

The Development and Exploitation of Breakthrough Inventions

Charles Connaughton

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

H. Kevin Steensma, Chair

Warren Boeker

Abhinav Gupta

Program Authorized to Offer Degree:

Foster School of Business, Department of Management and Organization

©Copyright 2021
Charles Connaughton

University of Washington

Abstract

The Development and Exploitation of Breakthrough Inventions

Charles Connaughton

Chair of the Supervisory Committee:

H. Kevin Steensma

Foster School of Business, Department of Management and Organization

I use a cognitive lens to examine how firms achieve and benefit from breakthrough inventions. I propose a life cycle where firms initially use knowledge creatively before routinely exploiting it, but mastery from use renews creativity. Firms are better able to exploit their breakthrough inventions with experience but are impeded by bureaucracy. Utilizing machine learning on 5.5 million USPTO patent abstracts, I identify and analyze 1,644 breakthrough inventions. I find that firms are most likely to produce breakthrough inventions when combining prior art that they are either intimately familiar or completely unfamiliar with. My findings therefore suggest coexisting but incompatible mechanisms for creating breakthrough inventions. Interestingly, firms that are old and small are best able to benefit from their breakthroughs. This finding suggests that the accumulated growth of old firms comes at a cost of a decreased ability to benefit from breakthrough inventions.

INTRODUCTION

Breakthrough inventions disrupt existing paradigms, launch new technological trajectories, and in doing so, establish fresh vocabularies and cognitive maps (Kuhn, 1962; Dosi, 1982). These extraordinary inventions fuel the engine of creative destruction (Schumpeter, 1942) and potentially provide sustainable competitive advantages for the firms that invent them (Barney, 1991). While previous studies have documented the disruptiveness of certain high-profile breakthrough inventions (Dosi, 1982; Arthur, 2007), there are few generalizable insights on how firms develop them, and which firms are inclined to exploit the ones they develop. Specifically, what role does a firm's technical expertise play in the development of breakthrough inventions? Once developed, which firms have the wisdom and wherewithal to exploit their breakthrough inventions with follow-on inventions as opposed to abdicating these opportunities to other firms?

There are alternative perspectives to these questions. On one hand, breakthrough inventions may emerge when firms deploy their in-depth expertise by recombining technologies with which they are intimately familiar in highly creative and disruptive ways (Henderson and Clark, 1990). On the other hand, expertise and associated recombinatorial routines can constrain creativity (Henderson, 1993); breakthrough inventions may occur when firms apply fresh perspectives unjaded by expertise and routine to recombine technologies that are relatively unfamiliar to them (Ahuja and Lampert, 2001).

Developing a breakthrough invention is merely the first step in the evolution of a new technological trajectory. To fully exploit the potential of breakthrough inventions, complementary follow-on inventions that build on the respective breakthrough are often essential (Rosenberg, 1976). Firms that develop breakthrough inventions, however, are not always the

ones to develop the necessary follow-on inventions (Teece, 1986). Xerox famously initiated the computer revolution with its initial breakthrough technology only to relinquish the trajectory of this technology to rival firms. Similarly, Eastman Kodak pioneered digital photography technology, yet failed to capitalize on its breakthrough invention and ultimately went bankrupt (Lucas and Goh, 2009). There is competing speculation as to whether established or nascent firms more effectively exploit their respective breakthrough inventions (Mueller and Tilton, 1969; Henderson, 1993; Arthur, 2007). On one hand, entrenched bureaucracy that is prevalent in established incumbent firms may impede their ability to pivot resources toward exploiting the breakthrough inventions they develop (Nelson and Winter, 1982). On the other hand, established incumbent firms are more likely to have experienced prior paradigm shifts and thus be wise to the potential of their breakthrough inventions (Teece, Pisano, and Shuen, 1997).

Few studies have considered these competing views systematically because it is difficult to precisely distinguish breakthrough from non-breakthrough inventions across a broad array of firms. The inability to do so prevents the use of rigorous research designs that contrast the characteristics of firms' breakthrough inventions with their otherwise comparable non-breakthrough inventions, as well as firms' exploitation of their breakthroughs (relative to other firms exploiting the focal breakthrough) with their exploitation of comparable non-breakthrough inventions (relative to other firms exploiting the comparable non-breakthrough). I address these challenges by using natural language processing to analyze the written abstracts of 5.5 million patents granted by the USPTO between 1969 and 2017, and conservatively identify 1,644 breakthrough inventions across a wide range of industries. I then carefully match each breakthrough invention with similar non-breakthrough inventions developed by the same firm as had developed the focal breakthrough invention.

My initial hypothesis and analysis suggest that breakthrough inventions tend to be recombinations of prior art technologies that are either highly familiar or highly unfamiliar to the inventing firm. Breakthrough inventions are less likely to result from firms recombining prior art that is moderately familiar to them. These results suggest a lifecycle of creative output that firms can anticipate from the prior art technologies they integrate into their research pursuits. Initially, firms enter a creative honeymoon as they apply fresh perspectives to recombine prior art technologies that are relatively unfamiliar to them. As a firm becomes increasingly familiar with the prior art in question, however, they enter a creative lull; entrenched recombinatorial routines that evolve from preceding recombinations of the specific prior art technologies hamper the creativity of subsequent recombinations. Nonetheless, if firms continue to recombine these prior art technologies, they may ultimately experience a creative revival; as firms develop in-depth expertise and insights regarding specific prior art, they start to question and discount the recombinatorial routines that previously constrained their creativity. Such a lifecycle aligns with the notion that firms ideally explore unfamiliar terrain while continuing to exploit what they already know (O'Reilly and Tushman, 2004). Firms may benefit from patient on-going R&D investment in technical domains to work their way out of creative lulls and into creative revivals.

My subsequent hypotheses and analysis disentangle the influence of a firm's age and size on its tendency to exploit the breakthrough inventions it develops with follow-on inventions as opposed to other firms doing so. I find that older firms are more inclined to build on their own breakthrough inventions than are their younger counterparts. All else equal, older firms have had greater opportunities to learn both experientially and vicariously how breakthrough inventions shift paradigms and launch new technological trajectories. Such lessons provide firms with perspective to recognize the potential of the breakthrough inventions they create. I also find,

however, that larger firms are less adept than their smaller counterparts at exploiting their breakthrough inventions relative to their ability to exploit their comparable non-breakthrough inventions. As firms grow, their increasing bureaucracy, rigid communication routines, and internal political rivalries hinder their ability to pivot toward their breakthrough inventions and redeploy resources to develop follow-on inventions. While smaller firms may have fewer resources overall, they are more agile in redeploying them toward exploiting the breakthrough inventions they develop.

These findings provide nuance to our understanding of how maturity and growth influence firms' capacities to exploit the breakthrough inventions they develop. Overall, older, smaller firms may be best positioned to exploit their own breakthrough inventions. They possess valuable wisdom that enables them to recognize the potential of their own breakthrough inventions yet lack the bureaucracy that inhibits their ability to pivot resources toward exploiting them. Consequently, firm age and growth have countervailing effects on a firm's capacity to exploit its breakthrough inventions. While growth is beneficial in terms of economies of scale and scope (Chandler, 1962), these benefits come at a cost - a decreased ability to reap the rewards of the breakthrough inventions they develop. Although managers have little agency over the maturation of their firms, our results suggest they need to carefully weigh the advantages and disadvantages of growth and how it may impede their ability to exploit their own breakthrough inventions.

TECHNOLOGICAL PARADIGMS AND BREAKTHROUGH INVENTIONS

Technical knowledge is initially difficult to understand and communicate. Concepts are complex, unintuitive, or interact with each other in intricate ways. Over time, however, precise

vocabularies and frameworks evolve to form the basis of a technological paradigm. Such paradigms help inventors make sense of concepts and relationships between them and develop derivative inventions that predictably advance along a technological trajectory. On rare occasions, however, an invention disrupts the prevailing technological trajectory. These breakthrough inventions are particularly creative and extraordinary in that there are no established vocabularies to adequately describe them, and repurposing existing vocabularies misrepresents their potential. They challenge current frameworks and assumptions associated with the respective trajectory (Kaplan and Vakili, 2015). Concepts that had been considered critical may be less relevant. Previously discarded concepts that had been assumed to be of little value may reemerge as promising. Consequently, breakthrough inventions cause prevailing paradigms to shift, branch, or be dislodged, and launch distinctly new technological trajectories (Dosi, 1982).

Inventions typically recombine knowledge embodied in previous inventions (Nelson, 1959; Arrow, 1972). Whether an invention is a breakthrough in terms of disrupting a prevailing paradigm is determined more by how underlying prior art technologies are recombined than by what prior art is being recombined (Kaplan and Tripsas, 2008; Kuhn, 1962). As such, characterizations of inventions based on their backward and forward citations (e.g., novelty, value) fail to effectively distinguish breakthrough inventions. Figure 1 plots the breakthrough inventions that I identified by analyzing patent abstracts using natural language processing (described more fully in the methods section below) along the dimensions of their backward and forward citations. The range of backward and forward citations associated with these breakthrough inventions is extensive.

Insert Figure 1 about here

Inventions that build on little prior art are often characterized as novel (Ahuja and Lampert, 2001). Although such novel inventions may trigger paradigm shifts and initiate new trajectories, most do not (Fleming, 2001). Breakthrough inventions that lack backward citations to prior art technologies are frequently those that initiate technological trajectories through academic scientific discovery. For example, Genentech's influential hybrid immunoglobulin patent invention (US5116964) cited over 100 scientific journal articles and no patented inventions.

Inventions need not be novel to initiate new technological trajectories. For example, Royal Dutch Shell's invention underlying hydraulic fracking (US6581684) patented in 2001 revolutionized the oil industry by recombining insights from 353 patented inventions originating from as far back as 1885 (see Figure 1). Similarly, International Game Technology introduced a client-server model for gambling machines (US7568973) in 2006 that disrupted the industry. This architectural change in gambling machines built on 511 previously patented inventions and quickly became the dominant design.

Inventions that are highly cited by subsequent inventions (i.e., forward citations) are often characterized as valuable (Albert et al, 1991). Because breakthrough inventions provide the foundations for new trajectories, they tend to be cited extensively by subsequent inventions that build on them. Nonetheless, highly cited patented inventions are not necessarily breakthrough inventions. When part of an established trajectory, non-breakthrough inventions are easily recombined with complementary inventions within the same trajectory and thus are often highly

cited as well. While Xerox's breakthrough digital copy protection scheme (US5629980) is the most highly cited breakthrough invention in our sample, Xerox has a multitude of other patented inventions with similar rates of forward citations that have not initiated new technological trajectories. For example, its patented invention (US5715403) introducing the acceptance of a user agreement before accessing copy protected works accumulated roughly the same number of forward citations as its digital copy protection scheme over the same time-period; however, it was not particularly disruptive.

