Categories, Search, and the Influence of Patented Innovations

Abstract

In this paper, we argue that an innovation's influence is shaped by the social structures that it is embedded within; specifically, the categories that are used to organize items in a given domain. Our approach builds novel theory at the intersection of research on categories, search, and attention, and argues that an innovation's influence is related to: a) its position in a given category; b) that category's internal properties, and; c) the category's position within a broader classification system. We show that these factors help to explain an innovation's overall influence, as well as the degree to which this influence is broad versus narrow. Analysis is based on a quasi-experiment that compares the citations received by different members of the same patent family (i.e., patents for the same innovation in different patent systems). This controls for an innovation's features and allows us to isolate category effects. Results support our arguments, and show that an innovation's influence is shaped by its categorization, regardless of its underlying features.

In 2001, Honda Motor Company developed an innovation to increase the efficiency of fuel cell technology. The innovation was patented in both the United States (US 6916563) and Japan (JP 3904191). Three years later, the Japanese patent had been cited 17 times, while the U.S. patent had not been cited even once, despite being for the same exact innovation.

Similarly, in 2001, Nissan Motor Company developed an innovation to limit the harm of auto accidents to pedestrians. Again, the innovation was patented in the U.S. (US 6182782) and Japan (JP 3687418), and again, the two patents were very differently cited. In this case, the U.S. patent was cited 20 times, while the Japanese patent was cited twice.

As with other areas of creative endeavor, a patent's value is closely related to the influence that it has on subsequent innovations (Singh and Fleming, 2010). Studies have shown that citation counts are highly correlated with expert assessments of a patent's value and quality (Albert et al., 1991); that highly cited patents are more likely to be renewed by their holders (Harhoff et al., 1999; Guimmo, 2003); and that such patents create surplus value for both end-consumers (Trajtenberg, 1990) and the patent-holder (Hall, Jaffe, and Trajtenberg, 2005).

Given the importance of these outcomes, scholars have endeavored to understand why some patents are more influential, and thus valuable, than others. A key premise is that citations result from a process where actors search for items that are relevant to their efforts, and choose a select few to build-upon (Singh and Fleming, 2010). To date, most studies have approached this by showing that patents with certain features are more likely to be noticed and built-upon than others. Patents that draw on diverse knowledge are thought to attract attention because they stand-out as novel (Trajtenberg et al., 1997) and are relevant to multiple technology areas (Rosenkopf and Nerkar, 2001). Others have argued that patents that build on scientific knowledge are more likely to be noticed due to their quality and broad relevance (Podolny and Stuart, 1995; Gittleman and Kogut, 2003), and that actors who work in teams and within organizations are more likely to create novel, high-quality innovations than those who work alone (Singh and Fleming, 2010; Leahey, Beckman, and Stanko, 2017).

Yet these explanations cannot account for the examples we've raised. Why would patents for the same innovation be cited so differently in different patent systems? By definition, each builds on the same knowledge, has the same inventor(s), is equally novel, and is of the same quality. By holding patent-features constant while varying the context, we suggest that these examples point to a larger gap in the literature; namely, we know little about how a patent's influence is shaped by the broader social structures that it is embedded within. To address this, we argue it is important to recognize that innovations are often organized into category-based knowledge structures such as genres (Hsu, 2006), disciplines (Leahey, Beckman, and Stanko, 2017), and patent classes (Lo and Kennedy, 2014), and that this may affect what actors notice, and what they overlook, as they search for items to build-on in their own work.

Our theoretical approach is based on a novel integration of research on categories and the behavioral theory of the firm. Both assume that actors are boundedly rational, with limited attention spans and information processing capacities; yet they offer different explanations for how actors deal with these limitations. Behavioral theory focuses on cognitive processes, and argues that decisions about where to search for new information are guided by an actor's specific knowledge, goals, and schemas (Gavetti et al., 2012). Related work has shown that the density and visibility of alternate stimuli affects the likelihood that any given item will receive attention in such searches (Ocasio, 2011). Categories theory offers a complementary, socio-cognitive view. This work argues that actors rely on categories to group items in ways that distill the vast differences among stimuli into cognitively manageable chunks (Rosch and Lloyd, 1978). As a result, categories support shared understandings about the similarity and distinctiveness of different items, and thus reduce the cognitive burden of searching for and assessing potential alternatives (Zuckerman, 1999; Vergne and Wry, 2014).

Bridging these two perspectives, we suggest that search and attention are consequentially shaped by the categories (or classes) that are used to organize items in a given domain.[1] Specifically, we argue that, by grouping items together in meaningful ways, categories help actors determine where to search for items that are relevant to their efforts. Categories also create bounded areas to search within and, in doing so, define the features of this terrain. For example, a class may be variously crowded, and contain members that are more or less representative: both factors may affect the likelihood that particular items will stand-out and receive attention (Rosch, 1975; Hansen and Hass, 2001). Categories also exist within classification systems, and are variously similar to other classes therein (Leahey, Beckman,

[1] Note: we use the terms "category" and "class" interchangeably through this paper.

and Stanko, 2017). The contrast between categories may affect the likelihood that actors who are working in one class will search for items in another (Rao, Monin, and Durand, 2005).

In short, we argue that categories provide a socio-cognitive infrastructure that guides actors' decisions about where to search, and defines the terrain on which attentional processes operate (Ocasio, 2011). Based on this, we predict that an innovation's influence is shaped by: a) its position in a given category; b) that category's internal properties, and; c) the category's position within a broader classification system. We further predict that these factors will differently affect an innovation's influence within its own, versus other classes.

In developing our arguments, we discuss relevant insights from categories theory, link these to research on search and attention, and discuss how these dynamics might play-out in our specific context. Following the innovation literature, we recognize that search may focus in a single class, or expand to cover many (Kaplan and Vakili, 2015). When actors search in their "home" class (i.e., the category they are innovating within), we expect attention will focus on representative (or typical) items,[2] as these are likely to be the most relevant for innovating in that class (Tversky, 1977). Yet when search expands to other classes, we expect that attention will fall to atypical items, as these have features that may be germane to innovating in other domains (Rosch, 1975). We thus predict that typicality has a diametric effect: atypical items will be less influential in their own category, but are more likely to be cited by actors who are searching across-classes. We expect this pattern to be amplified in densely populated classes, as crowding leads actors to focus more intently on the items they expect will be most relevant to their efforts (Hansen and Hass, 2001). Patents in classes that are proximate to others within a classification system should also be subject to more cross-class searches, further benefiting atypical innovations. Overall, we suggest that a patent will be most influential—regardless of its technical features—when it is an atypical member of a category that is proximate to others in a classification system.

Analysis is based on a quasi-experiment that holds an innovation's features constant so as to isolate category effects. Specifically, we study patent families, which are patents for the same innovation issued in different jurisdictions (Gittelman and Kogut, 2003); in our case, Japan,

[2] Per the categories literature, we define "typicality" in relation to the category that an item sits within.

Germany, and United States (U.S.). Each nation uses a different patent classification system. As such, the same innovation may be more or less typical of the class it is placed into in each country, and these classes may be more or less crowded and proximate to others. Variance in citation norms can be addressed econometrically, allowing us to meaningfully compare the citations that the same innovation receives in each patent system. We enrich our arguments with data from interviews with 20 patent examiners, inventors, and attorneys representing each nation in our analysis.[3] Results support our predictions and have implications for the study of categories, innovation, and search.

THEORY AND HYPOTHESES

Categories and Search-Selection

To understand the conditions under which an innovation is likely to be noticed and built-upon, it is necessary to consider how actors search for new ideas to integrate into their own endeavors (Singh and Fleming, 2010). In this regard, search is the controlled process of attending to and evaluating stimuli, and is necessary because actors lack perfect knowledge and have a limited capacity to absorb information (Cohen and Levinthal, 1990; Gavetti and Levinthal, 2000). Actors cannot pay attention to all possible choice alternatives, and need to decide where to search for new knowledge. These choices are important, as the information that is available for an actor to notice and build-upon is determined by where they choose to search (Li et al., 2013).

To this end, studies rooted in the behavioral theory of the firm argue that search is guided by an actor’s specific goals and schemas, and that these create expectations about where to best search for relevant information (Gavetti et al., 2012). Most search takes place in the vicinity of an actor’s existing expertise (i.e. local search), as information here is easy to interpret and assimilate, and it is reasonable to expect that it will be relevant to one’s efforts (Levinthal and March, 1993; Podolny and Stuart, 1995). If local search fails to yield sufficient or satisfactory information, actors will begin to consider progressively more distant, less familiar areas of the search terrain (Cyert and March, 1963). In short, existing research characterizes search as

[3] Details about our sampling approach and interviewees are available in the online data appendix.

an actor-centric processes that takes place on a smooth, continuous terrain that is defined in relation to an actor's expertise.

This approach has proven very useful for studying innovation processes and outcomes at the actor- and organization-levels (Rosenkopf and Nerkar, 2001; Katila and Ahuja, 2002). Yet the potential for search to unfold in patterned ways across actors is largely neglected, creating blind spots in the theory's ability to predict which innovations will be more or less likely to be noticed and built upon.[4] There is precedent for a more socialized view, however.

Studies at the intersection of culture and cognition have argued that, in order to understand what actors pay attention to, it is important to view social structures and individual cognition as interpenetrating (Ocasio, 1997; Nigam and Ocasio, 2010; Ocasio, 2011). As Ocasio (1997: 193) notes, "cultural and institutional perspectives highlight the treatment of cognitive schemas, symbols, and systems of meaning as external and objective to individual decision-makers." For example, in their study of industry events, Hoffman and Ocasio (2001) showed that shared beliefs, as reflected in collective identities, led groups of actors to attend to similar environmental issues. Recent research on institutional logics (Thornton, Ocasio, and Lounsbury, 2012) has also argued that shared meaning systems within communities and fields channel attention in patterned ways among the actors therein (Marquis and Lounsbury, 2007; Pahnke, Katila, and Eisenhardt, 2015).

