markets.

Comparing the left and right halves of fig. 5b shows that mid-quality producers play a much larger role in responsive threshold markets than they do in fixed threshold markets. They are responsible for nearly all innovative output in markets with unfavorable innovation opportunities, and for a substantial stable proportion of innovative output in markets with favorable opportunities. More importantly, high-middle quality producers, those just above the threshold, play a small but persistent role in producing innovation in responsive threshold markets, whereas they play no innovative role in fixed threshold markets. fig. 5b confirms that responsive thresholds generate a persistent pattern of mid-quality innovation while fixed thresholds generate a pattern consistent with middle-status conformity.

Fig. 5c expands this story by breaking down the population of each market across the four segments. The bottom of the figure shows that the low-middle quality producers, who play the major innovative role in markets with unfavorable innovation opportunities, quickly disappear. But while low-middle producers quickly split into low and high-middle producers in fixed threshold markets, successful innovation by low-middle producers in responsive threshold markets tends to pull high-middle producers below the threshold, sustaining a longer overall period of innovation. The top of fig. 5c tells a different story: in fixed threshold markets, the population of high-middle quality producers balloons over time, as they become trapped in a situation that forces them to avoid innovation. In contrast, the population of mid-quality producers forms a relatively stable fraction of the population of responsive threshold markets, as the responsive threshold dynamic forces such producers to repeatedly churn across the threshold. Finally, a comparison of figures 5b and 5c shows the disproportionate innovative role played by low-middle quality producers, who make not only a large proportion of overall innovation decisions, but who also induce a degree of innovation in high-middle quality producers in responsive threshold markets.

Before concluding, it is worth discussing the robustness of this model. The model presented above relies on a number of assumptions about actor utility, the innovation decision process, the exact specification of thresholds, and the nature of the innovative risk distribution. The specific decisions were made for analytical or explanatory convenience, but the major patterns predicted by the typology persist in almost all alternative specifications of these details. The online appendix steps through a number of alternatives in greater detail. There are few notable differences as a

result of these checks: alternative utility functions determine whether high- and low-quality actors are too risk averse to pursue favorable innovation opportunities or not; alternative specifications for the mean and variance of the innovation distribution show that results persist so long as the variance is large enough to capture some number of actors near the threshold—too small, and not enough actors are in jumping distance of the threshold for it to affect them. Otherwise, innovation patterns are robust to changes to the actor decision process, the shape of the risk distribution (i.e. skewness or bimodality, c.f. Eliashberg et al., 2006; Rossman and Schilke, 2014), or to more imprecise thresholds. The difference between fixed and responsive quality thresholds represents a qualitative difference in market dynamics.

Discussion

One question drives this research: where does innovation come from? I focus on markets with a particular set of features. Producers in the market face some audience—consumers and stakeholders—that places demands on the performance of the producers. Producers within the market fall along some quality distribution, so that those higher in the distribution prove more appealing to the audience, or at least find it easier to carry out their work. Higher quality producers earn greater reward from the audience. Innovation is the process of trying to move up in the world—it is a producer’s attempt to raise her quality by changing the offer or the process of producing it. Such innovation is uncertain. The question I ask is why only some producers choose to innovate? Why do some markets see innovation coming only from the middle-quality producers, and others only from the extremes? Why do some markets see no innovation?

The answer I propose highlights two features of such markets that drive differences in innovation: (1) the responsiveness of the quality thresholds that the audience sets for producers; and (2) the favorability of innovation opportunities available to producers. Audiences create thresholds of quality that producers must meet in order to win audience attention and rewards. These thresholds may be fixed—they may take the form of formal certifications, licenses, degrees, ratings, or well-established schema for behavior. These thresholds may be responsive—they may take the form of top-10 lists, rankings, best-of-breed awards, or even emergent processes such as appearance in online searches or recommendations. Through such thresholds, markets sort producers into the best and the rest. But

the responsiveness of these thresholds determines whether producers just past the cutoff are safe or whether they remain vulnerable as their rivals innovate.