Moreover, breakthrough inventions are not necessarily highly cited. Monsanto's patented breakthrough invention (US8198509) in 2008 introduced a soybean resistant to the southern root knot nematode. Gene sequences are not readily recombined in subsequent inventions, however, and not heavily cited by subsequent inventions. Breakthrough patents may also occur before their time and fail to gain early traction in terms of forward citations. Hitachi's patented invention (US6388986) in 2001 entails how to record data radially within a disc drive. Although this patented invention was not cited by follow-on inventions initially, it eventually formed the basis for a new technological trajectory.

In sum, accounting for the prior art that a respective invention recombines, or the extent to which subsequent inventions build on the respective invention fails to distinguish between inventions that disrupt prevailing technological paradigms and those that do not. Rather, breakthrough inventions are distinct due to the creative manner in which prior art technologies are combined (Kaplan and Tripsas, 2008). How creatively prior art technologies are recombined will depend on the inventing firm's familiarity with them and how entrenched are their respective recombinatorial routines.

RECOMBINATORIAL ROUTINES AND THE DEVELOPMENT OF BREAKTHROUGH INVENTIONS

As various activities within firms such as production and distribution are routinized through repetition, they become highly efficient and predictable (Becker, 2004; Levitt and March, 1988). Nonetheless, routines are rigid because individuals tend to resist changing the behaviors that brought them past success and discount alternative perspectives and behaviors (Madsen and Desai, 2010; March and Shapira, 1987). While adhering to routines may be beneficial in the short term due to their efficiencies and predictability, firms may fall into competency traps; ongoing adherence to routines further enhances their efficiency making it increasingly difficult to deviate from them (March, 1991). Such competency traps can eventually lead to suboptimal performance.

Although all inventions are unique to varying degrees, like other activities within the firm, the process of recombining prior art technologies to develop technical solutions can become routine. As firms grow increasingly familiar with specific prior art technologies from having repeatedly recombined them in their past inventive efforts, they develop heuristics and short cuts to efficiently mix and match the respective prior art with highly predictable outcomes (Nelson and Winter, 1982). Such recombinatorial routines, however, discourage firms from deploying fresh perspectives for recombining prior art technologies; particularly creative and potentially disruptive changes in technological architecture go unheeded (Henderson, 1993). As such, entrenched recombinatorial routines are more likely to give rise to inventions that sustain prevailing paradigms than inventions that initiate fresh technological trajectories. In contrast, when firms recombine prior art technologies that are new to them, they do so unconstrained by an established playbook of heuristics and recombinatorial routines. Firms recombine unfamiliar

prior art with the naivete that spawns particularly creative inventions and disrupts existing technological paradigms. Only as firms become increasingly familiar with specific prior art from recombining it repeatedly do creativity-hindering recombinatorial routines become entrenched.

Although routine behaviors are generally resistant to change, they need not persist indefinitely (Feldman, 2000). At some point, firms grow dissatisfied with the outcomes from their routines and are open to experimentation (Holmqvist, 2004). As firms garner ever more experience performing a specific activity, they eventually develop deep expertise that underscores limitations in how they have done things in the past. In the words often attributed to Aristotle, “The more you know, the more you know you don’t know.” In-depth expertise can engender the humility necessary for applying fresh perspectives that dislodge entrenched routines. Likewise, the constrictions of firms’ recombinatorial routines are likely to reach a tipping point. As a firm attains greater familiarity with specific elements of prior art from having extensively recombined them in the past, the limitations of the heuristics and short cuts that have guided their recombinations become evident. In-depth expertise breeds the confidence to recombine prior art in sophisticated and creative ways (Amabile, 1983; Taggar, 2002).

Accordingly, firms transition through phases of creativity as they repeatedly recombine specific elements of prior art in their inventive processes. At the onset, they enter a particularly creative honeymoon phase where recombinations are uninhibited by routine. A particularly striking example from my sample of breakthrough inventions is Broadcom’s multi-stage integrated circuit packaging invention as patented (US6853070) in 2001. As a new entrant to a well-established technological trajectory, Broadcom developed a breakthrough process improvement by recombining 104 prior art technologies on semiconductor packaging, none of which Broadcom had previously built on in their previous patented inventions.

The honeymoon phase, however, is followed by a creative lull as recombinatorial routines take hold. Firms that recombine prior art technologies that are neither highly familiar nor unfamiliar lack the naivete that allow them to deploy fresh perspectives to combine the prior art in question in particularly new and creative ways; they also lack the in-depth expertise from having extensively recombined the respective prior art that highlight the limitations of their entrenched routines and enable them to see past the prevailing paradigms in which they are working.

Nonetheless, as firms continue to develop deep expertise and familiarity with the respective prior art technologies, a creative revival may ensue where they begin to discount their entrenched recombinatorial routines. An example from my sample is Dexcom's breakthrough signal processing invention (US8005524) patented in 2010 that entails a method of consistently measure blood sugar levels despite instability in biological measurements. This breakthrough invention represented the culmination of Dexcom's research on continuous glucose monitoring and built on 264 prior art elements that Dexcom had previously recombined 25,233 times in the five years prior to their breakthrough invention.

All told, the relationship between a firm's familiarity with the prior art technologies it recombines to generate an invention and the likelihood of the invention being a breakthrough (vis-à-vis a non-breakthrough) is U-shaped. Inventions that recombine prior art that is either highly familiar or unfamiliar to the inventing firm are more likely to be breakthrough inventions; those inventions that recombine prior art that is moderately familiar to the inventing firm are least likely to be breakthrough inventions.

H1: The relationship between a firm's familiarity with the prior art technologies being recombined and the likelihood of the ensuing invention being breakthrough is U-shaped; the likelihood is greatest when familiarity is relatively high or low.

EXPLOITING BREAKTHROUGH INVENTIONS WITH FOLLOW-ON INVENTIONS

Developing breakthrough inventions is merely the first step. The technological trajectory initiated by a breakthrough invention typically requires the development of many complementary follow-on inventions building upon the breakthrough (Arthur, 2007). Although various firms may develop breakthrough inventions, some will be better than others at recognizing the potential of their breakthrough inventions, and subsequently pivoting their resources toward creating follow-on inventions in support of breakthrough inventions.

Age, wisdom, and firms' abilities to recognize their inventions to be breakthroughs

The knowledge of an organization is comprised of more than its granular technical knowledge and the routines it uses to recombine such technology. Through both experiential and vicarious learning, firms accumulate higher-level strategic wisdom on how the world works and cause-and-effect relationships (Huber, 1991).

As firms pursue relatively rare, yet highly consequential activities such as corporate acquisitions and international market expansions, they develop wisdom regarding these activities (Delios and Henisz, 2003; Finkelstein and Halebian, 1999). To make sense of the outcomes from their various pursuits or perhaps missed opportunities, firms compose narratives that are stored in their memories through standard operating procedures and stories. These narratives provide a means of communicating basic lessons within the firm that can be drawn on when related situations arise (Garud, Dunbar, and Bartel, 2011; Huber, 1991). Firms develop strategic wisdom not only through their own experiences, but also vicariously by observing the experiences of other organizations (Beckman and Haunschild, 2002; Denrell, 2003). Vicarious learning is particularly effective when firms observe the failures of others as these failures are

often publicized and salient (Madsen and Desai, 2010); they serve as wake-up calls that highlight how such missteps can threaten one's survival (Kim and Miner, 2007). Because organizational memory evolving from experiential and vicarious learning is not tied to specific individuals, it is resilient to staff turnover (Cyert and March, 1963; Weick and Gilfillan, 1971).

The history-dependent wisdom that firms acquire either through their own experiences or observing those of others influence how they respond to changing environments (Hannan and Freeman, 1984). Such wisdom will influence a firm's ability to recognize the technological breakthroughs it develops for what they are. Over the course of a firm's lifetime, it develops perspective regarding how technological trajectories emerge and how to navigate paradigm shifts. In the past, a firm may have successfully exploited its own breakthrough or may have erroneously passed over opportunities to do so. Regardless, such experiential lessons become seared in a firm's memory. With the benefit of time, firms also have opportunities to observe the strategic missteps of their counterparts when they have foregone opportunities to exploit the breakthrough inventions they have developed.

With wisdom developed over time, the myopia of adhering to prevailing technological paradigms and erroneously disregarding one's own breakthrough inventions is held in check. As such, older firms are wiser to the potential of the breakthrough inventions they develop internally and will be more inclined to develop follow-on inventions than will firms that lack the memory of missed opportunities and wisdom that comes with age. Toshiba and Lucent Technology's handling of their respective breakthrough inventions exemplify the benefits of age. The highly seasoned Toshiba quickly built on its digital bookmarks (US6360057) breakthrough invention it patented in 2000 with follow-on inventions that secured its leadership in this area. In contrast,

newly formed Lucent did little with its magnetic integrated circuit breakthrough (US 6118351) it invented in 1997 and subsequently relinquished this technology to Agere Systems.

H2: Older firms are more effective at exploiting their own breakthrough inventions by developing follow-on inventions than are younger firms.

Size, bureaucracy and firms' inabilities to pivot resources towards their respective breakthrough inventions

Exploiting one's own breakthrough inventions entails firms not only recognizing their potential, but also pivoting resources toward developing follow-on inventions. Larger firms will generally have greater difficulty doing so due to their entrenched bureaucracy.

Bureaucracy is the alignment of highly formalized procedures, specialized departments, vertical hierarchies, and centralized decision making (Pugh, Hickson, and Hinings, 1969). Despite the changing nature of work, such alignment within large organizations has remained pervasive over the decades (Walton, 2005). Formalized bureaucratic structures are particularly effective for scaling activities (Simon, 1976). To streamline internal communication and contend with ever increasing demands on employees' attention, firms develop communication routines, composed of clearly defined channels and standardized vocabularies to quickly distill and share important information within their growing organizations (Lewin, Massini, and Peeters, 2011).

Despite these benefits, such bureaucracy hampers firms from exploiting the breakthrough inventions it develops. While a streamlined vocabulary allows the firm to spread information quickly, it comes at the cost of richness. Standardized vocabularies and defined channels of communication are prone to filtering out insights that are inconsistent with existing technological trajectories. The nuances of breakthrough inventions and their potential will be lost when

ingrained vocabularies that are based on existing technological paradigms are used to diffuse information on them throughout the organization (Henderson, 1993). Thus, even if individuals within the firm recognize and advocate for the potential of breakthrough inventions developed within, it will be relatively difficult to gain the broad support necessary to secure the resources to build on them.