Applied to innovation, the implication is that decisions about where to search are unlikely to be wholly idiosyncratic. Broader social structures may shape the thinking of commonly embedded actors, such that they predictably focus their attention on similar types of items (Ocasio, 2011). In this spirit, we argue that search is likely guided by the category systems that are used to organize and order innovations in a given domain (Lo and Kennedy, 2014; Leahey, Beckman, and Stanko, 2017).

[4] This is not to say that the search and attention literature is not interested in how impactful knowledge is created. However, efforts to date have focused on the processes through which actors and organizations generate innovations with features that are associated with greater impact (e.g., greater novelty, higher quality, etc…).

As with the behavioral theory of the firm, categories theory assumes that humans are boundedly rational, have limited information-processing capacities, and must decide how to allocate attention among a potentially vast array of stimuli (Rosch and Lloyd, 1978; Vergne and Wry, 2014). Yet rather than focusing on actor-level cognitions, this research advances a socio-cognitive approach that emphasizes the role of shared knowledge structures. The key insight is that, while individual knowledge profiles may vary, categories shape cognition by lumping items into distinct and recognizable clusters. In so doing, categories support shared understandings about the similarity and distinctiveness of different items, and reduce the burden of searching for and assessing potential alternatives. Indeed, studies have shown that actors do not compare among infinite alternatives when making decisions. Rather, they rely on categories to provide bounded consideration-sets that comprise like items (Ruef and Patterson, 2009; Cattani, Porac, and Thomas, 2017).

From this perspective, search does not take place on a smooth and open terrain. Rather, it follows a process where actors pick a category to focus on, and compare items therein based on the expectation that these will be similar to each other and different from those in other classes (Zuckerman, 1999). Notably, such expectations are driven by category effects more so than the actual similarities among items. There are many ways that the same set of items can be grouped, and different attributes can form the basis for classification (Bowers and Prato, 2017). Still, a desire for cognitive economy motivates actors to view classes as distinct and as clear-cut as possible (Tversky, 1977). In turn, this creates a strong reliance on the category systems available in a given milieu to guide comparative processes. Illustrating this, Zhao (2005) showed that wine producers were accorded different status in France and the U.S. because these nations used different systems to classify wines. Bowers and Prato (2017) similarly showed that the status of equities analysts changed, regardless of their underlying quality, when the categories used to evaluate their performance changed. Speaking to the power of such processes, experiments have shown that subjects rely on categories to define consideration-sets even when the classes in question are artificially created (Rosch, Simpson, and Miller, 1976).

In the organizational literature, these insights have primarily been used to understand how an item's position within and across classes affect assessments of its value. Studies have shown that audiences judge organizations (Zuckerman, 1999), products (Hsu, 2006), and proposed innovations (Elsbach and Kramer, 2003; Boudreau et al., 2016; Criscuolo et al., 2017) by comparing them to others in the same category. A common finding is that moderately distinct items are viewed favorably, but those that do not fit neatly within a particular category are devalued because it is unclear how they should be evaluated, and they may be of lower quality (but see Wry, Lousnsbury, and Jennings, 2014; Paolella and Durand, 2016).

Our approach builds on the same foundations as this work, but adapts categories theory to explain search and attention in the context of innovation. As with previous research, we argue that innovators use categories to determine where to search for items that are relevant to their endeavors. Yet unlike competitive, market-based contexts where actors evaluate items to discern relative quality and identify the best alternative, innovators seek-out and cite multiple items that are relevant to their work. Thus, while influence is a key indicator of a patent's value (Harhoff et al., 1999; Guimmo, 2003), we argue that this reflects different processes than are discussed in the extant categories research. To be cited, an innovation does not need to be evaluated superior to alternatives; it needs to be noticed and perceived as relevant.

Moreover, actors may look for relevant items within a single class, or expand their search to cover many (Kaplan and Vakili, 2015). So, while our approach suggests that categories parse the terrain on which search takes place, we embrace the potential for this process to unfold across multiple categories. Akin to arguments about local search, we expect that actors will begin by searching the category that best aligns with the innovation they are pursuing. If this fails to yield sufficient information, search will likely expand to other classes. However, we argue that actors will attend to systematically different types of items when searching within- versus across-classes, and that this will create patterned variance in the likelihood of different innovations becoming influential.

Patent Classes and the Search for Prior Art

We develop our arguments in the context of patented innovations. In each patent system worldwide—including the three in our analysis—innovations are organized according to category systems that group items based on predefined sets of attributes. The specific classes in each system differ but, in each case, their primary function is to enable inventors and patent examiners to search for previous patents (i.e., prior art) that are relevant to a focal innovation (Jaffe, Trajtenberg, and Henderson, 1993; Bacciocchi and Montobbio, 2009).

Consistent with the insight that actors "acquire membership in a given discipline or social group by becoming comfortable with the language, categories, and things used by the group" (Star, 2005: 174), inventors are advised to acquaint themselves with the patent classification system in their nation, and to use it as a resource when developing their ideas. To this end, guidelines published by the German, Japanese, and U.S. patent offices suggest that inventors should perform searches early in the innovation process where they: 1) enter keywords that describe their idea into a patent search tool; 2) scan the returned patents to determine which patent class best captures their particular idea; 3) work through that class to determine which patents are relevant to their efforts, and; 4) expand their search to other potentially relevant classes (see Bellis, 2014; JPO, 2017; USPTO, 2017).

While it is difficult to observe the degree to which this advice is followed, Maggitti et al., (2013) have noted that many innovators scan the patent landscape routinely to update their knowledge in the areas where they are active. Fleming (2001) has suggested that this becomes more evident when an actor becomes serious about filing a patent application. The informants that we interviewed made similar claims. For example, a German inventor told us:

> *I do basic keyword searches to confirm where to search…you start with what you know and then dive into that class.*

Expanding on this theme, an American patent attorney noted:

> *You would get lost if you had to look at everything. You need to search one category at a time if you want to find relevant [prior] art. I tell inventors to start with the class that has innovations that align with what they're working on.*

The same basic process is formalized when an application is received by a national patent office. Before a patent is issued (or not), it is reviewed by an examiner. As with other expert evaluators, patent examiners have expertise that is tailored to a specific area, typically in the form of an advanced degree supplemented by specialized on-the-job training (USPTO, 2017). Yet, unlike evaluators who make comparative quality assessments of different items, a patent examiner's job is to assess whether or not an innovation is sufficiently novel and non-obvious to warrant legal protection. This involves searching for relevant prior art, and adding citations to a patent application. Based on their expertise, examiners are assigned to "art units" that are responsible for reviewing applications that fit specific patent classes (Alcacer and Gittelman, 2006; Yanagisawa, 2016). When an examiner receives an application, they establish a "field of search" that guides their effort to identify related prior art. The focal class is prioritized as the most relevant, and other classes are added when the examiner suspects they may contain patents that are relevant to an application (USPTO, 2012; Yanagisowa, 2016). Speaking to this, a member of the Office of the Chief Economist at the U.S. Patent Office told us:

> *The corpus of information that [examiners] have on patented prior art is just humungous. Their search is dictated by classes... They use the primary classification to start the search [and then] branch out into other classes as necessary.*

A patent attorney with knowledge of each patent system in our analysis similarly reported:

> *In Japan and Germany…the process is the same. You start with keywords, find the class that aligns with what you're doing, and then start digging…if you find there are claims [in the focal innovation] that aren't covered by the main class, you start looking in other classes.*

Reflecting the similarity in the process that inventors and examiners use to search for prior art, there tends to be a high correlation in the types of patents that these actors cite (Alcacer and Gittelman, 2006). Also, while citations are not a perfect measure, studies have shown that they meaningfully reflect the types of knowledge an innovation builds-upon (Jaffe and Trajtenberg, 1996; Jaffe, Trajtenberg, and Fogarty, 2000). Indeed, there is a well-established precedent in the literature for using prior art citations to proxy for search results (e.g., Stuart and Podolny, 1996; Rosenkopf and Nerkar, 2001; Katila, 2002; Katila and Ahuja, 2002).

Within-Category Influence

Consistent with the theoretical perspective that we are advancing, we expect that search will begin in the category that most closely aligns with an actor's current efforts and expertise.

Moreover, categories theory suggests that certain items are systematically more likely to stand out on this search terrain (Vergne and Wry, 2014). We extend these insights to predict which innovations are the most likely to be noticed and built-on by others in the same class.

Once established, categories not only create delimited sets, but also structure these items in ways that are apparent to knowledgeable actors (Rosch and Lloyd, 1978). In this regard, a key finding is that most categories comprise a structure where items are arrayed based on their typicality (i.e., the extent to which they have features in common with members of the same class, and few in common with those in other classes) (Rosch, 1975; Tversky, 1977). For instance, most people recognize a Robin or Blue Jay as a typical member of the "bird" category, whereas chickens, penguins, and ostriches are variously atypical (Rosch, 1975).

In turn, this internal structuring affects which items in a class are the most likely to receive attention. Typical items have high cue validity and are viewed as the clearest examples of a category (Rosch 1973, 1975; Tversky, 1977). By virtue of these features, experiments have shown that actors rely on typical items as referents in search and comparison tasks, and recall them more quickly and easily than other category members (Rosch, 1975; Rips, 1975). Typical items are also routinely used as comparators that help actors to determine how unfamiliar objects should be classified (Rosch, 1975).

In organizational contexts, the relationship between typicality and attention can be seen in studies that have shown that audiences evaluate firms and products in relation to those that exemplify a class's defining features (Elsbach and Kramer, 2003; Rao, Monin, and Durand, 2003; Bowers, 2015). Related work has found that typical items are evaluated relatively quickly and favorably, while atypical items are more likely to be overlooked (Zuckerman, 1999; Hsu, 2006; Lo and Kennedy, 2014). Entrepreneurs also tend to describe their firms in relation to the most typical members of a market category, and communicate this in the stories they tell to resource providers (Lounsbury and Glynn, 2001; Martens, Jennings, and Jennings, 2007; Wry, Lounsbury, and Glynn, 2011).