At the same time, producers' attempts to innovate are inherently risky but predictably so. Producers have limited knowledge of audience desires. They have limited knowledge of their own capability to carry out an innovation (Hannan et al., 2003). The process of implementation may result in unexpected employee reactions or unexpected synergies with their existing offer. Law firms, for instance, may not know how adopting family law would affect their desirability to future recruits. Law schools may not know whether their attempts to cheat the rankings will be caught. Yet producers can make an educated guess as to whether the change will likely improve or reduce their perceived quality. This expected value of innovation interacts with the quality thresholds imposed by audiences: quality-improving innovations are good, but for the threat of pulling a producer below the quality threshold in case of failure; quality-reducing innovations are bad but for the opportunity of pushing a producer above the threshold in case of success.

These two dimensions combine to form four patterns of innovation behavior. Favorable innovation risks and a responsive threshold generate a Red Queen dynamic, where producers at the top of a market innovate to exploit the desirable opportunities available to them and producers near the threshold innovate to stay ahead of the pack. With favorable risks and a fixed threshold, the setting produces a middle-status conformity pattern—the highest quality producers innovate to take advantage of available opportunities, while those below them hunker down to avoid displeasing the audience and falling below the threshold. Unfavorable risks and a responsive threshold, in contrast, generate middle-status innovation—only producers in the middle innovate, hoping to get lucky and beat the crowd. With unfavorable risks and a fixed threshold, the world evolves towards stasis, as producers just below the threshold winnow themselves out trying to reach it.

These patterns represent ideal market types, and, in practice, markets may feature multiple simultaneous thresholds. Price competition in particular appears to function as a responsive threshold in all markets insofar as producers must strive to meet or beat a prevailing market price. Though the model above does not delve into this issue, it does offer a natural guide for thinking through multiple threshold markets. Quality thresholds exert a local effect on nearby actors in the quality distribution and the size of this effect is commensurate with the bonus afforded by the threshold. When multiple thresholds operate in a market, they will exert multiple local effects,

creating additional bands of innovation and stability in a market.

This framework is useful for both predictive and explanatory purposes. Researchers armed with a priori knowledge of audience standards and industry opportunities can predict the pattern of innovation behavior by producers. On the other hand, a particular pattern of innovation behavior sheds light on the dynamics of its setting: Innovation by high- and low-quality producers reveals the quality of opportunity available in a setting. Innovation by middle-quality producers reveals whether audience standards are stable or responsive.

The ability of this framework to guide inference about audience standards may provide a powerful tool to advance empirical work on market categories. Empirical support for category theory has relied on direct measurement of the audience reaction to producer offers—researchers have looked at ratings, reviews, and sales. Phillips and Zuckerman (2001) suggested that producer behavior could be used to infer the categorical dynamics within a market. The typology described here expands that ability by describing how additional market patterns emerge from underlying audience standards and market circumstances. Stability within law firms points to the stability of audience thresholds. Persistent churn in the smartphone market points to the flexibility of the audience standard in that setting. Observation of producer behavior is an observation of audience standards for producers.

More broadly, the distinction between fixed and responsive thresholds works to emphasize the key role that categorizations play in markets. While price dynamics play an important role in competitive behavior, markets often feature substantial information asymmetries or informational shortcomings that cannot be adequately addressed by the price system (Akerlof, 1970; Spence, 1973; Podolny, 2008). Categorizations both within and between markets relieve audiences of these informational burdens. By highlighting the difference between fixed and responsive threshold dynamics, this research emphasizes the breadth of market consequences that emerge from this solution to information problems within markets.

In its reliance on the bounded rationality of the audience, this framework presents a mirror to performance feedback theory. The limited information available to the audience forces them to create means for categorizing producers, turning a complex marketplace into a binary threshold. Seeing thresholds in the market, producers react as they would to their own private aspirations, but for two differences. First, the shared threshold imposes a shared aspiration level on all producers;

second, competitive threat under responsive thresholds induces producers to innovate defensively even when they find themselves above the threshold.

Following another vein, it is critical to point out that this analysis describes only a ‘slice of life’ in the course of a market’s development. Two factors motivate innovation: first, the inherent desirability of innovation due to the presence of favorable opportunities; and second, pressure to meet and stay above an audience quality threshold. It is possible to conceive of markets in which both features persist over time—modern technology markets, for instance, appear to be infinitely capable of producing opportunities, and the food industry appears to be persistently flooded with short-lived entrepreneurs. In general, however, both motivations are likely to be temporary: producers can exhaust the opportunities available to them, and low-quality innovators will winnow themselves out over time as they either become established or leave the setting.