Furthermore, as argued by Weber (1958), bureaucracy is an iron cage that promotes efficiency at the price of alienation. Specialization can lead to internal political rivalries that may impede firms from further building on the breakthrough inventions they develop (Benner and Tushman, 2003; Crozier, 1964). Breakthrough inventions can be organizationally disruptive as they displace prevailing technological trajectories around which existing networks of relationships and bureaucratic power structures had been built. Individuals may be reluctant to relinquish the power and status they derive from the prevailing technological trajectory that the firm was on and work against diverting resources toward building on breakthrough inventions.

Smaller firms lack the bureaucracy and rigid communication routines of their larger counterparts. With fewer employees to coordinate, there is less need for extensive bureaucracy and parsimony to coordinate the flow of information, making it easier for individuals to disseminate the knowledge comprising a breakthrough invention. Smaller firms lack the political rivalries between specialized units that may impede diverting resources toward the development of breakthrough inventions. Consequently, in contrast with their smaller competitors, larger organizations are less capable of pivoting resources toward exploiting their own breakthrough inventions.

Energy Conversion Devices and General Motor's handling of breakthrough inventions exemplify the effects of size. The small photovoltaics manufacturer Energy Conversion Devices

invented a breakthrough phase change memory system (US5166758) in 1991 and quickly developed the first six follow-on inventions based on its breakthrough. As a smaller firm, it was able to focus attention and resources on developing their breakthrough invention. General Motors' breakthrough technique that integrated metallic and ceramic elements into stereolithography components (US5705117) created a new class of integrated circuit devices. Because the applicability of this breakthrough for motor vehicles was unclear, General Motors was unable to pivot resources to further develop its breakthrough invention. Although this breakthrough invention has accumulated over 140 forward citations, General Motors has produced no follow-on inventions from its breakthrough.

H3: Larger firms are less effective at exploiting their own breakthrough inventions by developing follow-on inventions than are smaller firms.

METHODS

Sample and Data

To identify breakthrough inventions, I deployed topic modeling techniques on the abstracts of patented inventions. Topic modeling has been used increasingly in the social sciences to analyze the content of text and is a powerful tool for identifying distinct topics within patent abstracts (e.g., Kaplan and Vakili, 2015, Arts, Cassiman, and Gomez, 2018). Similar measures utilizing language to identify breakthrough inventions have been validated on inventions from as early as 1840 (Kelly, et al, 2020). Patent abstracts are clear and concise summaries of the claims to intellectual property made by each patented invention and are vetted by both the patent applicant and USPTO patent examiner. As applied, topic modeling assumes there are common topics across a collection of patent abstracts, and that the language within each abstract is based on these topics (Blei, 2012; Wilkerson and Casas, 2017). The algorithm

identifies sets of words that frequently appear together within the abstracts as distinct topics, and then determines which topic or combination of topics best describes each abstract.

I compiled the 5.5 million patent abstracts associated with the entire population of patents granted by the USPTO between 1969 and 2017. Prior to 1969, granted patents did not include abstracts. Analyzing the entire population of abstracts is essential not only to identify a broad range of breakthrough inventions, but to confidently distinguish what are true breakthroughs from non-breakthrough inventions. If my sample were limited to a specific industry, topic modeling techniques would overlook the fertilization of knowledge across industry. If true breakthrough inventions occur outside of the industry of interest and their underlying knowledge is later introduced to that industry through derivative inventions, these derivatives could be miscategorized as breakthrough inventions. Analyzing the entire population of patent abstracts protects against these false positives.

I analyzed my sample of patent abstracts using a combination of Latent Dirichlet Allocation (LDA) and Hierarchical Dirichlet Process (HDP) models. LDA identifies common characteristics of a specified number of topics to generate the best fitting topics from the patent abstracts. Individual patent abstracts are treated as combination of those topics. HDP works in a similar fashion but attempts to identify both the number of topics and the characteristics of those topics simultaneously (Wang, Paisley, and Blei, 2011). Although LDA is the most effective tool available to perform topic modeling, it has limitations. First, the number of topics to identify within a set of documents must be specified prior to analysis. I utilized topic density (Cao et al, 2009) to compare the quality of models with differing numbers of topics. Nonetheless, the a priori choice of how many topics to specify is subjective. Second, LDA can distinguish topics only if the set of documents being evaluated fall within a narrow domain. Applying an LDA

model to a diverse set of documents will only cleanly identify the most significant divisions in the data. Even within a set of documents from a narrow domain, an LDA model may find topics that do not map neatly onto the construct of interest. In addition to topics describing distinct technologies, an LDA model may also identify topics based upon differences in tone, syntax, or other idiosyncrasies of language. The content of discovered topics must be monitored carefully to ensure their validity. Third, LDA's complexity increases geometrically with the specified number of topics. Models attempting to identify more than 200 distinct topics are typically infeasible to compute.

To address these limitations, I developed a multi-stage methodology (Figure 2). Knowledge generally has a hierarchical structure (Hall, Jaffe, and Trajtenberg, 2001), with patented inventions belonging to broad categories (e.g., chemicals, electronics), more granular sub-categories, and precise topics of comprised of related inventions. However, technological categorization schemes typically focus on contemporaneous organization of capabilities (Strumsky, Lobo, and Van der Leeuw, 2012), rather than on carefully tracing the lineage of its knowledge. Instead of relying upon prior categorization efforts, I applied a series of LDA and HDP models to organically discover this hierarchical structure within patent abstracts.

Insert Figure 2 about here

I initially normalized, tokenized and stemmed each of the 5.5 million abstracts¹. For the first stage of my analysis, I estimated multiple LDA models on the entire population of abstracts

¹ Normalizing all documents to use identical forms of words is essential for text analysis. Important steps include removing accents from non-english words, use of number words, and addressing inconsistent hyphenation and use of compound words across documents. A normalized text is then tokenized, removing all punctuation and reducing a document to a set of words. Those words are then stemmed, which reduces different forms (plural, tense) of the

using various pre-specified numbers of topics. My objective was to identify the smallest number of topics where each topic matched a clearly identifiable category of inventions. In addition to utilizing statistical measures of model fit, I examined the keywords associated with each topic, representative abstracts on each topic, and the consistency of the identified topics across models with different numbers of topics. Consistency between those approaches led me to conservatively choose 21 broad categories, listed in Table 1.

Insert Table 1 about here

My chosen LDA model described each patent abstract as a mixture of these 21 broad topics. To simplify this for further analysis, I assigned each patent to a category corresponding with the topic from which it drew its largest proportion of text (Kaplan and Vakili, 2015).²

Each of the 21 categories were then further analyzed using an HDP model to identify sub-categories within that category. Unlike LDA, which assumes all topics are equally weighted within the data set, HDP assumes some topics are more prevalent than others, and based upon that assumption returns a rank ordered set of topics. I examined the results of the HDP model and identified the first topic that substantively overlapped with a previous topic. This overlap indicates a cut-off that determines the number of substantial topics contained within that

same word to its root. Additionally, common words with little independent informational value are removed to increase the efficiency of the model.

² One concern are the clusters of multi-disciplinary patents that do not fit cleanly into any one category. Subsequent analysis, however, showed that these clusters largely remained intact. Patent abstracts contained within each cluster were all assigned to the same broad category. Each of these clusters were ultimately assigned to unique finer-grained topics in our final stage analysis. As such, we are confident that the results of our overall analysis are not sensitive to the initial categorization of different clusters of multi-disciplinary patents in the first stage model because their uniqueness is accounted for in subsequent stages.

category. For example, from the category of vehicle patents, the first six topics identified by the HDP model all corresponded to distinct sets of patents that did not substantially overlap with each other. However, patents using the seventh topic overlapped almost entirely with patents using the first topic. Thus, I concluded that this category contained patents on six major topics. I then used the number of topics determined this way for each category to inform a new LDA model on each category. Finally, I then utilized that model and the previous methodology to assign each patent to a sub-category within its category.

Applying this methodology again to each sub-category identified a smaller number of further divisions, which were taken as further refinements of our sub-categorization. After this second pass of sub-categorization, HDP found only one substantial topic within each sub-category, suggesting that I had identified a parsimonious set of 682 sub-categories for further analysis. To confirm, I subsequently worked with another graduate student and a team of research assistants to examine the keywords and representative patents of each sub-category. We found that each of these 682 sub-categories consisted of patents that used generally similar language to describe a broadly similar set of inventions. For example, within the category of vehicle patents, one sub-category contained automobile transmission patents, while another covered engine patents.

Each sub-category was then analyzed one final time using LDA to identify specific topics within those sub-categories. For this last stage of analysis, I allowed for relatively large numbers of topics within each sub-category, as I wanted to capture the full scope of the variance within the data. A large number of topics additionally has the benefit of returning precisely defined topics (those corresponding to streams of knowledge). This precision allows me to distinguish trends arising from technological developments from idiosyncrasies related to syntax, tone, or

the like. Correspondingly, I selected a number of topics for each sub-category such that the addition of more topics would not substantially decrease the topic density (measured using cosine similarity).

I then reduced my sample to those patents with a strong model fit. Patents with low information content in their abstracts (less than one standard deviation below the mean for their category) were discarded. These are patents with short, non-descriptive abstracts that are poorly fit by the model. Additional analysis revealed that patents with short abstracts tended to be low value and low impact; inventors with substantial claims tended to write more substantial abstracts. As an additional check, I discarded patents that did not have at least a 30% fit with a single topic. Patents below that threshold lacked significant fit with any of the discovered topics, indicating they did not fit with any of the major streams of knowledge within the sub-category and could safely be eliminated as a breakthrough invention.

I further reduced my sample to only topics with many patents following a clear theme. Topics with fewer than ten patents primarily on that topic were discarded, as those topics had not established a sustained trajectory. I also eliminated all topics that lacked compact vocabularies of less than 6 bits per word. A clear and precise vocabulary was found to be strongly associated with a clear trajectory, and topics with more expansive vocabularies were much more likely to be capturing syntactic or idiosyncratic elements, rather than technological elements, of the patent abstracts. Together, these ensured that our sample of topics of interest were both well identified and demonstrated ex-post importance.

The first patented invention on each remaining topic was deemed to be a breakthrough invention. However, I was concerned about false positives – non-breakthrough inventions wrongly identified as breakthroughs due to censoring in the data. Patent abstracts were not

required in patent applications before 1974, so it would be possible for the true breakthrough on a topic to have occurred before 1974, with my sample consisting only of derivative patents on that topic. To address this censoring, I restricted my sample to only topics where the first patent was applied for in 1989 or later. This ensured a 15-year window preceding each identified breakthrough where no patents on its topic were filed, giving me confidence that I was observing the entire population of patents on that topic. With this restriction, I identified 3569 topics with breakthroughs occurring within the sample period.