Building on this, we expect that within-class searches will begin with a focus on typical items. Indeed, when an actor searches for items that are relevant for innovating in a category, it makes sense to start with those that have high cue validity, exemplify the class, and are readily available for recall (Rosch, 1975). Pragmatically, it also stands to reason that actors will focus their attention on items that are, by definition, the most likely to be relevant to other innovations in the same category (Tversky, 1977). Subsequently, search will likely progress to consider less typical items, as these may be germane to particular lines of advance within a class, even if they are less generally relevant. Given that items can be atypical in a variety of ways, though, not all are likely to be relevant to an actor's particular efforts. As such, we expect that within-class searches will not be selective, rather than exhaustive.

Our theoretical argument is consistent with the responses we received from our informants. For instance, a former German patent examiner told us:

> *Even in a class that you know well, there are a lot of patents to search…you want to focus on the most typical I.P. [intellectual property] first because this is [likely to be] the most relevant.*

This sentiment was echoed by an American patent attorney who told us:

> *It's hard to be exhaustive because there's so much art in these classes. My approach has always been to find something that's really representative of the class, and then follow the breadcrumbs to find other things that are relevant.*

An economist at the United States Patent Office offered a similar account, while also suggesting that a focus on typical innovations becomes reinforced when an actor searches the same category repeatedly:

> *You look in the usual spots first and then go out from there…Examiners have a sense of what the most relevant bits of prior art are in their [home class]. They come across these quite a lot because they're relevant to so many other patents that hit that class. This is where you want to start…You get very lost if you go broad first.*

Thus, while some degree of atypicality may be beneficial for an innovation in the context of competitive evaluation processes (Lounsbury and Glynn, 2001; Boudreau et al., 2016; Criscuolo et al., 2017), we expect that the overall effect of atypicality will be negative as it relates to within-category influence. As such, we predict:

Hypothesis 1 (H1): Atypicality is negatively related to an innovation's within-category influence.

In addition to typicality effects, there is variance in the number of items that populate different categories. As with previous studies of search and attention, we expect that this has implications for an innovation's within-class influence. Yet while this work suggests that crowding makes it less likely that individual items will stand-out on the search terrain, as there are many items competing for attention (Hansen and Haas, 2001; Sullivan, 2010; Piezunka and Dahlander, 2015), we expect that this will be offset in contexts like ours where crowding also implies that a large audience is searching within a class.

Studies have shown that, even in situations where items are evaluated against each other, there are benefits to being in a class that has a large audience. For instance, movies benefit from being in popular genres (Hsu, 2006), as do products in popular eBay categories (Hsu, Hannan, and Koçak, 2009) and startups in "hot" technology areas (Wry, Lounsbury, and Jennings, 2014; Grodal, 2017). We expect that an innovation will similarly benefit when there is a large audience searching within its class (Podolny and Stuart, 1995).

Concerns about crowding are further attenuated in our context by the fact that actors search for multiple items that are relevant to their own efforts. Indeed, there are incentives to search persistently through a densely populated class where one expects to find relevant items: identifying such innovations can help to advance one's own efforts, while their neglect may lead to sanctions.[5] This effort should be further enabled to the extent that actors have domain-specific knowledge that eases the burden of considering a large number of offerings in the class where they are innovating (Zuckerman, 1999; Negro and Leung, 2013). As such, we expect that within-class citations will be higher when a patent is situated in a densely populated class (Podonly and Stuart, 1995).

[5] For example, a failure to cite all relevant prior art in a patent application can result in the invalidation of claims in the face of post-issuance lawsuits (Allison and Lemley 1998). Scholarly research is also less likely to be published (Pfeffer, 1993) or impactful (Mukherjee et al., 2017) if it neglects the relevant literature.

That said, we expect that being in a densely populated class will be especially beneficial for typical category members. While actors have incentives and ability to consider many items in their "home" class, there is a consistent finding that actors narrow their attention when faced with a crowded search terrain, or other decision-making pressures (Payne, Bettman, and Luce, 1996). Organizational research has studied this at the firm-level, examining attention to items such as internal documents (Hansen and Haas, 2001), problems (Haas, Criscuolo, and George, 2015), and externally solicited suggestions (Piezunka and Dahlander, 2015). There is a consistent finding that crowding leads actors to focus on the items that they expect will be most relevant and useful to their efforts. We believe that similar dynamics apply to within-category searches. Given that typical items are the most likely to be relevant to others in the same class, we anticipate that actors will narrow their attention and rely more heavily on these items, as compared to atypical ones, when a class is crowded.

Again, interview data complement our arguments. For example, a U.S. patent attorney noted:

> *Even when it's [a class] you're familiar with, you really have to rely on the salient stuff that you know is likely to be relevant and [that you] have experience with...It's hard to identify what [all is] relevant in a big, crowded space.*

This was echoed by a former Japanese patent examiner, who told us:

> *It can be hard to spot what's relevant in a big technology area…realistically, there's so much stuff that you can't look at everything... You need to rely on quick and dirty tools and focus on the usual suspects.*

Based on this, we predict:

Hypothesis 2a (H2a): Category density is positively related to an innovation's within-category influence.

Hypothesis 2b (H2b): The positive relationship between category density and within-category influence is weaker for atypical innovations.

Across-category Influence

Innovations may also serve as building blocks for development outside of their home class (Rosenkopf and Nerkar, 2001; Kaplan and Vakili, 2015). To understand overall influence, it is thus necessary to account for the conditions under which an innovation is likely to be noticed and built-upon by those who are innovating in other categories. However, while being a typical member of a densely populated class may be conducive to within-category influence, we expect that these same factors make it less likely that an innovation will be noticed and built-on by actors who are searching across category boundaries.

In addition to their grouping function, categories distinguish between different types of items (Ruef and Patterson, 2009). Cognitive efficiencies are gained because categories convey not only within-class similarities, but also between-class differences (Rosch and Lloyd, 1978; Zuckerman, 1999). Further, while there may be overlaps in some situations, categories tend to be associated with different knowledge and expertise (Negro and Leung, 2013), making it likely that an actor will encounter cognitively distant information when searching across categories (Cohen and Levinthal, 1990; Stuart and Podolny, 1996). Based on this, we argue that actors will focus on atypical items when searching across category boundaries. It is rational for an actor to expect that atypical items will be the most relevant for innovating in a different class; by definition, these items have features in common with other categories (Rosch, 1975). Atypical items are also likely to stand-out relative to others in a class, even among actors who lack the knowledge to discern nuanced differences between category members (Kovács and Hannan, 2010). Taken together, this suggest that atypical innovations are the most likely to be noticed and seen as relevant in cross-category searches. Supporting this, Li and colleagues (2013) found that managers were more likely to attend to novel and distinctive information when they engaged in distant search. This is also consistent with the finding that innovations that blend insights from multiple perspectives often receive broad-based attention within scientific and technological communities (Rosenkopf and Nerkar, 2001; Leahey, Beckman, and Stanko, 2017).

Likewise, our informants reported that they expect atypical patents to be the most relevant to their own cross-category searches. As one German patent attorney reported:

> *As you broaden the search, you're looking for those areas where another class might overlap a little bit with relevant [prior art] in your own area of expertise.*

This was echoed by a former Japanese patent examiner who told us:

> *When you're looking at more tangential classes, you cross those [with your home class] to get a narrower pool of patents… [you] look for art in other classes that stands out as relevant for what [you're] doing. This tends to be peripheral stuff.*

In the American context, a former patent examiner similarly noted that:

> *A human can't focus on everything when search expands beyond a focal class. A conscientious examiner tries to take everything into account, but there's just so much art you can't look at it all. It requires some real effort. You start at the margins of the [other classes], and find the art that covers a portion of what you're looking for.*

It is also unlikely that cross-category searches will proceed significantly beyond the category-periphery, further strengthening the attentional advantages of atypical items. Indeed, to the extent that typical items are perceived as maximally different from the members of other classes (Rosch, 1975), it is unlikely that actors will see much value in considering these when they search across classes. Further, searching across classes puts an actor on terrain where they are more likely to encounter information that is challenging to understand and assimilate (Cohen and Levinthal, 1990; Stuart and Podolny, 1996). In such situations, more time is needed to apprehend the nature of different items, increasing the cognitive burden of search without increasing the expected payoff. We expect that this will lead to satisficing behavior, where an actor stops searching after they find a few satisfactory items (Cyert and March, 1963). Again, this aligns with our interviews. Per one American inventor:

> *[You] want to know what generally is out there… [but I'm] not interested in being totally exhaustive when I get outside of the narrow search area [and start searching in other classes]...You can't cover everything.*

Likewise, two members of the Office of the Chief Economist at the United States Patent Office respectively reported to us:

> *When you get a few hits outside the focal class, you don't need to chase down a bunch [more] leads. Satisficing is fine.*
>
> *Sometimes you'll need to get into different areas…but people aren't searching these classes nearly as intensively. It's just not worth spending that much time to sift through the noise.*

Based on this, we predict:

Hypothesis 3 (H3): Atypicality is positively related to an innovation's across-category influence.

Density may also affect the likelihood that items will be noticed by actors who are searching across classes. As a baseline, we expect that being in a densely populated class will make it less likely for an innovation to be noticed. Attention constraints are less relevant when the search terrain contains fewer items, and offerings are more likely to be noticed when there are fewer others competing for attention (Sullivan, 2010). By corollary, it is tough to stand-out in crowded spaces (Hansen and Haas, 2001), and actors who are not motivated to search intensively often overlook items (Levitt and March, 1988). This is especially true when search takes place on unfamiliar terrain, or when actors are motivated to satisfice (Piezunka and Dahlander, 2015). Further, unlike within-category search, density does not imply that a class will be subject to a greater number of across-class searches. As an American patent examiner told us: "big class, small class, it doesn't matter; you search where you think you're going to find relevant art."