Of the four quadrants of innovation described in this article, the most common may be stasis—the total absence of innovation throughout a market. While stasis forms only one quadrant of the typology, the presence of innovation in the others depends critically on the continuing presence of innovation opportunities or on innovative pressure from lower-quality competitors. Once these disappear, all markets will tend toward stasis.

Yet the nature of quality thresholds and the range of innovation opportunities are impermanent qualities of a market. If either feature changes, the pattern of innovation in the market will change. The taxi industry appeared to be essentially static for decades until the sudden appearance of Uber, Lyft and similar startups introduced novel ways to hail rides and skirt regulations. These novel entrants exploited emerging technological opportunities and a belief in the underlying willingness of taxi-rides to accept alternative approaches to the problem: the apparently static market hid a responsive audience threshold that ride-sharing startups exploited. In practice, most markets may contain this kind of latent potential, as even the most categorically-bound audiences may harbor hope for some better alternative.

Research on disruptive innovation (Christensen and Bower, 1996), for example, appears to describe this very dynamic, showing that risky innovation by low-quality, peripheral players can reliably destroy markets by suddenly changing the audience’s threshold for quality. Insofar as incumbents avoid such innovations, disruptive innovations may more often than not be quality-reducing innovations, which researchers observe as disruptive only when they succeed. As such, there

may be a vast body of failed disruptions that merit study. Such study may reveal that attempts to disrupt fail more reliably than they succeed, so that incumbents' best strategy is to mount no defense against them, and simply hope that innovators do not get lucky. More importantly, innovation may disrupt less through the success of disruptors themselves and more by spurring defensive innovation by incumbents—Microsoft Bing and Google Plus may have played larger roles in Google Search and Facebook than commonly believed; Napster may have done more to destroy existing record companies than to ensure its own survival.

These concerns raise the further possibility that research on innovation may have been hampered by a form of selective sampling of empirical settings. The innovative character of a market or industry stems not from the observed consequences of innovation within it, but from the structures that cause producers to seek or avoid innovation. Markets may be innovative even if innovation leads to no major changes in observed products or the success of the players within them, so long as the structure of the market induces some set of producers to pursue innovation or puts them in a position where a small change in circumstances will push them to pursue it.

This discussion leads naturally to limitations of the current model that could be profitably addressed in subsequent work. One straightforward improvement would look at the dynamic effect of fixed and responsive thresholds. While I have speculated on what these effects might be, a full analysis of the evolution of a population of innovating producers in the presence of quality thresholds may reveal unexpected and novel market dynamics. A second would be to explore the effect of variable risk portfolios, showing how producers take on different magnitudes of risk depending on their position relative to the threshold.

A more important extension would involve more direct modeling of the feature space, the innovation problem in this space, and the audience response to offers. Much of the logic underlying this article implicitly assumes that innovation takes place in some abstract feature space which hosts both producer offers and audience standards for producers, and which translates unproblematically into a linear quality spectrum. Producers are able to stake out novel positions within this space, and audiences may likewise be able to learn what their own preferences over positions in this space may be. Modeling the former process would formally pin down a way to interpret the quality of risks in a setting. Modeling the latter process would pin down the nature and responsiveness of quality thresholds. The model of innovation risks and quality thresholds described here is relatively

abstract—it may be made less so without losing its essential insights. A deeper specification of the feature space and audience reaction would also help build a stronger bridge to category theory and aid the flow of insight between work on social classification and work on market strategy.

Finally, this work is naturally limited to settings in which the notion of quality differentiation makes sense, and in which the behavior of audiences functions to create thresholds in this quality distribution. The formal model described here is highly stylized—most markets likely feature multiple simultaneous mechanisms that act to create multiple and multiple types of thresholds constraining producer behavior. Alternatively, some markets may have no clear dividing line between high- and low-quality producers. At the same time, innovation in many markets may primarily involve efforts to horizontally differentiate to find a niche audience (c.f. the differentiation of microbrewed beers, Carroll and Swaminathan, 2000). The model described here hopes to offer some insight into such situations, insofar as it guides intuition on how multiple constraints may combine together, or helps researchers look for signs of audience constraint. The growing literature on market categories (Zuckerman, 1999; Zuckerman et al., 2003; Hsu, 2006b) should suggest that audience-induced quality thresholds may be fairly common throughout the economy, and older organizational literature may suggest that audiences may impose idiosyncratic standards on even idiosyncratic producers (Hannan and Freeman, 1984). Nevertheless, the framework described here leaves room for alternative drivers of innovation.