To streamline the collection and ensure consistency of additional data required for my analyses, I further restricted my sample to the 1652 breakthrough inventions that had been created by publicly traded firms. To determine patent ownership, I utilized the NBER patent assignee database for patents between 1974 and 2006 (Bessen, 2009). This database accounts for mergers and acquisitions to determine ownership, and disambiguates the differing naming conventions for the same firm across patents. Because this database ends in 2006, I utilized similar methodology to extend this database through 2017 to cover the period of my sample. This allowed me to identify potential matching patents by the same firm for our control group, as well as determine which non-breakthrough patented inventions were created by the same firm.

I then matched each of these 1652 breakthrough inventions to highly similar patented inventions created by the same firm as the focal breakthrough invention. Patents identified as breakthrough inventions by my algorithm were often assigned to several patent classes and subclasses, with several assigned to a unique combination of classes. Consequently, matching on the exact class combination of a breakthrough patent was impossible, while matching any patent with at least one class in common with the breakthrough returned large numbers of very weak matched control patents. To resolve this, I identified patents with the most similar patent class

combinations to the focal patent, filed within two years of the patent application date of the breakthrough invention. This created an extremely conservative set of matches for my inferences – the matched sample includes many patents created not just by the same firm, but by the same inventors on the same day as the breakthrough invention. This matched sample controls for unobserved heterogeneity between firms, as well as heterogeneity between technological domains with patented inventions that may vary in their general propensity to generate citations from subsequent inventions. To reduce the impact of spurious correlations in the matched sample, three matches were found for each breakthrough patented invention. Of the original sample of 1652 breakthrough patents, only eight did not have a set of three reasonable matches (same firm, at least one overlapping subclass within a two-year period). These were dropped from my sample, and 1644 sets of breakthrough - control patented inventions from 357 unique firms remained for analysis (6576 total patents).

H1: The development of breakthrough inventions

Dependent Variable

Breakthrough invention. A dummy variable was used to indicate whether a patented invention was a breakthrough invention (1) or part of the control set (0).

Independent Variable

Prior art familiarity. To measure the focal firm's familiarity with the prior art cited by the focal patent, I added the number of times the focal firm had cited each prior art element in their previous inventions in the five years prior to the application of the focal patent. For example, if a focal patent cites two elements of prior art, one which has been previously cited twice, and one which has been previously cited three times by the focal firm in the previous five

years, the value of this measure would be five. If a prior art element had not been cited within the five-year time frame, the focal firm is assumed to no longer be familiar with the respective art (Katila and Ahuja, 2002).

To address extreme values, I transformed this measure using inverse hyperbolic sine (IHS). This transformation performs similarly to a logarithmic transformation but naturally transforms values of zero (Burbidge, Magee, and Robb, 1988).

Control Variables

Because I match each breakthrough invention with control inventions based on firm, subclass, and date, all firm, technology class, and temporal heterogeneity is controlled for by the research design. Any remaining measurable extraneous heterogeneity will be limited to characteristics of the focal patent's prior art.

Total prior art. To scale my measure of *prior art familiarity*, I control for the count of prior art elements cited by the focal patent. This measure was transformed using IHS.

First-time use of prior art. To control for the effect of exploring new prior art technologies and isolate the effect of a focal firm's *prior art familiarity* on the likelihood of *breakthrough invention*, I count the number of prior art elements cited by the focal patent that have not been previously cited by the focal firm in the previous five years. This measure was transformed using IHS.

Analysis and Results

For hypothesis 1, my unit of analysis is the patented invention. I use a linear probability model with control observations (non-breakthrough inventions) weighted on a 1:1 basis with the

breakthrough invention observations to predict the likelihood of each patented invention in my sample being a breakthrough invention. Because of the 1:1 basis, predicted probabilities will center around fifty percent. As such, the advantage of non-linear models for handling extreme probability events is not particularly germane.

Table 2 provides the descriptive statistics of my variables and their correlation coefficients. Table 3 provides my empirical models.

Insert Tables 2 and 3 about here

Model 1 of Table 3 is the baseline model and is comprised of all control variables. In Model 2, I added the main effect of *prior art familiarity*. In Model 3, I added the squared term of *prior art familiarity*.

Hypothesis 1 predicted that the relationship between a firm's familiarity with the prior art technologies being recombined and the likelihood of the ensuing invention being breakthrough is U-shaped; the likelihood is greatest when familiarity is relatively high or low. The squared term of prior art familiarity is positive and significant in Model 3 ($p < .01$). To better understand the relationship between *prior art familiarity* and the likelihood of *breakthrough invention*, I plotted it using Model 3 and holding all other variables at their sample mean (Figure 3).

Insert Figure 3 about here

Figure 3 shows that a patented invention in my sample is more likely to be a breakthrough invention when the inventing firm is either highly unfamiliar or highly familiar with the prior art technologies that it is recombining. Patented inventions in my sample were less

likely to be breakthrough inventions under moderate levels of *prior art familiarity*; the inflection point of this curve is approximately 45 prior citations (*prior art familiarity* = 4.5) of the prior art elements cited by the focal patented invention. For patented inventions built on prior art elements that had been previously cited by the focal firm over 10,000 times (*prior art familiarity* > 9.22), the additional predicted probability of it being a breakthrough invention is triple that of patented inventions in our sample built on no prior art elements (i.e., a novel invention). Hypothesis 1 is supported.

Notably, the coefficient associated with *first-time use of prior art* is also positive and significant suggesting that firms' integrating never-before-used prior art is generally beneficial for their generating breakthrough inventions.

I provide additional robustness checks on my results in Models 4, 5, and 6 of Table 3. Model 4 provides a robustness check on this study's matched design. It adds a fixed effect for each matched breakthrough-control set used in Model 3. This model controls for any residual confounding factors that were not controlled for by the matching. From this table, I observe that the point estimate for my independent variable has roughly the same coefficient in both Models 3 and 4. This suggests that the matching was effective in eliminating confounding factors, and that the pooled model is unbiased and efficient in estimating its effects.

Model 5 provides a robustness check on the linear approximation of Model 3. It is a pooled logistic model with otherwise identical specifications to Model 3. While the coefficients in Model 5 are log-odds coefficients from the logistic specification, their marginal effects at the mean are effectively identical to the coefficients from Model 3. Additionally, the standard errors and significance of the model are effectively unchanged from Model 3. This supports the use of a linear model to minimize assumptions and ease interpretation of my results.

Model 6 provides an additional robustness check on both the pooling and linearity assumptions of Model 3. It is a conditional logistic regression, conditioning on each breakthrough-control set. This is a conservative model that imposes a logistic transformation on a within-matched set design. While this design lacks the power of the previous models, its results are also consistent with Models 3, 4 and 5. This provides further confidence in the robustness of the results in Model 3.

H2 and H3: The exploitation of breakthrough inventions

Dependent Variable

Exploitation by inventing firm. I assess the extent to which the inventing firm exploited the focal patented invention by counting the subsequent patented inventions of the inventing firm that cite the focal patented invention within a 15-year time window of its application (Griliches, 1991).

Independent Variables

Breakthrough invention. A dummy variable was used to indicate whether a patented invention was a breakthrough invention (1) or part of the control set (0).

Firm age. The age of each inventing firm is calculated as the difference between the application date of the focal patented invention and the inventing firm's date of its initial public offering. To address extreme values, firm age was transformed using IHS.

Firm size. The size of each inventing firm is measured as the total sales of the firm for the year in which the patent for the focal invention was applied for. To address extreme values, total sales was transformed using IHS.

Control Variables

Total exploitation. I control for the extent to which each patented invention in our sample is exploited regardless of whether by the inventing firm or other firms by counting the number of times that the focal invention is cited by subsequent patented inventions within a 15-year time window of its application. Doing so allows me to infer the extent to which the inventing firm exploits its own invention (breakthrough or control) relative to other firms exploiting the focal patented invention. To address extreme values, this count was transformed using IHS.

Inventing firm total patents. Firms with greater numbers of patented inventions may have a higher capacity to absorb external knowledge (Cohen and Levinthal, 1990), and possibly a greater propensity to exploit their own inventions. I control for the count of all patented inventions produced by the inventing firm in the five years prior to the application year of the focal patented invention, including application year. To address extreme values, this count was transformed using IHS. I also control for the interaction of *inventing firm total patents* and *breakthrough invention* to account for the varying effect of the inventing firm's existing knowledge on its exploitation of breakthrough and non-breakthrough inventions.

Breadth of inventing firm patents. Firms with a broad base of technical knowledge have a wider variety of knowledge to draw on when exploiting their own inventions. I control for the scope of the focal firm's patents as measured by the inverse Herfindahl index of the patent classes of the firm's patented inventions in the five years prior to the application year of the focal patented invention, including application year. I also control for the interaction of *breadth of inventing firm patents* and *breakthrough invention*.

Firm performance. High performing firms may have greater ability to pursue subsequent development of their inventions in general. I control for the firm's return on assets for the year in which the focal patented invention was applied for. This variable was transformed using the IHS transformation. I also control for the interaction between *firm performance* and *breakthrough invention* to control for differing incentives between further development of breakthrough and non-breakthrough inventions.

Additionally, I control for industry fixed effects (based on 4-digit SIC code). Different industries face structurally different levels of competition and propensities to create follow-on inventions, leading to differences in expected behavior with respect to exploiting breakthrough and derivative inventions through the development of subsequent inventions. While my matched sample addresses structural biases in the entire population due to this confounding factor, it still explains a substantial portion of the variance in the dependent variable.

Analysis and results

For hypotheses 2 and 3, my unit of analysis is the patented invention. I use a weighted negative binomial model with control observations (non-breakthrough inventions) weighted on a 1:1 basis with the breakthrough inventions to predict *exploitation by inventing firm*. All coefficients are interpreted as average treatment effects of a focal invention being a breakthrough (Gelman and King, 1990).

Table 5 provides the descriptive statistics of my variables and their correlation coefficients. Table 6 provides my empirical models.

Insert Tables 4 and 5 about here

Model 1 of Table 5 is the baseline model and is comprised of all control variables. In Model 2, I added the main effects of *firm age* and *firm size*. In Model 3, I added the interaction terms of *firm age* and *firm size* with *breakthrough invention*.

Hypothesis 2 predicted that all else equal, older firms will produce more follow-on inventions from breakthrough inventions than will younger firms. This interaction term is positive and significant in Model 3 ($p < 0.05$). To better understand the relationship between *exploitation by the inventing firm* and the interaction of *firm age* and *breakthrough invention*, I plotted the relationship based on Model 3 in Figure 4, holding all other variables at their sample mean.