Still, to the extent that crowding leads actors to rely more heavily on their expectations about which items are the most likely to be relevant to their efforts (Hansen and Haas, 2001), the negative effects of crowding may be attenuated for atypical items. Indeed, as Piezunka and Dahlander (2015) found, actors respond to large amounts of distant information by narrowing their attention, applying strong filters, and considering only a small subset of ideas. Speaking to this, an American patent attorney that we interviewed reported:

> *There are more and less crowded spaces [and it gets] really hard to identify what's relevant when you're searching a big space that you don't know well…you need to really zero-in those pieces of art that you think are going to be relevant.*

Based on this, we predict:

Hypothesis 4a (H4a): Category density is negatively related to an innovation's across-category influence.

Hypothesis 4b (H4b): The negative relationship between across-category influence and density is weaker for atypical innovations.

A Category's Position Relative to others in a Classification System

Our final hypotheses predict that an innovation's across-class influence will vary based on the position that its category occupies relative to others within a classification system.

Studies have long recognized that categories are not isolated, but rather sit within broader knowledge structures (Vergne and Wry, 2014). Early studies showed that category systems are taxonomic, and comprise hierarchically nested classes: high-level categories are more inclusive, while lower-level classes make finer-grained distinctions among items (Rosch et al., 1976). A key insight is that actors rely on "basic-level" classes (i.e., those that maximize in-group similarities and also convey meaningful out-group differences) to guide cognition and decision-making (Porac and Thomas, 1990). Recent work has extended this to consider the relationships between basic-level classes, such as those that we focus on in this study. For instance, Porac and Thomas, 1990) found that patent classes for different nanotechnology inventions became more or less similar over time, thus affecting the perceived coherence of startup firms' patent portfolios. Similarly, Leahey and colleagues (2017) modelled the knowledge-distance between academic disciplines, and showed that collaborations were more likely among scholars in adjacent, versus distant domains. In short, while all basic-level categories distinguish between items, these classes are not equally different from each other.

This is noteworthy because, even if search extends beyond a focal class, it is not infinitely broad (Monahan, 1982). As search crosses greater distance, cognitive challenges increase and success prospects decline (Cohen and Levinthal, 1990). It is also more straight-forward to see how items in different categories relate to each other when there are underlying similarities or symbioses between the classes (Wry, Lounsbury, and Jennings, 2014; Paolella and Durand, 2016). Indeed, there is evidence that actors begin with proximate domains when searching across categories, and only rarely engage ideas from far-flung areas (Leahey, Beckman, and Stanko, 2017; Mukherjee et al., 2017). For instance, psychologists and sociologists often look for ideas in the other's literature, but bibliometric analyses show they are unlikely to expand their searches into distant domains like engineering and botany (Jacobs and Frickel, 2009).

Taken together, this suggests that, while within- and across-class searches likely differ, there may also be substantial variance in the latter. Categories that are closer to others should fall within the potential radius of a greater number of cross-class searches. We thus expect that, on average, items in such categories will receive more attention than those in more distant classes, thus producing a general lift in across-category influence. Following the logic of our previous hypotheses, we also expect that this increased search activity will be most beneficial for innovations that stand-out within a category. To wit, the effect of atypicality on across-category influence should be amplified when a category is more proximate to others.

Finally, while our theory thus far implies tradeoffs between within- versus across-category influence, it does not predict overall influence. Typicality and category density may affect the distribution of citations across- versus within-classes without changing an item's overall influence. When a category is subject to more across-class searches, however, the associated rise in citations should result in innovations becoming more influential, overall. This should be particularly evident among atypical items, as these are the most likely to receive attention from actors who are engaged in across-class search. Based on this, we predict:

Hypothesis 5 (H5): A category's proximity to others in a classification system is positively related to an innovation's (a) across-category and (b) overall influence.

Hypothesis 6 (H6): The positive relationship between a category's proximity to others in a classification system and an innovation's (a) across-category and (b) overall influence is stronger for atypical innovations.

DATA AND METHODS

Empirical Context

We test our hypotheses by examining the forward citations (i.e., the number of times a patent is cited by subsequent innovations) for patent families where the same innovation is patented in each of the U.S., Germany, and Japan. Classification and search follow similar principles in each nation, but there are differences in the categories that comprise each system and, of course, in the specific patents issued in each (Jaffe, Trajtenberg, and Henderson, 1993; Goto and Motohashi, 2007). As such, the same innovation may be more or less typical of the class it is placed into in different nations, and the category itself may be variously crowded and proximate to others. We leverage this variance to analyze differences in the citations received by the same exact innovation in different patent systems. This effectively holds constant factors such as patent quality, team composition, number of claims, inventor characteristics, knowledge inputs, and others that have been shown to affect citation rates (Rosenkopf and Nerkar, 2001; Gittelman and Kogut, 2003; Singh and Fleming, 2010; Boudreau et al., 2015), thus allowing us to isolate category effects.

Per our theory, each country's patent office organizes inventions using a category system that groups items together, and distinguishes them from others based on particular attributes and functions (Goto and Motohashi, 2007). Still, classifications emerge and ossify within unique cultural and historical contexts (Bowker and Star, 1999), and this is reflected in differences among national patent systems. Based on their varied resource endowments, cultures, and institutional frameworks (Cantwell, 1989; Bartholomew, 1997), the U.S., Germany, and Japan each developed their own classification system.[6] Despite efforts to harmonize the systems, there are persistent differences in the classes that comprise each, and in the patterns of relationship among these classes (Goto and Motohashi, 2007; Bacchiocchi and Montobbio, 2009). Each national patent office conducts its own examination of a patent application: inventors can apply for intellectual property protection simultaneously across jurisdictions, but there is no such thing as an "international patent."

Despite these differences, the primary reason that patent classes exist in each system is to guide search. There are multiple, hierarchical levels of classification within each system. Yet as with other category systems, search and comparison are primarily guided by basic-level categories that define groups in ways that are not too broad, nor too specific (Rosch and

[6] Germany and Japan use differently modified versions of the International Patent Classification (IPC) system, while the U.S. uses a completely different system.

Lloyd, 1978). In the U.S., this is 3-digit technology classes (Hall, Jaffe, Trajtenberg, 2001); in Germany and Japan, it is 4-digit IPC classes (Bacchiocchi and Montobbio, 2009).

Identifying Patent Families

We identified patent families using the European Patent Office's Autumn 2014 edition of the Worldwide Patent Statistical Database (i.e., PATSTAT). PATSTAT is a comprehensive database of patent applications and grants from over 100 countries, and includes information on more than 35 million granted patents. From this, we collected all patent applications and granted patents in the U.S., Germany, and Japan from 1995 to 2013. We ended our data collection at 2013 to avoid potential changes that may have arisen when national patent offices began to implement the Cooperative Patent Classification (CPC) system.[7] Our analysis is thus limited to patents granted on or before 2010 so as to observe a full 3-year window of forward citations for all patents in the sample.

An advantage of using PATSTAT is that it tracks patent families. Each family represents the set of documents (i.e., applications and grants) that are issued in different nations (i.e., different patent jurisdictions) to protect a single innovation. Families are identified by linking patents to a common first application, called a priority, which establishes the date when intellectual property protection begins. The Paris Convention for the Protection of Industrial Property allows twelve months to file applications in other member nations while claiming the same priority date by linking to the initial priority. In some cases, however, applications may be linked to multiple priorities. This can be due to the timing of the applications (e.g., simultaneous applications are made in two nations followed by a later application in a third), or when applications make different claims. Families can thus be defined in different ways, and may comprise patents that are more or less closely related to each other. Our analysis uses the Espacenet approach, which is the most restrictive, and thus conservative way to define families. Espacenet requires that all documents (i.e., applications and grants) are

[7] Beginning in 2010 efforts were initiated to harmonize across systems using a newly developed CPC system, which is a modified version of the IPC system. The CPC was implemented alongside the U.S., German, and Japanese patent classification systems in 2013.

linked to the exact same priority, or set of priorities. In total, 7,288 patent families in our analysis window spanned the U.S., Germany, and Japan and met this criterion.

Dependent Variables

Our dependent variables track the overall examiner-added citations received by the same innovation in the U.S., Germany, and Japan, as well as the extent to which this comprises within- versus across-class citations. As with previous studies, we use a patent's 3-digit U.S. class and its 4-digit German and Japanese classes to distinguish between within- and across-class citations (e.g., Trajtenberg, Henderson, and Jaffe, 1997; Alcacer and Gittelman, 2006). In each system, a patent has a single mandatory original classification that corresponds to its "controlling" (or most prominent) claim (USPTO, 2012). When a patent is cited by another with the same original class, we code this as a within-class citation. Citations from patents with a different original class are across-category citations.

Regardless of the system being analyzed, citation data is a common proxy for search behavior (Rosenkopf and Nerkar, 2001; Bacchiocchi and Montobbio, 2009), and highly cited patents are more influential, and have greater economic value to their owners (Trajtenberg, 1990; Albert et al., 1991; Harhoff et al., 1999; Guimmo, 2003; Hall, Jaffe, and Trajtenberg, 2005). As with previous studies, we examine the three-year count of forward citations for each patent in each system (Kaplan and Vakili, 2015). Results are consistent when using a five-year window, but this reduces the number of observations in our sample.