Conclusions

This article proposes a typology of innovation behavior in vertically differentiated markets. Producers make offers to an audience that rewards them for the quality of their efforts. This quality measure sorts the producers in a market, with those of the highest quality earning the largest rewards. Innovation is risky, as it involves making changes with unknown costs and payoffs. Audiences impose quality thresholds on producers in the setting, with those crossing the threshold winning outsize rewards—fear of falling below these thresholds generates fear of innovation. The typology derives from two features of this market setting: the innovation opportunities producers consider can be on average quality-improving or quality reducing, favorable or unfavorable; and the quality thresholds that audiences impose may be fixed or may be responsive to producers offers. Different combinations

of these two factors activate innovation at different sites in the quality hierarchy: we may observe innovation from the best producers, from mid-quality producers near the market threshold, from the best and the middle, or from nowhere in the market.

This typology makes several contributions to past work. First, it expands the logic behind the middle-status conformity argument to a broader scope, suggesting that the pattern of a static middle and dynamism at the top and bottom of a hierarchy can describe hierarchies not just of status but of arbitrary quality standards, and that the behavior this logic predicts can apply to general innovation behavior. The mid-status conformity pattern, with stability in the middle and innovation at the extremes, may be a feature of many different market settings.

Second, this typology naturally extends to a much broader set of markets, predicts a broader set of innovation patterns, and hints at how thin the boundaries between different patterns may be. The specific pattern of innovation in a market depends on the quality of available risks and the responsiveness of quality thresholds to innovation. But the quality of innovation risks changes over time, and audience standards may become or may reveal themselves to be more or less stable over time. A market with a stable middle can easily transform into an unstable crucible. A market with an innovative top can easily transition into stasis.

The most important contribution of this framework may be in its ability to build a bridge between category theory and the work on strategic exploration in markets. The framework draws on insights from both fields: On the one hand, audiences set quality thresholds in order to standardize the offers producers bring to market. On the other hand, neither audiences nor producers know what the real landscape of possible things may be. Producers respond to the explicit standard set them by their audience, but they also respond to the potential of this veiled landscape. Audiences attempt to regulate and schematize an otherwise chaotic market, but leave some room to learn what's out there. Incorporating strategic insights provides a way to deepen our understanding of how audiences perceive and respond to markets. Incorporating audience explains how social expectations shape the incentives of markets and the producers within them. A bridge between these two fields improves both.

References

- Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84(3), 488.
- Bapat, R. B. and M. I. Beg (1989). Order Statistics for Nonidentically Distributed Variables and Permanents. *Sankhya: The Indian Journal of Statistics* 51(1), 79–93.
- Barnett, W. P. (2008). *The Red Queen*. Princeton: Princeton University Press.
- Baum, J. A. C., T. J. Rowley, A. V. Shipilov, and Y.-T. Chuang (2005). Dancing with Strangers: Aspiration Performance and the Search for Underwriting Syndicate Partners. *Administrative Science Quarterly* 50(4), 536–575.
- Bothner, M. S., J.-H. Kang, and T. E. Stuart (2007). Competitive Crowding and Risk Taking in a Tournament: Evidence from NASCAR Racing. *Administrative Science Quarterly* 52, 208–247.
- Bowman, N. A. and M. N. Bastedo (2009). Getting on the front page: Organizational reputation, status signals, and the impact of U.S. news and world report on student decisions. *Research in Higher Education* 50(5), 415–436.
- Boyle, E. and Z. Shapira (2012). The Liability of Leading: Battling Aspiration and Survival Goals in the Jeopardy! Tournament of Champions. *Organization Science* 23(4), 1100–1113.
- Briscoe, F. and S. Safford (2008). The Nixon-in-China Effect: Activism, Imitation, and the Institutionalization of Contentious Practices. *Administrative Science Quarterly* 53, 460–491.
- Brown, K. C., W. Harlow, and L. T. Starks (1996). Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry. *The Journal of Finance* 51(1), 85–110.
- Cabral, L. M. B. (2003). R&D Competition When Firms Choose Variance. *Journal of Economics & Management Strategy* 12(1), 139–150.
- Carroll, G. R. and A. Swaminathan (2000). Why the Microbrewery Movement? Organizational Dynamics of Resource Partitioning in the US Brewing Industry. *American Journal of Sociology* 106(3), 715–762.