Insert Figure 4 about here

Figure 4 shows that older firms in my sample produced more follow-on inventions to their breakthrough inventions than younger firms. In contrast, there was not a significant difference between the exploitation of non-breakthrough inventions of young and old firms. The effect size is substantial; an older firm (age = 47 years, 90th sample percentile) produces approximately 35% more follow-on inventions from a breakthrough invention than a young firm (age = 4 years, 10th sample percentile). Mature firms with more than 20 years of experience post-IPO tended to produce more follow-on inventions from breakthrough inventions than from non-breakthrough inventions. The youngest firms (age = 0), on the other hand, produced 25% fewer follow-on inventions from their breakthroughs than from their non-breakthrough inventions. This supports Hypothesis 2.

Hypothesis 3 predicted that larger firms will produce fewer follow-on inventions from breakthrough inventions than smaller firms. This interaction term is negative and significant in Model 3 ($p < 0.05$). To better understand the relationship between *exploitation by the inventing firm* and the interaction of *firm size* and *breakthrough invention*, I plotted the relationship based on Model 3 in Figure 5.

Insert Figure 5 about here

Based on Model 3, relatively small firms (~100 million US in annual revenues) produce 75 percent more follow-on inventions from their breakthrough inventions than from their comparable non-breakthrough inventions. In contrast, for relatively large firms (~100 billion US in annual revenues), there is a negligible difference between follow-on inventions produced from their breakthrough inventions and their comparable non-breakthrough inventions. These findings highlight the relative agility of smaller firms in diverting resources toward exploiting their breakthrough inventions. Hypothesis 3 is supported.

I provide an additional robustness check on my result in Model 4 of Table 5. Model 4 uses a Poisson specification with fixed effects for each breakthrough-control matched set. While this specification is unable to estimate direct effects for firm age and firm size due to matching – and thus is more difficult to interpret – it provides a robustness check on the interaction terms from my hypotheses. In Model 4 the terms for the interaction between firm age and breakthrough invention, and the interaction between firm size and breakthrough invention, are both significant and similar to the terms in Model 3. This suggests that the findings in Model 3 are robust, and it

can be used to understand the joint effect of the direct and interaction terms for firm age and firm size.

DISCUSSION

My objective with this study was to investigate how firms create and exploit breakthrough inventions. I did this by first examining the prior art recombined by a breakthrough invention, and whether the inventing firm was exploiting familiar knowledge or exploring new knowledge in its creation. I then examined whether those breakthrough inventions were subsequently exploited by the inventing firm or by other organizations. These important questions have been difficult to address due to the difficulty of identifying breakthrough inventions without the use of patent citations. To solve this, I made use of the new technique of topic modeling to identify patents that introduce new language, and thus new ways of thinking.

My initial study finds that firms are most likely to create breakthrough inventions when they either combine wholly unfamiliar knowledge or exploit knowledge they are intimately familiar with. Inventions that combine familiar and unfamiliar knowledge elements are less likely to produce breakthrough inventions. This suggests a knowledge lifecycle, where new knowledge is first used creatively before being exploited routinely – but then revived creatively as the firm achieves mastery.

These findings have important implications. They suggest that the specific combination of prior art being combined in an invention is insufficient for understanding creative outcomes. I suggest that there are multiple ways a combination of prior art can be combined to create derivative inventions, and some of these ways are more useful than others. A firm's familiarity with a combination of prior art affects its perspective on that combination, which influences the

sorts of inventions they will derive from it. Due to this, some firms are more likely to produce breakthrough inventions from a particular combination of prior art than others.

These results are also suggestive to practitioners. In contrast with research suggesting that balancing the use of familiar and unfamiliar knowledge improves the value of inventions on both the firm (Fang, Lee, and Schilling, 2010) and patent (Arts and Veugelers, 2015) levels, my work suggest that firms looking to maximize the creativity of their inventions are better off segregating these activities and relying on either exploitation of familiar prior art, or eschewing its use in a search for new knowledge.

My subsequent studies suggest a nuanced take on the capacity of firms to develop derivative inventions from their breakthroughs. Experienced firms have a clear competitive advantage in the future development of their own breakthrough inventions, creating substantially more derivative inventions from them than their competitors. Firms lacking that experience will tend to see more of the development of their inventive breakthroughs performed by other firms. However, experienced firms tend to be larger, and that size can wipe out the advantages their experience confers. Large firms have difficulty identifying their breakthrough inventions, impeding their ability to benefit from their experience in its future development.

This finding draws an important distinction between breakthrough inventions and regular inventions with respect to knowledge spillovers. In general, inventions by well-established firms are more likely to be exploited by rivals than inventions by nascent firms (Park, Kim, and Choi, 2018); in contrast, the benefits of breakthrough inventions are disproportionately captured by the inventing firm when the inventing firm is older. Managers of older firms may capture a larger share of the benefits of their R&D activities when they focus their efforts on breakthrough, rather than incremental, inventions.

While this study provides a broad view on the creation and exploitation of breakthrough inventions, that comes at the cost of nuance. I hypothesize about the interplay between expertise and routines on a firm's ability to create and exploit breakthrough inventions, but those characteristics are inferred indirectly from broader traits of the inventing firm. Further studies will be necessary to better identify how expertise and routines interact to shape the creation of breakthrough inventions.

I have undertaken great effort to include a wide range of technologies in my study. However, there are important limitations on its scope posed by the sample and the methodology. The selection model used was able to identify breakthrough inventions across a variety of technological domains, but not all. My methodology depends upon shared language between similar patents to be able to detect clear themes. Stand-alone inventions without broader context, such as many agricultural and chemical patents, could not be identified well enough for inclusion. More targeted studies would be needed to establish the validity of my results for those classes of technologies.

The study is also limited by the period comprising the sample. Industries and technologies have their own life cycles, and I naturally weighted industries undergoing rapid technological change more heavily than those based upon mature technologies. There are many automotive patents in my sample, for instance, but relatively few of them were identified as breakthrough inventions. Information technology provided a substantial number of breakthrough inventions. Re-evaluating historical breakthroughs in the context of my results could help us better understand the evolution of those industries. Future research will also be necessary to understand the applicability of my results to new and emerging technologies, and under which conditions my results apply.

REFERENCES

Ahuja, G., & Lampert, C. M.

2001 "Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions." *Strategic Management Journal*, 22(6-7), 521-543.

Albert, M. B., D. Avery, F. Narin, and P. McAllister

1991 "Direct validation of citation counts as indicators of industrially important patents." *Research Policy*, 20(3), 251-259.

Amabile, T. M.

1983 "The social psychology of creativity: A componential conceptualization." *Journal of Personality and Social Psychology*, 45(2), 357.

Arrow, K. J.

1972 "Economic welfare and the allocation of resources for invention." In *Readings in Industrial Economics* (pp. 219-236). Palgrave, London.

Arthur, W. B.

2007 "The structure of invention." *Research Policy*, 36(2), 274-287.

Arts, S., and R. Veugelers

2015 "Technology familiarity, recombinant novelty, and breakthrough invention." *Industrial and Corporate Change*, 24(6), 1215-1246.

Arts, S., B. Cassiman, and J. C. Gomez

2018 "Text matching to measure patent similarity." *Strategic Management Journal*, 39(1), 62-84.

Barney, J.

1991 "Firm resources and sustained competitive advantage." *Journal of Management*, 17(1), 99-120.

Becker, M. C.

2004 "Organizational routines: a review of the literature." *Industrial and Corporate Change*, 13(4), 643-678.

Beckman, C. M., and P. R. Haunschild

2002 "Network learning: The effects of partners' heterogeneity of experience on corporate acquisitions." *Administrative Science Quarterly*, 47(1), 92-124.

Benner, M. J., and M. L. Tushman

- 2003 “Exploitation, exploration, and process management: The productivity dilemma revisited.” *Academy of Management Review*, 28(2), 238-256.
- Bessen, J.
- 2009 “NBER PDP Project User Documentation.” National Bureau of Economic Research.
- Blei, D. M.
- 2012 “Probabilistic topic models.” *Communications of the ACM*, 55(4), 77-84.
- Burbidge, J. B., L. Magee, and A. L. Robb
- 1988 “Alternative transformations to handle extreme values of the dependent variable.” *Journal of the American Statistical Association*, 83(401), 123-127.
- Cao, J., T. Xia, J. Li, Y. Zhang, and S. Tang
- 2009 “A density-based method for adaptive LDA model selection.” *Neurocomputing*, 72(7-9), 1775-1781.
- Chandler, A. D.
- 1962 “Strategy and structure: chapters in the history of American industrial enterprises.” Cambridge, Mass.: MIT Press, 14, 16.
- Cohen, W. M., and D. A. Levinthal
- 1990 “Absorptive capacity: A new perspective on learning and innovation.” *Administrative Science Quarterly*, 128-152.
- Crozier, M.
- 1964 “Le Phénomène Bureaucratique The Bureaucratic Phenomenon.” Tavistock Publications.
- Cyert, R. M., and J. G. March
- 1963 “A behavioral theory of the firm.” Englewood Cliffs, NJ, 2(4), 169-187.
- Delios, A., and W. J. Henisz
- 2003 “Political hazards, experience, and sequential entry strategies: The international expansion of Japanese firms, 1980–1998.” *Strategic Management Journal*, 24(11), 1153-1164.
- Denrell, J.
- 2003 “Vicarious learning, undersampling of failure, and the myths of management.” *Organization Science*, 14(3), 227-243.
- Dosi, G.
- 1982 “Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change.” *Research Policy*, 11(3), 147-162.

Fang, C., J. Lee, and M. A. Schilling

2010 “Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning.” *Organization Science*, 21(3), 625-642.

Feldman, M. S.

2000 “Organizational routines as a source of continuous change.” *Organization Science*, 11(6), 611-629.

Fleming, L.

2001 “Recombinant uncertainty in technological search.” *Management Science*, 47(1), 117-132.

Garud, R., R. L. Dunbar, and C. A. Bartel

2011 “Dealing with unusual experiences: A narrative perspective on organizational learning.” *Organization Science*, 22(3), 587-601.

Gelman, A., and G. King

1990 “Estimating incumbency advantage without bias.” *American Journal of Political Science*, 1142-1164.

Griliches, Z.

1991 “The search for R&D spillovers (No. w3768).” National Bureau of Economic Research.

Haleblian, J., and S. Finkelstein

1999 “The influence of organizational acquisition experience on acquisition performance: A behavioral learning perspective.” *Administrative Science Quarterly*, 44(1), 29-56.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg

2001 “The NBER patent citation data file: Lessons, insights and methodological tools (No. w8498).” National Bureau of Economic Research.

Hannan, M. T., and J. Freeman

1984 “Structural inertia and organizational change.” *American Sociological Review*, 149-164.

Henderson, R.

1993 “Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry.” *The RAND Journal of Economics*, 248-270.

Henderson, R. M., and K. B. Clark

1990 “Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms.” *Administrative Science Quarterly*, 9-30.

Holmqvist, M.

2004 “Experiential learning processes of exploitation and exploration within and between organizations: An empirical study of product development.” *Organization Science*, 15(1), 70-81.