We took a number of steps to ensure that our variables captured meaningful cross-system comparisons. First, we corrected for variation in citation norms across systems. One way we do this is by including citations to both the granted patent and its associated application. This is important, as citation practices vary somewhat across systems on this particular dimension. For instance, U.S. examiners primarily cite granted patents (in our PATSTAT data, 91% of U.S. backward citations are to granted patents). Yet examiners in Germany and Japan cite yet-to-be-granted patent applications more frequently (59% of citations in Germany and 98% in Japan are to patent applications). By including both types of cites, we capture all citation information and create a comparable measure across the nations in our sample. We exclude cross-system citations (i.e., those from patents in other systems). For example, when counting citations to a Japanese patent, we only include citations from other Japanese patents and do not include citations from German or U.S. patents. We do this to isolate cross-system variance, in keeping with our theoretical arguments and identification strategy. As we discuss in the Robustness Checks section, results are similar when we include cross-system citations.

Prior research has also shown that average citation rates differ across nations (Yasukawa and Kano, 2015). We thus need to account for systemic differences in citation frequencies over time and across patent systems. One approach, as suggested by Henderson and colleagues (2005), is to divide each patent's citation count by the mean citations of its cohort (i.e., all patents from the same patent system year). Yet while this standardizes citation means across patent cohorts, it doesn't address citation heterogeneity in cohorts across systems and over time. Thus, we instead used citation z-scores, calculated by subtracting the cohort mean from the patent's citation count, and then dividing by the cohort standard deviation. For example, the mean of all examiner-added citations within three years of publication for all U.S. patents granted in 2005 is 0.50, and the standard deviation is 1.10. A patent from this cohort with 2 forward citations would have a z-score of 1.364 (i.e., (2-0.50)/1.10). Using z-scores instead of raw counts creates variables that are comparable across systems over time. Our variables are thus standardized counts of all citations; within-class citations, and; across-class citations. Positive (negative) values indicate that a focal patent has received more (fewer) citations as compared to all other patents granted in the same patent-system-year.

Independent Variables

Our first independent variable, atypicality, considers the extent to which an innovation is typical of the category that is placed within. Per the categories literature, we assess this based on the extent to which an innovation's attributes fit with a single category (Rosch, 1975; Tversky, 1977). Following Lo and Kennedy (2014), our measure is based on the "secondary" classes listed on a patent. While a patent's original class reflects its most prominent inventive features, secondary classes are used to account for other claims that are not covered by the original class. These may be sub-classes of the original class, or different classes altogether.

Our specific measure integrates two approaches that are commonly used to assess a patent's fit with a given technology class. One is the Herfindahl measure suggested by Hall and colleagues (2001) that captures the concentration of classes listed on a patent. The other is the measure suggested by Trajtenberg and colleagues (1997) that captures the distance of the classes being spanned.[8] As with Leahey, Beckman, and Stanko (2017), we incorporate both measures using Porter and colleagues (2007) inter-disciplinarity measure. This provides a robust assessment of category fit by accounting for both the portion of secondary classes that are different from the original class, as well as the distance between these secondary classes and the original class. The measure is calculated in two steps.

First, we represent the technology landscape and the relationships among patent classes as a weighted network. In network terminology, nodes are patent classes and the edges connecting these nodes signify the degree to which connected classes are technologically related. We calculate the degree of relatedness, or edge weight, of two focal patent classes as a proportion of the relative frequency by which the two classes are listed simultaneously on granted patents. The degree of relatedness, P_{ij}, for technology classes *i* and *j* is calculated as:

$$P_{ij} = \frac{1}{2}\left(\frac{c_{ij}}{c_i} + \frac{c_{ij}}{c_j}\right)$$

Where C_{ij} is a count of the number of times patent classes *i* and *j* are co-listed on a patent in the prior year. C_i and C_j are the counts of the number of patents that classes *i* and *j* are listed on in the prior year, respectively. By construction, the relatedness of any two classes falls somewhere on the unit interval. This is the same measure used by Lo and Kennedy (2014).

Second, we calculate atypicality by applying Porter and colleagues' (2007) formula to patents and patent technology classes. The formula for a given patent is:

$$Atypicality = 1 - \sum_{for\ all\ i,j} P_{ij} * s_i * s_j$$

[8] These measures use the backward citations listed on a patent, as opposed to secondary classifications. Still, the two approaches are conceptually similar, as they both reflect the areas of innovation that a focal invention builds upon. We use the latter in our analysis because secondary classification practices are much more comparable across countries than are citation practices, resulting in a more reliable measure for cross-country comparisons.

For each technology class *i* and *j*, where P*ij* is the degree of relatedness between technology classes *i* and *j*, s_i and s_j are the percent of secondary classes listed on the given patent from technology classes *i* and *j*, respectively. Possible values lie on the unit interval. Larger values indicate that a patent has many secondary classes that are not related to the primary class.

Our second independent variable is category density. We operationalize this by counting the number of patents granted in the same class the year prior to a patent's publication. We standardized this measure by converting raw counts into z-scores because the number of granted patents varies across patent systems (i.e., many more patents are granted in the U.S. versus Germany or Japan) and over time (i.e., the number of patents has generally followed an increasing trend in all three systems). As density is a patent class-level construct, though, we use the mean and standard deviation for all of the classes in a given system-year instead of using patent cohorts: this yields a measure that is comparable across systems and over time. As noted in our Robustness Checks, we also ran models that used different lags to differentiate between category popularity (i.e., the number of subsequent patents issued in a class, which indicates the size of the audience that might notice a focal patent, per H2a) and crowding (i.e., the number of existing patents in a class, which indicates the number of items that are competing with a focal patent for attention, per H2b and H4a,b). We used a single measure in our main analysis for the sake of parsimony: the alternate measures produced almost identical results.

Our last independent variable, proximity, captures the extent to which a given patent class is proximate to others in a classification system. For each system, we calculated the relative proximity of a focal class to all others using weighted degree centrality. In the network we constructed, the nodes are technology classes and the weights are the degrees of relatedness between two classes, defined as the relative frequency with which the two are listed together on granted patents. In a binary (un-weighted) network, degree centrality for a node is the number of edges or ties to the node (i.e., a count of connected nodes). In a weighted network, degree centrality is the sum of all of the weights of the edges or ties to the node. Based on this, we use Newman's (2005) weighted network extension of degree centrality, calculated as the sum of all of the weights for edges connected to the given node. We thus calculate proximity$_i$ — or the relative proximity of a focal technology class *i* to all others *j* —as:

$$Proximity_j = \sum_{for\ all\ i} P_{ij}$$

As the formula shows, proximity, is a patent-class level construct. Since our weights (degrees of relatedness between classes) are values between 0 and 1, the sum of these values produce a continuous measure greater than or equal to 0. A value of 0 indicates that a class is isolated. Moderate values indicate that a class is very proximate to a handful of others, or somewhat proximate to many. Large values indicate that a class is very proximate to many others.

Control Variables

Given that we are analyzing variance in forward citations within patent families—i.e., for the same innovation in different patent systems—our approach naturally controls for observable factors such as team composition, inventor characteristics, and knowledge inputs (Rosenkopf and Nerkar, 2001; Singh and Fleming, 2010), as well as unobservable factors, such as patent quality, that may affect citation rates. Controls are thus limited to variables that differ across patents belonging to the same family, and have been identified in prior studies as affecting forward citations. In terms of patent features, we control for the number of secondary classes, backward citations to prior patents, and unique claims listed on each patent (Podolny and Stuart, 1995; Singh and Fleming, 2010). We also control for the time in years (days/365) between application and grant dates, as our dependent variables measure citations from the date a patent was applied for (Lo and Kennedy, 2014).

Innovations may also be evaluated differently depending on the identity of the originator. For example, inventions from actors who are prominent (Trapido, 2015), or who have a focused patent portfolio (Hansen and Haas, 2001) in a given system may draw more attention and be more favorably viewed than others. Such considerations may apply to both inventors (i.e., the actors who create an innovation) and assignees (i.e., the patent-holding organization).[9] We control for these potential effects in three ways. First, we include dummy variables set to "1" for inventors and assignees that are among the top 10% of patenters in a given patent-system (results are the same when 1% or 5% are used as the cutoff). Second, we include the log number of prior patents granted to each inventor and assignee in each system. Third, we use a Herfindahl measure to assess the degree to which each inventor and assignee's patents are focused in a narrow range of classes in a given system, or dispersed among many classes.

To control for changes in examiner workload across art units over time, we include a measure that captures the percent change in applications from the month prior to the given month that the given patent is examined. We also include country fixed-effects with dummy variables for patents issued to German and Japanese inventors, making U.S. the baseline.

Estimation Strategy

Our empirical approach is a quasi-experiment that isolates the influence of categories on forward citations. We accomplish this by using patent family fixed-effects so that models analyze and explain cross-system variance for the same innovation. Put another way, our estimation strategy examines why patents for the same innovation are cited differently across patent systems. Our approach is thus similar to an experiment where random assignment into treatment groups ensures that subsample populations are similar. Instead of randomization, we use innovations that are represented in three patent systems and compare the differences among them. This approach is unique in both the innovation and categories literatures, where studies have focused almost exclusively on single jurisdictions and category systems.

We include a number of interaction terms among our main independent variables (e.g., proximity, density, and atypicality). All of these variables are z-scores based on their respective patent cohorts (i.e., same system and year). In the final sample of patent families spanning U.S., Germany, and Japan, we again converted these variables by mean-centering them and dividing by their standard deviations across the entire final sample. This ensures

[9] See the online appendix at http://categoriesandsearch.wordpress.com/ for details on the name-matching procedure we used to identify inventors and assignees.

that the mean is zero and the standard deviations are equal to 1. This is done to separate the main effects from the interaction effects and aid in model coefficient interpretation.

As our dependent variables are continuous z-scores (standardized 3-year forward citation measures), we ran models in Stata 14 using the xtreg command. All models include patent family fixed-effects and robust standard errors. Our dataset includes 7,288 patent families, giving us 21,864 (3 x 7,288) observations; one observation for each patent family in each of the three patent systems.