- Christensen, C. M. and J. L. Bower (1996). Customer Power, Strategic Investment, and the Failure of Leading Firms. *Strategic Management Journal* 17(3), 197–218.
- Cyert, R. and J. G. March (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.
- Deephouse, D. L. (1999). To Be Different, or to Be the Same? It’s a Question (And Theory) of Strategic Balance. *Strategic Management Journal* 20(2), 147–166.
- Edelman, L. B. (1990). Legal Environments and Organizational Governance: The Expansion of Due Process in the American Workplace. *American Journal of Sociology* 95(6), 1401–1440.
- Eliashberg, J., A. Elberse, and M. A. A. M. Leenders (2006). The Motion Picture Industry: Critical Issues in Practice, Current Research, and New Research Directions. *Marketing Science* 25(6), 638–661.
- Glaser, R. (1963). Instructional technology and the measurement of learning outcomes: Some questions. *American Psychologist* 118, 519–521.
- Gould, R. V. (2002). The Origins of Status Hierarchies: A Formal Theory and Empirical Test. *American Journal of Sociology* 107(5), 1143–1178.
- Greve, H. R. . (1998). Performance, Aspirations, and Risky Organizational Change. *Administrative Science Quarterly* 43(1), 58–86.
- Greve, H. R. (2003a). A Behavioral Theory of R&D Expenditures and Innovations: Evidence from Shipbuilding. *Academy of Management Journal* 46(6), 685–702.
- Greve, H. R. (2003b). *Organizational Learning from Performance Feedback: A Behavioral Perspective on Innovation and Change*. Cambridge: Cambridge University Press.
- Hannan, M. T. and J. Freeman (1984). Structural Inertia and Organizational Change. *American Sociological Review* 49(2), 149–164.
- Hannan, M. T., L. Pólos, and G. R. Carroll (2003). The Fog of Change: Opacity and Asperity in Organizations. *Administrative Science Quarterly* 48(3), 399–432.

- Hannan, M. T., L. Pólos, and G. R. Carroll (2007). *Logics of Organization Theory: Audiences, Codes, and Ecologies*. Princeton: Princeton University Press.
- Hörner, J. (2004). A Perpetual Race to Stay Ahead. *Review of Economic Studies* 71(4), 1065–1088.
- Hsu, G. (2006a). Evaluative schemas and the attention of critics in the US film industry. *Industrial and Corporate Change* 15(3), 467–496.
- Hsu, G. (2006b). Jacks of all trades and masters of none: Audiences’ reactions to spanning genres in feature film production. *Administrative Science Quarterly* 51(3), 420–450.
- Iyer, D. N. and K. D. Miller (2008). Performance feedback, slack, and the timing of acquisitions. *Academy of Management Journal* 51(4), 808–822.
- Kacperczyk, A., C. M. Beckman, and T. P. Moliterno (2015). Disentangling Risk and Change: Internal and External Social Comparison in the Mutual Fund Industry. *Administrative Science Quarterly* 60(2), 228–262.
- Kahneman, D. and A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47(2), 263–292.
- Kaniel, R. and R. Parham (2017). WSJ Category Kings - the impact of media attention on consumer and mutual fund investment decisions. *Journal of Financial Economics* 123(2), 337–356.
- King, B. G. and N. A. Pearce (2010). The Contentiousness of Markets: Politics, Social Movements, and Institutional Change in Markets. *Annual Review of Sociology* 36(1), 249–267.
- Kovács, B. and M. T. Hannan (2015). The Space of Categories and the Consequences of Category Spanning. *Sociological Science* 2, 252–286.
- Kovács, B. and C. Liu (2016). Audience structure and status multiplicity. *Social Networks* 44, 36–49.
- Lehman, D. W., B. Kovács, and G. R. Carroll (2014). Conflicting Social Codes and Organizations: Hygiene and Authenticity in Consumer Evaluations of Restaurants. *Organization Science* 60(10), 2602–2617.