Huber, G. P.

1991 “Organizational learning: The contributing processes and the literatures.” *Organization Science*, 2(1), 88-115.

Kaplan, S., and M. Tripsas

2008 “Thinking about technology: Applying a cognitive lens to technical change.” *Research Policy*, 37(5), 790-805.

Kaplan, S., and K. Vakili

2015 “The double-edged sword of recombination in breakthrough innovation.” *Strategic Management Journal*, 36(10), 1435-1457.

Katila, R., and G. Ahuja

2002 “Something old, something new: A longitudinal study of search behavior and new product introduction.” *Academy of Management Journal*, 45(6), 1183-1194.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy

2020 “Measuring technological innovation over the long run (No. w25266).” National Bureau of Economic Research.

Kim, J. Y., and A. S. Miner

2007 “Vicarious learning from the failures and near-failures of others: Evidence from the US commercial banking industry.” *Academy of Management Journal*, 50(3), 687-714.

Kuhn, T. S.

1962 “The structure of scientific revolutions.” University of Chicago press. Original edition.

Levitt, B., and J. G. March

1988 “Organizational learning.” *Annual Review of Sociology*, 14(1), 319-338.

Lewin, A. Y., S. Massini, and C. Peeters

2011 “Microfoundations of internal and external absorptive capacity routines.” *Organization Science*, 22(1), 81-98.

Lucas Jr, H. C., and J. M. Goh

2009 “Disruptive technology: How Kodak missed the digital photography revolution.” *The Journal of Strategic Information Systems*, 18(1), 46-55.

Madsen, P. M., and V. Desai

2010 “Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry.” *Academy of Management Journal*, 53(3), 451-476.

March, J. G.

1991 “Exploration and exploitation in organizational learning.” *Organization Science*, 2(1), 71-87.

March, J. G., and Z. Shapira

1987 “Managerial perspectives on risk and risk taking.” *Management Science*, 33(11), 1404-1418.

Mueller, D.C., and J. E. Tilton

1969 “Research and development costs as a barrier to entry.” *The Canadian Journal of Economics / Revue canadienne d'Economie*, 2 (4), 570-579.

Nelson, R. R.

1959 “The simple economics of basic scientific research.” *Journal of Political Economy*, 67(3), 297-306.

Nelson, R. R., and S. G. Winter

1982 “An evolutionary theory of economic change.” Cambridge, Mass. and London, Belknap Harvard.

O'Reilly III, C. A., and M. L. Tushman

2004 “The ambidextrous organization.” *Harvard Business Review*, 82(4), 74.

Park, S., S. Kim, and S. Choi

2018 “Knowledge Spill-Ins as Double-Edged Sword: How Knowledge Spill-Ins by Originating Firms Affect Entrepreneurial Recipients.” Working Paper.

Pugh, D. S., D. J. Hickson, and C. R. Hinings

1969 “An empirical taxonomy of structures of work organizations.” *Administrative Science Quarterly*, 115-126.

Rosenberg, N.

1976 “On technological expectations.” *The Economic Journal*, 86(343), 523-535.

Schumpeter, J. A.

1942 “Socialism, capitalism and democracy.” Harper and Brothers.

Simon, H. A.

1976 “From substantive to procedural rationality. In 25 years of economic theory” (pp. 65-86). Springer, Boston, MA.

Strumsky, D., J. Lobo, and S. Van der Leeuw

2012 “Using patent technology codes to study technological change.” *Economics of Innovation and New Technology*, 21(3), 267-286.

Taggar, S.

2002 “Individual creativity and group ability to utilize individual creative resources: A multilevel model.” *Academy of Management Journal*, 45(2), 315-330.

Teece, D. J.

1986 “Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy.” *The Transfer and Licensing of Know-How and Intellectual Property: Understanding the Multinational Enterprise in the Modern World*, 15, 67-88.

Teece, D. J., G. Pisano, and A. Shuen

1997 “Dynamic capabilities and strategic management.” *Strategic Management Journal*, 18(7), 509-533.

Walton, E. J.

2005 “The persistence of bureaucracy: A meta-analysis of Weber’s model of bureaucratic control.” *Organization Studies*, 26(4), 569-600.

Wang, C., J. Paisley, and D. Blei

2011 “Online variational inference for the hierarchical Dirichlet process.” In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (pp. 752-760).

Weber, M.

1958 “The Protestant Ethic and the Spirit of Capitalism, translated by T. Parsons.” New York: Charles Scribner’s Sons.

Weick, K. E., and D. P. Gilfillan

1971 “Fate of arbitrary traditions in a laboratory microculture.” *Journal of Personality and Social Psychology*, 17(2), 179.

Wilkerson, J., and A. Casas

2017 “Large-scale computerized text analysis in political science: Opportunities and challenges.” *Annual Review of Political Science*, 20, 529-544.

Table 1. LDA Patent Categories

| Category | Number of Patents |
|-----------------------|--------------------------|
| Digital Data | 752107 |
| Transponders | 427197 |
| Pharmaceuticals | 354498 |
| Mechanical Tools | 350235 |
| Photonics | 344345 |
| Semiconductors | 330224 |
| Structural Components | 323409 |
| Power Supplies | 296306 |
| Rotators | 289005 |
| Polymers & Ceramics | 283305 |
| Liquid Chemicals | 206373 |
| Medical Devices | 205897 |
| Electric Circuits | 197579 |
| Hydraulics | 192978 |
| Mechanical Structures | 189421 |
| Metallurgy | 186706 |
| Heating & Ventilation | 151513 |
| Organic Chemicals | 134055 |
| Vehicles | 132695 |
| Physical Media | 71563 |
| Plants & Animals | 68768 |

Table 2. H1 Variable Correlation Table

| # | Variable | Mean | Std Dev | (1) | (2) |
|---|-----------------------------|-------|---------|--------|-------|
| 1 | Prior art familiarity | 2.169 | 1.985 | | |
| 2 | First-time use of prior art | 2.001 | 1.082 | -0.105 | |
| 3 | Total prior art | 2.740 | 1.032 | 0.648 | 0.549 |

Table 3. H1: Probability of Patent Being a Breakthrough

| Variable | (1) | (2) | (3) | (4) | (5) | (6) |
|------------------------------------|--------------------------|------------------------|--------------------------|--------------------------|------------------------|--------------------------|
| Prior art familiarity | | -0.0156 * (0.00692) | -0.0346 *** (0.00979) | -.0829 *** (0.0156) | -0.138 *** (0.0396) | -0.164 *** (0.0312) |
| Prior art familiarity (squared) | | | 0.00391 ** (0.00147) | .00471 * (.00243) | 0.0157 ** (0.00592) | 0.00921 * (0.00507) |
| First-time use of prior art | 0.0405 *** (0.00833) | 0.0211 * (0.012) | 0.0285 * (0.0123) | .0374 * (0.0156) | 0.116 * (0.0506) | 0.0845 * (0.0344) |
| Total Prior Art | -0.0404 *** (0.00875) | -0.00979 (0.0162) | -0.0202 (0.0165) | -0.0166 (0.0234) | -0.0824 (0.0673) | -0.0441 (0.0495) |
| Intercept | 0.528 *** (0.0204) | 0.517 *** (0.021) | 0.537 *** (0.0221) | | 0.150 + (0.0890) | |
| Model Type | Linear, Pooled | Linear, Pooled | Linear, Pooled | Linear, Fixed Effects | Logistic, Pooled | Logistic, Conditional |
| Total Observations | 6576 | 6576 | 6576 | 6576 | 6576 | 6576 |

p < 0.05 = *, p < 0.01 = **, p < 0.001 = ***. One tailed p-values for hypothesized terms, two tailed p-values for control terms.

Table 4. H2-H3 Variable Correlation Table

| # | Variable | Mean | Std Dev | (1) | (2) | (3) | (4) | (5) | (6) |
|---|-----------------------------------|-------|---------|--------|--------|--------|--------|-------|--------|
| 1 | Exploitation by Inventing Firm | 3.828 | 12.92 | | | | | | |
| 2 | Firm Age | 3.658 | 1.233 | -0.032 | | | | | |
| 3 | Firm Sales | 9.836 | 2.099 | -0.094 | 0.607 | | | | |
| 4 | Total Exploitation | 3.625 | 1.400 | 0.539 | -0.121 | -0.169 | | | |
| 5 | Inventing Firm Total Patents | 3.118 | 3.673 | -0.084 | 0.392 | 0.539 | -0.121 | | |
| 6 | Breadth of Inventing Firm Patents | 25.79 | 17.66 | -0.096 | 0.421 | 0.639 | -0.070 | 0.376 | |
| 7 | Firm Performance | 0.13 | 0.136 | 0.045 | 0.179 | 0.337 | 0.020 | 0.008 | -0.023 |

Table 5. H2-H3: Exploitation by Inventing Firm: Negative Binomial Model

| Variable | (1) | (2) | (3) | (4) |
|---|--------------------------|--------------------------|--------------------------|------------------------|
| Firm Age | | 0.0754 ** (0.0245) | 0.0319 (0.0248) | |
| Firm Size | | 0.034 (0.0263) | 0.0810 ** (0.0263) | |
| Firm Age x Breakthrough Invention | | | 0.0871 * (0.0405) | 0.0823 ** (0.0351) |
| Firm Size x Breakthrough Invention | | | -0.0914 ** (0.0367) | -0.107 *** (0.0235) |
| Breakthrough Invention | 0.158 (0.119) | 0.159 (0.121) | 0.570 * (0.263) | 0.583 *** (0.149) |
| Total Exploitation | 0.855 *** (0.0182) | 0.864 *** (0.0183) | 0.865 *** (0.0182) | 0.887 *** (0.0285) |
| Inventing Firm Total Patents | -0.0433 *** (0.00821) | -0.0544 *** (0.00918) | -0.0597 *** (0.00929) | -0.0751 (0.0484) |
| Breadth of Inventing Firm Patents | .00742 ** (0.00243) | 0.00449 (0.00276) | 0.00248 (0.00276) | 0.00740 (0.0171) |
| Firm Performance | 0.389 (0.316) | 0.0748 (0.359) | -0.0712 (0.370) | -0.352 (0.382) |
| Inventing Firm Patents x Breakthrough Invention | 0.00541 (0.0114) | 0.00518 (0.0116) | 0.0159 (0.0132) | 0.0339 ** (0.0117) |
| Breadth of Inventing Firm Patents x Breakthrough Invention | -0.00244 (0.00263) | -0.00239 (0.00264) | 0.00132 (0.00348) | 0.00340 (0.00265) |
| Firm Performance x Breakthrough Invention | -0.0589 (0.519) | -0.0871 (0.530) | 0.195 (0.583) | 0.553 * (0.204) |
| Model Type | N. Binomial | N. Binomial | N. Binomial | Poisson |
| Fixed Effects | Industry | Industry | Industry | BT-Control |
| Total Observations | 6576 | 6576 | 6576 | 6576 |