RESULTS

Table 1 lists descriptive statistics for all variables used in the model. Note that these include the transformed/standardized versions where appropriate. Table 2 lists the correlations among all variables in the model. The highest Variance Inflation Factor for any variable in any model is 7.53, which is below the generally accepted threshold of 10 (Aiken and West, 1991). As such, there should be no multicolinearity issues in our analysis. We also note that patent class proximity and density are positively correlated (0.65) (i.e., more central classes also tend to be densely populated). If the two measures were negatively correlated, particularly at a high level, this might suggest that the different patent systems were simply carving up similar patents in different ways (e.g., in one system, 20 patents might be in a single class, while in another system, a similar 20 patents might fall into 5 proximate classes). This would affect the interpretation of our results. Table 2 shows that this is not a concern.

-------Insert Tables 1 and 2 about here-------

Tables 3, 4, and 5 show the results of our analysis of within-category, across-category, and overall citations, respectively. Models 1, 4, and 7 are baseline models and only include control variables. Results are generally consistent with what we would expect based on the extant literature. Backward citations are positively related to all types of citations. The log of applicant prior patents and inventor prior patents are either positive or near zero for all types of forward citations. Additionally, the count of secondary classes is positively related to across-category and overall citations, but negatively related to within-class citations.

Approval time is positive related to all types of citations. We observe inconsistent effects for German and Japanese patents. It might be that these effects are picking up country level differences in types of innovation, level of recombination, scope of claims or other factors that might influence forward citations (Jaffe and Trajtenberg, 1999).

Hypotheses 1 and 2a predicted that patent atypicality would have a negative effect on within-class citations, while category density would have a positive effect. As predicted, model 3 (table 3) shows that the coefficient for atypicality is negative and statistically significant, thus providing support for H1. The coefficient of -0.068 indicates that a 1 standard deviation increase in the measure of atypicality leads to a 0.068 standard deviation reduction in the number of relative within-class forward citations (i.e., 2.9% fewer within-class citations for the average US patent in 2005). Model 3 also shows that the coefficient for category density is positive and statistically significant, which supports H2a. The coefficient of 0.127 indicates that a 1 standard deviation increase in the size of an innovation's primary technology class will increase the relative number of within-class forward citations by 0.127 standard deviations (i.e., 16.5% more within-class citations). Together, these findings lend credence to the idea that predictions gleaned from extant categories theory are germane to understanding an innovation's within-category influence.

-------Insert Table 3 about here-------

Hypotheses 3 and 4a predicted that tradeoffs exist, such that category density and patent atypicality will have the opposite effects on an innovation's influence across- versus within-categories. To wit, we expect that cross-category citations will be greater when a patent is more atypical and in a less dense technology class. Results in model 6 (table 4) support both predictions. The positive and statistically significant coefficient for atypicality supports H3. The coefficient of 0.077 shows that a 1 standard deviation increase in atypicality will increase the relative number of cross-category citations by 0.077 standard deviations (i.e., 2.43% more citations). We also observe that the coefficient for density is negative and significant, thus supporting H4a. The coefficient of -.106 indicates that a 1 standard deviation increase in the size of a patent's primary technology class will decrease the relative number of cross-class citations that it receives by 0.106 standard deviations (i.e., 14.2% fewer cross-class citations).

-------Insert Table 4 about here-------

Hypotheses 2b and 4b predicted that the relationship between density and both within- and across-category influence is moderated by patent atypicality. Hypothesis 2b predicted that the positive relationship between category density and within-category influence is weaker for atypical innovations. The positive and statistically significant coefficient for the interaction of density and atypicality along with a positive coefficient for density in Model 6 (table 4) supports H2b. The interaction is plotted in Figure 1. Hypothesis 4b predicted that the negative relationship between category density and across-category influence is weaker for atypical innovations. The interaction of density and atypicality reported in Model 6 is in the expected direction, but it is not statistically significant. This lack of significance may be due to the fact that, while the typical members of a category are by definition highly similar to each other—and may thus focus attention on a narrow set of stimuli—patents can be atypical in various ways. As category density increases, actors who are searching across classes may focus more intently on atypical items—consistent with the direction of our observed result—but different actors will likely focus on different types of atypical items.

-------Insert Figure 1 about here-------

The preceding analysis shows that category density and patent atypicality affect the pattern of citations that an innovation receives within- versus across-classes. Hypothesis 5 expanded this to consider how the proximity of a patent class to others might affect cross-category citations, and potentially an innovation's overall influence. H5 predicted that a patent will receive (a) more cross-category citations, and (b) more overall citations when it is in a more proximate technology class. Results in model 6 support H5a. The coefficient for proximity is .086, and significant. This means that a 1 standard deviation increase in the relative proximity of an innovation's original class to other categories improves its cross-category influence by 0.086 (i.e., 13.62% more cross-class citations). However, results in model 9 (table 5) do not lend strong support for H5b. The coefficient for proximity on overall citations is positive (0.017), but not statistically significant.

-------Insert Table 5 about here-------

Hypothesis 6 predicted that patent atypicality will magnify the positive effect of category proximity on an innovation's (a) cross-category influence and, (b) overall influence. We find support for both predictions. The positive, significant coefficient for the interaction of proximity and atypicality along with a positive coefficient for proximity in model 6 supports H6a, and shows that a patent has greater across-category influence when it is atypical and in a class that is proximate to others.

For overall influence, the positive and significant coefficient for the interaction of proximity and atypicality in model 9 along with the non-significant coefficient for proximity supports hypothesis H6b, though the interpretation is a bit more nuanced. As we mean-centered both the proximity and atypicality variables, the interpretation of the coefficient on proximity is the average effect of category proximity on forward citations. The results in model 9 thus suggest that category atypicality by itself shifts the distribution (i.e., within vs across-class), but not the overall amount, of citations a patent receives. Yet a patent is more likely to have greater (less) overall influence when it is atypical and in a class that is close to others (far from others). This indicates that the effects of atypicality on an innovation's overall influence are contingent on the class where it resides. Both interactions are plotted in Figures 2 and 3.

-------Insert Figures 2 and 3 about here-------

Robustness Checks

We ran a number of robustness checks and supplementary models to help rule out alternate explanations and bolster our findings. All models are available to view at http://categoriesandsearch.wordpress.com/.

Alternate variable measures. Our first set of additional models establish that our results are robust to different measures of our independent and dependent variables.

Our reported dependent variables include citations that are added to a patent by patent examiners. However, inventors may also add citations. Of note thought is that there is variance in the relative share of citations that are added by these groups in different patent systems. In Japan, only examiners add citations, and inventor-added citations are more common in the U.S. than in Germany. Our main models include only examiner-added citations to ensure that our results were not biased by these differences. However, for the sake of generalizability, we created alternate dependent variables that included both examiner-added and inventor-added citations. Results very closely matched our reported findings. Also, while it is uncommon for studies to include cross-system citations when calculating impact measures,[10] we recognize that such citations may affect a patent's overall influence. To address this, we created an alternate measure for overall influence that included foreign citations to each patent in our analysis (e.g., citations to U.S. patents from Japanese and German patents). Again, results were consistent with our reported findings.

We also took steps to ensure that our results are robust to alternate independent variable measures. Our reported density measure is based on patents that are issued in the same class-year as a focal patent. However, older patents may also be relevant to search and attention. As such, we created supplementary measures that track a category's cumulative, and rolling 5yr density. The first variable reflects the overall, accumulated activity in a category, and the second tracks how many patents are realistically competing for actors' attention in a given year (patents accumulate most of their citations within 5-yrs of being granted (Jaffe and Trajtenberg, 1999)). Unsurprisingly, all three density measures are highly correlated, and results are almost identical regardless of which one is included in our analysis.

We also considered alternate measures of atypicality. The first is a squared-measure that investigates if there are extra benefits or hazards to being very different from other patents in a class. Previous studies have found that moderate atypicality can help an item to stand-out in a category, but that problems ensue when an item is highly distinctive (e.g., Criscuolo et al.,

[10] Studies generally begin with a sample of patents from a single system, and calculate forward influence based on the number of times that a given patent is cited by others in the sample. Patents from multiple additional systems must be gathered to calculate a measure that includes cross-system citations.

2017). We see some evidence of this in our results, but only with regard to within-class influence. The coefficient for squared-atypicality is negative and statistically significant for within-class influence, but small and not statistically significant for across-class and overall influence. Our second measure follows previous studies (e.g., Hsu, 2006; Durand, Rao, and Monin, 2007), and considers the concentration of classes listed on a patent, regardless of how distant these classes are from each other. Our results hold when using this alternate measure. Our third alternate variable replicates the "tail novelty" measure that Uzzi et al. (2013) have found leads to greater influence among academic papers.[11] Items with high tail novelty are typical of a class, but integrate a few highly novel attributes. Results show that such patents are more likely to be cited by others in the same class. This is consistent with our argument that typical items are more likely to be built-upon by others in the same class, as well as with Uzzi et al.'s (2013) finding that tail novelty helps such items to stand out. We see no effect for across-class, nor overall citations, though. Our reported atypicality measure thus appears to be a better fit with our theoretical argument.

Our last independent variable is category proximity, which measures the degree to which a focal class is proximate to others in the same classification system. This shows us how an innovation's value is shaped by the structure of a classification system. However, it does not account for variance in the density of proximate classes. To check for any potential issues this might cause, we ran additional models using a weighted proximity measure. We calculated this by taking the original measure and multiplying the proximities between primary class i and all technology classes *j*, P*ij*, by the relative number of patents in the related class, Density*j*:

$$ProximityWeighted_j = \sum_{for\ all\ i} P_{ij} * Density_j$$

Again, results are consistent with our reported models.

[11] See the online appendix for a description of the steps that we used to calculate this variable.

Alternate mechanisms. Our second set of supplementary models help to establish that our results are not being driven by mechanisms other than those we have theorized.