- Levinthal, D. A. (1997). Adaptation on Rugged Landscapes. *Management Science* 43(7), 934–950.
- Lynn, F. B., J. M. Podolny, and L. Tao (2009). A Sociological (De)Construction of the Relationship between Status and Quality. *American Journal of Sociology* 115(3), 755–804.
- March, J. G. and H. A. Simon (1958). *Organizations*. New York: Wiley.
- Merton, R. K. (1968). The Matthew Effect in Science. *Science* 159(3810), 56–63.
- Miller, D. and M.-J. Chen (1994). Sources and Consequences of Competitive Inertia: A Study of the U.S. Airline Industry. *Administrative Science Quarterly* 39(1), 1–23.
- Mohn, K. (2008). Efforts and Efficiency in Oil Exploration: A Vector Error-Correction Approach. *The Energy Journal* 29(4), 53–78.
- Natriello, G. (1987). The Impact of Evaluation Processes on Students. *Educational Psychologist* 22(2), 155–175.
- Nelson, R. R. and S. G. Winter (1982). Organizational Capabilities and Behavior. In *An Evolutionary of Economic Change Theory*.
- Pfeffer, J. and G. Salancik (1978). *The External Control of Organizations*. New York: Harper & Row.
- Phillips, D. J., C. J. Turco, and E. W. Zuckerman (2013). Betrayal as Market Barrier: Identity-Based Limits to Diversification among High-Status Corporate Law Firms. *American Journal of Sociology* 118(4), 1–32.
- Phillips, D. J. and E. W. Zuckerman (2001). Middle-Status Conformity: Theoretical Restatement and Empirical Demonstration in Two Markets. *American Journal of Sociology* 107(2), 379–429.
- Podolny, J. M. (1993). A Status-Based Model of Market Competition. *American Journal of Sociology* 98(4), 829–872.
- Podolny, J. M. (2001). Networks as the Pipes and Prisms of the Market. *American Journal of Sociology* 107(1), 33–60.

- Podolny, J. M. (2008). *Status Signals: A Sociological Study of Market Competition*. Princeton: Princeton University Press.
- Pratt, J. W. (1964). Risk Aversion in the Small and in the Large. *Econometrica* 32(1), 122–136.
- Rao, H., P. Monin, and R. Durand (2005). Border Crossing: Bricolage and the Erosion of Categorical Boundaries in French Gastronomy. *American Sociological Review* 70(6), 968–991.
- Rivkin, J. W. (2000). Imitation of Complex Strategies. *Management Science* 46(6), 824–844.
- Rogers, E. M. (2003). *Diffusion of Innovations* (5 ed.). New York: Free Press.
- Rossman, G. and O. Schilke (2014). Close, But No Cigar: The Bimodal Rewards to Prize-Seeking. *American Sociological Review* 79(1), 86–108.
- Sauder, M. and W. N. Espeland (2009). The Discipline of Rankings: Tight Coupling and Organizational Change. *American Sociological Review* 74(1), 63–82.
- Simon, H. A. (1947). *Administrative Behavior*. New York: Macmillan.
- Simon, S. R., R. Kaushal, P. D. Cleary, C. A. Jenter, L. A. Volk, E. G. Poon, E. J. Orav, H. G. Lo, D. H. Williams, and D. W. Bates (2007). Correlates of Electronic Health Record Adoption in Office Practices: A Statewide Survey. *Journal of the American Medical Informatics Association* 14(1), 110–117.
- Simtowe, F., J. Mduma, A. Phiri, A. Thomas, and M. Zeller (2006). Can risk-aversion towards fertilizer explain part of the non-adoption puzzle for hybrid maize? Empirical evidence from Malawi. *Journal of Applied Sciences* 6(7), 1490–1498.
- Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics* 87(3), 355.
- Stewart, S. M., M. L. Gruys, and M. Storm (2010). Forced distribution performance evaluation systems: Advantages, disadvantages and keys to implementation. *Journal of Management & Organization* 16(1), 168–179.
- Vergne, J. P. and T. Wry (2014). Categorizing Categorization Research: Review, Integration, and Future Directions. *Journal of Management Studies* 51(1), 56–94.

- White, H. C. (1981). Where Do Markets Come From? *American Journal of Sociology* 87(3), 517–547.
- Zuckerman, E. W. (1999). The Categorical Imperative: Securities Analysts and the Illegitimacy Discount. *American Journal of Sociology* 104(5), 1398–1438.
- Zuckerman, E. W., T. Y. Kim, K. Ukanwa, and J. V. Rittmann (2003). Robust Identities or Nonentities? Typecasting in the Feature-Film Labor Market. *American Journal of Sociology* 108(5), 1018–1073.