Figure 1. Forward and Backward Citations for Breakthrough Patented Inventions

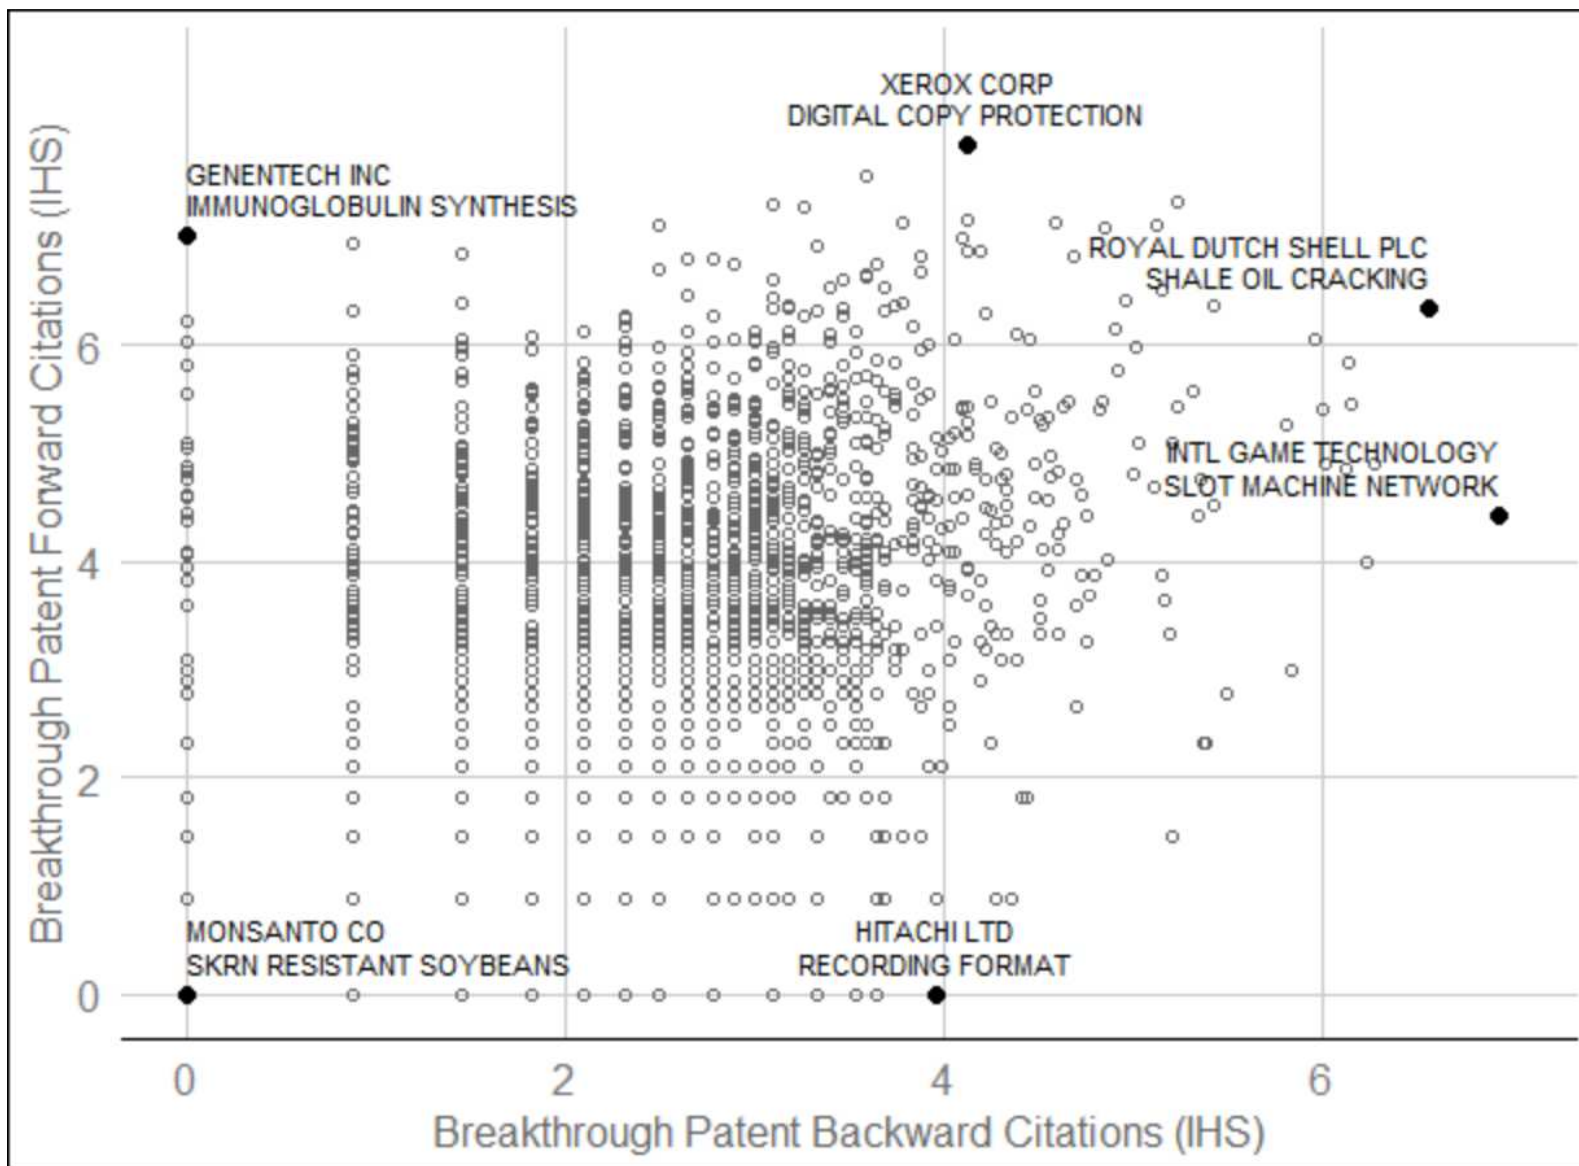


Figure 2. Text Categorization Process

Categorization

Individual documents (1969-2017) were sorted through multiple LDA and HDP testings, with the ultimate goal to narrow patents into more understandable topics.

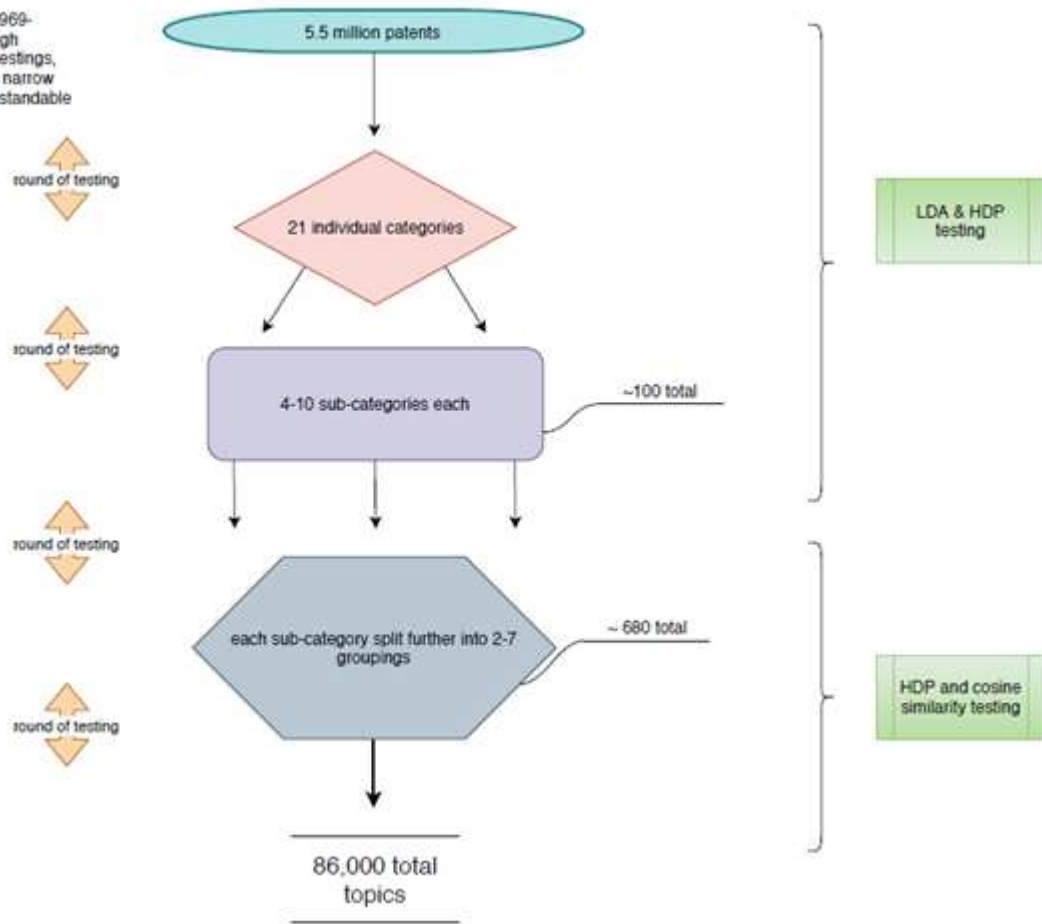


Figure 3. Effect of Familiarity with the Prior Art on Probability of Creating a Breakthrough