One potential concern is that our analysis does not directly account for the mix of innovations in each patent system that are subsequent to the patents we examine. Indeed, while we offer a socio-cognitive explanation for why certain patents are more or less likely to be noticed and built-upon, subsequently issued patents provide the basis for these searches. As such, the composition of future innovations in a patent system are an objective constraint on the processes we have theorized: shifts in the types of patents issued in different systems may bias our results. To address this, we examined changes in yearly density and proximity for each patent-class-year in our analysis. Year-by-year correlations for both measures are very high in each patent system. This suggests that our results are not being affected by temporal shifts in patent activity.

To further investigate, we created alternate measures for category density and proximity at *t+1 yr*. We could not include these in our main models, as they are very-highly correlated with our reported measures. Still, substituting one set of variables for the other did not affect our results. We also ran models that used density at *t+1 yr* to more directly test our argument for H2 (i.e., that patents benefit when there is a larger audience searching within a class). As compared to models where we opted for parsimony, and used the same density measure for all hypothesis-tests, the alternate variable provides even stronger support for our argument.

Also, while results support our argument that category density negatively affects a patent's across-category influence (H3), we are sensitive to the counter-argument that a class may draw more attention as its popularity grows, resulting in more cross-category searches. To check this alternate prediction, we ran two sets of supplementary models at the patent class-level (i.e., one observation per system, year, and class) that use density to predict across-class citations. One set includes patent class fixed effects, leaving only variation within classes over time, and one general set that doesn't. Results show that density actually has a negative effect on the frequency with which the patents in a class are cited by those in other classes.

This suggests that single patents may garner less outside attention when they reside in popular classes.

Another potential concern is that our results reflect patterns of underlying similarity among individual patents, and that categories do not shape search and attention in the ways that we theorized. Our interview data partly address this concern. Still, we ran a number of tests to validate our empirical arguments. As a basic check, we compared the likelihood that a focal patent would cite another in its home class versus one in a different class. Using the same process as Jaffe and Trajtenberg (1999), we see that patents are 327 times more likely, on average, to cite another in the same class: this suggests that categories are indeed relevant to search and citation. As a further check, we repeated this analysis for patents in the top decile for atypicality. If atypical patents—which are by definition more-similar to those in other classes—are equally likely to cite others in the same versus different classes, this would imply that search takes place on a smooth terrain that is agnostic to category boundaries. However, we see that atypical patents are 307 times more likely to cite others in the same class: our results thus suggest that such patents may stand-out when search proceeds beyond a focal class, but that an actor must decide to start searching in a new class for this to happen.

To further to this line of thinking, we used a matching analysis to examine the influence of categories on citation. To do this, we matched each of a random subset of U.S. patents (N = 44,587) to two other patents; one in the same class, and one in a different class. Analytically, we used nearest neighbor matching based on bibliographic coupling (i.e., the degree to which two patents cite the same prior art, and are thus technologically similar).[12] We then compared co-citation patterns for each matched pair (i.e., the focal patent and its match in the same class vs. the focal patent and its match in a different class). Co-citation is the extent to which two patents are cited by the same future patents. If search and attention are not guided by categories, co-citations should be similar for both matches. To wit very similar patents should be similarly cited, regardless of whether or not they are in the same class. This was not the

[12] We discuss the effectiveness of these matches, and the equivalence of a patent's same-class and different-class matches in the online appendix.

case. Matched patents in the same class have much higher co-citation measures than matched patents from different classes, and the difference is highly significant ($p < 0.001$). We also calculated the Pearson correlation in forward citation vectors (i.e., each entry representing the number of citations from each technology class) for each patent and its matched counterparts. The correlation in forward citation vectors is much higher for the same-class matches than for the different-class matches, and this difference is highly significant ($p < 0.001$).

Taken together, these tests provide strong and consistent evidence that categories guide search and attention in ways that are consistent with our theory.

Generalizability. There are also potential concerns about the generalizability of our findings. One issue is that our analysis is based on patent families: these patents may be systematically different from those that are not in families. If so, our findings would be limited to a small subset of innovations. To investigate, we examined how the patents in our analysis compared to all others in the same patent system. The comparison included each variable in our main models, as well as patent-level features that do not vary within families (i.e., assignee type (corporate = 1), number of assignees, and number of inventors). We also compared 3-yr forward citations to check for differences in average influence. Overall, patents in families are very similar to patents not in families. Among all of the comparisons we made, there were only a few notable differences: 1) in the U.S., the average patent in a family has ~3.2 more backward citations than the average non-family patent; 2) in Germany, the average patent in a family has one more inventor, but longer approval times than the average patent not in a family; and 3) the average patent in a family has more claims than the average patent not in a family in Japan and Germany, but fewer in the U.S.. Based on these results it seems unlikely that our findings, and their generalizability, are adversely affected by selection effects.

Another potential issue is that, in analyzing variance across three patent systems, our results may reflect the vagaries of one system, rather than a more general phenomenon. To address this, we ran supplementary models for each country-pair in our analysis. The overall pattern of results is consistent across all models. Some significance levels are lower when comparing Germany and Japan, but this is not surprising, as both nations use modified versions of the

International Patent Classification (IPC) system, while the U.S. uses an entirely different system. As such, there is less variance on our independent variables when comparing Japan to Germany than when either country is compared to the U.S.. This is consistent with what we were told by interviewees with knowledge of all three patent systems (i.e., attorneys who had prosecuted patents in each system, and U.S. Patent Office staff who participated in patent system harmonization initiatives). They noted that our results aligned with their expectations. For instance, one told us:

> *The USPC [United States Patent Classification] and the IPC are very different historically in their development...variants [used in Germany and Japan] are mostly similar in terms of classes and classification…the U.S. has historically done its own different thing.*

Another similarly noted that:

> *I'm not surprised that you see more-similar effects in Germany and Japan…the JPO [Japanese Patent Office] has the F-term…system that we never had in Europe, and there's going to be some different patents, obviously…but the classification systems are pretty similar.*

DISCUSSION

Studies to date have shown that innovations with certain features are systematically more likely to become influential, and thus valuable, than others (e.g., Podolny and Stuart, 1995; Rosenkopf and Nerkar, 2001; Singh and Fleming, 2010). Without discounting this work or its findings, we suggest that an innovation's influence is also related to how the members of a creative community search for ideas to build-on in their own endeavors. To this end, we developed a novel theoretical approach that combines insights about categories, search, and attention to predict that an innovation's influence is shaped by the position that it occupies within a broader knowledge structure. As compared to studies that suggest search takes place on a smooth terrain that is oriented around an actor's unique knowledge and experiences (Gavetti et al., 2012), we cast search as a socio-cognitive process that is guided by the classes that are used to organize items in a given domain. We further argued that categories create bounded areas for actors to search within, and define the features of this terrain such that some items are systematically more likely to stand-out and receive attention than others.

Based on this, we predicted that an innovation's influence varies in patterned ways based on: a) its position within a given category; b) that category's internal properties, and; c) the category's position within a broader classification system.

In laying out our arguments, we reasoned that actors begin by searching in the class where they are innovating, and focus on typical items therein before proceeding to consider less-typical items. When search progresses to other categories, we argued that attention focuses on atypical items based on the expectation that these are the most likely to be germane for innovating in other classes. We predicted that crowding amplifies these effects, as it leads actors to focus more intently on the items they expect to be most relevant to their efforts. We rounded out our arguments by predicting that innovations—especially the atypical variety—will have more across-class, and overall, influence when they are in a class that is proximate to others, as such classes should be subject to more cross-class searches. Our approach thus makes systematic predictions about the conditions under which an innovation is more or less likely to be noticed and built upon by others in the same versus different classes, and when this will lead to greater overall influence.

A quasi-experiment that compared how patents for the same innovation are cited in different category systems supported our arguments. We see that the forward influence of the exact same innovation is shaped by its categorization and that these effects are both economically and statistically significant. Our approach has implications for research on search and attention, categories, and innovation.

Implications for search and attention. The approach we have developed is faithful to the ontological foundations of existing search and attention research (i.e., bounded rationality, limited information processing capacities, and the need to focus attention on selected stimuli) but links these to socio-cognitive influences. Rather than conceptualizing search as an actor-centric process that unfolds on a smooth terrain comprising items that are variously distant from an actor's knowledge (Gavetti et al., 2012), we advance a categories-based view. Our approach embraces the insight that categories simplify comparison tasks by supporting shared understandings about the similarity and distinctiveness of different items: actors cannot consider every possible choice alternative, so they rely on categories to provide manageable consideration-sets (Zuckerman, 1999). From this, we suggest that search is not idiosyncratic, but rather unfolds in patterned ways on a lumpy, structured, and socially-constructed terrain. Actors search piecemeal through different categories, and focus on different types of items depending on where they're searching. As such, we highlight the utility of considering how actors' representations of the environment are shaped by broader, contextual factors, in addition to their unique knowledge and experiences (Gavetti and Levinthal, 2001).

Our approach articulates with recent efforts to build theory at the intersection of categories and strategy. To date, this work has mostly focused on how firms identify their competitors (see Cattani, Porac, and Thomas, 2017). We go beyond this to offer a fresh view on search that may have implications for questions about strategy formation, learning, and adaptation that have typically been studied using behavioral theory. In this regard, studies have shown that distant knowledge is vital for innovation (Katila and Ahuja, 2002), differentiation (Katila, Chen, and Piezunka, 2012), and avoiding competency traps (Levinthal and March, 1993). Yet our theory and findings suggest that similarly distant items may be spread across multiple classes. In practice, then, distant search may not follow a linear progression, but rather a two-step process where actors first decide to search beyond their home class, and then engage in more distant search by selecting other classes to explore (c.f. Zuckerman, 1999). The types of distant knowledge that an actor finds may thus be contingent on the classes where they choose to search. Potentially useful stimuli may be missed due to the way that search is structured. Future studies should build on this to explore how search-selection choices affect the types of distant knowledge that actors attend to, and with what results.