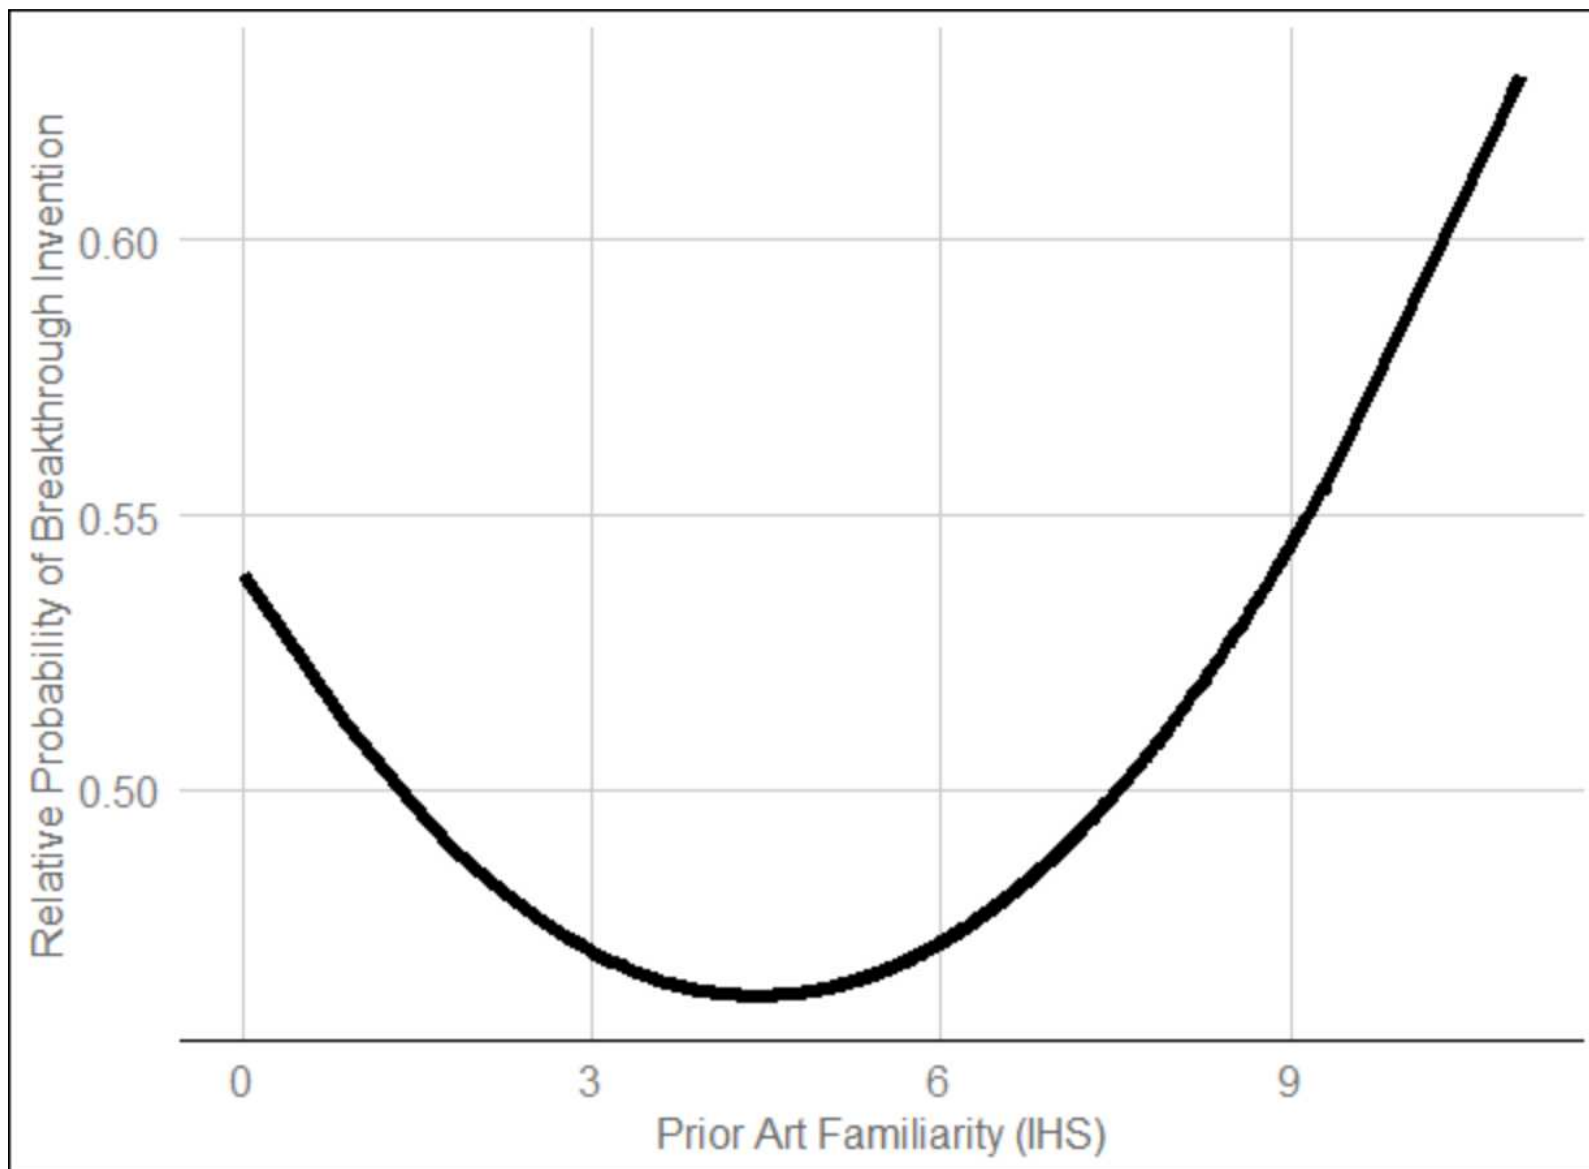


Figure 4: Effect of Firm Age on Exploitation by Inventing Firm

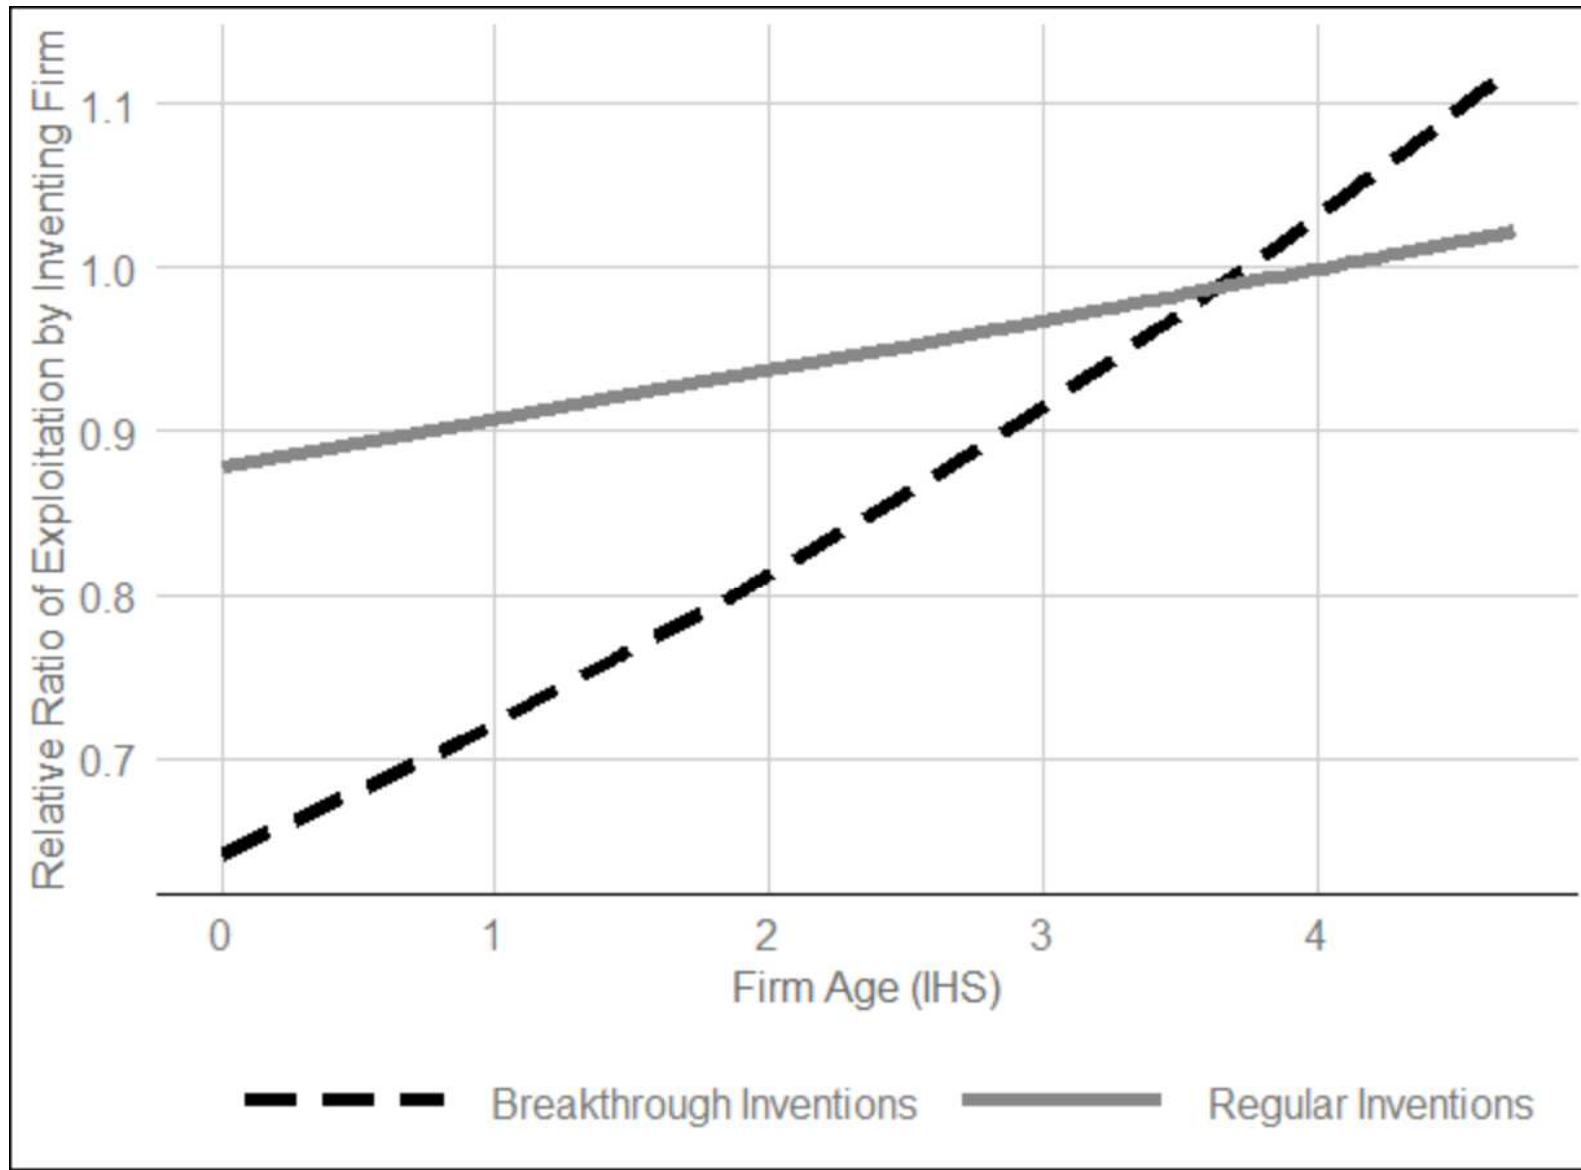


Figure 5: Effect of Firm Size on Exploitation by Inventing Firm