Our theory also contributes to research at the intersection of culture and cognition (Hoffman and Ocasio, 2001; Nigam and Ocasio, 2010). To this end, we show that categories theory can be used to help explain how search and attention unfold in patterned ways in a given context. In turn, this can be used to predict outcomes that reflect the aggregate behavior of individual actors. We used this to advance understanding about innovation-level outcomes, but the same approach could be used to study path emergence, innovation trajectories, or other outcomes that are based on the actions of multiple, distributed actors (e.g., Garud, Kumaraswamy, and Karnøe, 2010). Our approach also suggests that there is value in considering how variables such as the density and visibility of alternate stimuli shape attention in ways that are not unique to specific organizations (Hansen and Haas, 2001; Piezunka and Dahlander, 2014).

Implications for categories. Our approach makes a number of contributions to the categories literature. Most notably, we show that categories are relevant for understanding influence-based valuation processes. To date, most studies have focused on how categories enable comparisons that help actors to determine the best option among a set of alternatives (Vergne and Wry, 2014). This work has usefully shown that categories shape the success prospects of competing offerings, but it neglects contexts where value is based on an item's influence. In addition to innovation—where many studies have demonstrated the link between influence and value—forward influence is also an important outcome in domains such as art, music, architecture, gastronomy, and science (Rao, Monin, and Durand, 2003; Singh and Fleming, 2010; Jones et al., 2012; Leahey, Beckman, and Stanko, 2017). For instance, a film or a piece of art can be negatively evaluated (or a commercial flop), but it may still be a source of prestige and status if becomes influential in its field (Rao and Giorgi, 2006; Delacour and Leca, 2017). We extend categories theory to directly account for such contexts and outcomes.

By focusing on how categories shape search and attention, our approach also offers novel insight into how atypicality and category density affect valuation. Studies have shown that typical items tend to be recognized quickly and evaluated favorably, whereas those that do not fit cleanly within a class are overlooked or devalued (e.g., Zuckerman, 1999; Hsu, 2006). We show that similar insights apply when considering the influence that an innovation has in its own class. Yet because influence relies on perceived relevance rather than comparative rankings, our approach suggests that the usual corollary does not apply. Indeed, we find that atypicality has a diametric effect on within- versus across-class influence. Being atypical does not result in decreased appeal; rather, it changes the nature of an innovation's influence. Our approach also extends the typical argument that being in a popular category is beneficial (Hsu, 2006). We show that this applies to an innovation's influence in its own class, but not in others. When actors are searching across-classes, crowding implies that more items are competing for attention, without creating a concomitant bump in audience size. In either case, crowding appears to lead actors to focus more intently on items that are the most likely to be relevant to their efforts. We expect that future studies will examine how such attention-focusing affects traditional category-based valuation processes as well.

Our approach also contributes to categories research by directly considering the role of category systems. Most categories studies focus on how an item's position within and across classes affect its evaluation. There is a general recognition that categories exist within classification systems, but this has been engaged primarily to identify moderators of the relationship between category-spanning and evaluation. Indeed, there is evidence that the penalties for category-spanning disappear when the boundary between classes is fuzzy (Rao, Monin, and Durand, 2005; Ruef and Patterson, 2009) or when items blend features of similar classes (Wry and Lounsbury, 2013; Leahey, Beckman, and Stanko, 2017). We go beyond this to consider the benefits of being in a class that is proximate to others, especially for atypical innovations. Our theory and findings apply most directly to innovation, but we expect future research will continue to investigate the relevance of category systems for valuation in other contexts. This has intriguing implications when considered together with recent work on goal-based categorization, which suggests that actors often pursue their goals by assembling ad-hoc categories that comprise similar items from multiple classes (Durand and Paolella, 2013; Paolella and Durand, 2016).

Implications for innovation. Our findings also contribute to innovation research. Studies have aptly shown that innovations with certain features are more likely than others to be influential. However, aside from early attempts to model technology niches (Podolny and Stuart, 1995; Podolny, Stuart, and Hannan, 1996), there has been little consideration of how contextual factors shape an innovation's influence. We extend this work in a number of ways. Notably, we show that category-effects explain considerable variance in an innovation's value, as reflected in its influence. Regardless of quality, novelty, or objective relevance to other innovation domains, the likelihood that a patent will be noticed and built-upon is shaped by its classification. We thus depart from extant studies, and suggest that an innovation's influence may be loosely-coupled to its technical features. In this way, we also highlight the potential contingency of commonly used variables in the innovation literature. Studies generally treat novelty, knowledge-breadth, and others as objective features of an innovation (Rosenkopf and Nerkar, 2001; Kaplan and Vakili, 2015). Our approach suggests that such characteristics may be assessed in relation to broader social structures. For instance, an innovation that appears novel (or atypical) in one context, may seem mundane (or typical) in another.

Also, while we are not the first to distinguish between an innovation's narrow versus broad impact, we advance this thinking in a number of ways. For one, we link these outcomes to socio-cognitive influences, rather than to an innovation's features (c.f. Rosenkopf and Nerkar, 2001). Indeed, our empirical design shows that patents with the same objective relevance to other innovation domains are cited differently depending on how they are classified. We also show that the diametric effect of atypicality is amplified by category density, which further suggests that categories shape the depth and breadth of an innovation's influence. Most notably, though, we show that atypicality only results in greater overall influence when a patent is in a class that is proximate to others. This adds an important boundary condition to studies that have argued broad recombination is a recipe for impactful innovation (Ahuja and Lampert, 2001; Rosenkopf and Nerkar, 2001; Kaplan and Vakili, 2015).

Limitations

As with all studies, ours has limitations that point to opportunities for future research. One of the strengths of our paper is that we are able to econometrically isolate category effects while holding most other features of an innovation constant. However, in so doing, we are unable to observe cross-sectional differences in how categories affect the influence of different types of innovations. While our analysis suggests that patents in families do not differ systematically from other patents, we cannot rule out the possibility that the effects we observe might vary among innovations with different features (Ferguson and Carnabuci, 2017). Our findings might differ for innovations that are of high or low quality; that are produced by teams versus individuals (Singh and Fleming, 2010); or that are assigned to different types of organizations (Trajtenberg, Henderson, and Jaffe, 1997). Future studies should address these issues.

Also, while our theory suggests that categories shape forward influence by guiding search behavior, we cannot observe this directly. We partially addressed this by interviewing patent examiners and inventors, but it is possible that search behavior varies among actors with different expertise, time pressures, and other constraints (Payne, Bettman, and Luce, 1996). As with other studies of innovation-level outcomes, we assume that these variables are normally distributed and do not bias our results (Podolny and Stuart, 1995; Hansen and Hass, 2001; Piezunka and Dahlander, 2015). Future research should investigate this directly and work to more cleanly isolate the relationship between categories and search.

Finally, while patent classes have historically guided search and examination, this may be changing as patent systems move toward technology-enabled search. Categories simplify the world so as to allow actors to process vast information quickly and efficiently (Rosch, 1975). To the extent that technology is unencumbered by the same cognitive limits, the use of formal, static categories may wane. The temporal generalizability of our findings may thus be limited. However, this points to opportunities to study how valuation changes with the use of technology-mediated categorization and search processes. Indeed, this may have implications well-beyond our context, as firms are increasingly relying on algorithms to dynamically generate considerations-sets of offerings based on specific, user-generated features.

Conclusion

In this study, we developed and tested novel theoretical predictions about how categories shape search and attention, such that some innovations are systematically more likely be noticed and built-upon than others. Using a quasi-experimental design that allowed us to isolate category effects, we found that an innovation’s influence is shaped by its position within a broader knowledge-structure, regardless of its underlying features. In so doing, we extend categories theory to account for valuation processes that are based forward influence, as opposed to comparative evaluation, and in the process demonstrate the value of a socio-cognitive approach to search and attention.

REFERENCES

Ahuja, G. and C. Lampert
2001 “Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions.” Strategic Management Journal, 22(6-7): 521-543.

Aiken, L. S., S. G. West, and R. R. Reno
1991 Multiple Regression: Testing and Interpreting Interactions: Sage.

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.

Alcacer, J. and M. Gittelman
2006 "Patent citations as a measure of knowledge flows: The influence of examiner citations." The Review of Economics and Statistics, 88: 774-779.

Allison, J. R. and M. A. Lemley
1998 "Empirical evidence on the validity of litigated patents."

Bacchiocchi, E. and F. Montobbio
2009 "Knowledge diffusion from university and public research. A comparison between US, Japan and Europe using patent citations." The Journal of Technology Transfer, 34: 169-181.

Bartholomew, S.
1997 "National systems of biotechnology innovation: complex interdependence in the global system." Journal of International Business Studies: 241-266.

Bellis, M
2014 “Patent application tips: Tips on writing the claims for a patent application.” Thought Co. Accessed at https://www.thoughtco.com/writing-patent-claims-tips.

Boudreau, K. , E. Guinan, K. Lakhani, and C. Riedl
2016 “Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science.” Management Science, 62(10): 2765-2783.

Bowers, A.
2014 "Relative comparison and category membership: The case of equity analysts." Organization Science, 26: 571-583.

Bowers, A. and M. Prato
2018 “The structural origins of unearned status: How arbitrary changes in categories affect status position and market impact.” Administrative Science Quarterly, 63(3): 668-699.

Bowker, G. and S. L. Star

1999 Sorting Things Out: Classification and its Consequences. Cambridge, MA: MIT Press.

Cantwell, J.
1989 Technological Innovation and Multinational Corporations. Cambridge, MA: B. Blackwell.

Cattani, G., J. Porac, and H. Thomas
2017. “Categories and competition.” Strategic Management Journal, 38(1): 64-92.

Cohen, W. M. and D. A. Levinthal
1990 "Absorptive capacity: A new perspective on learning and innovation." Administrative Science Quarterly: 128-152.

Cyert, R. M. and J. G. March
1963 A Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice-Hall.

Delacour, H. and B. Leca
2017 “The paradox of controversial innovation: Insights from the rise of impressionism.” Organization Studies, 38(5): 597-618.

Durand, R., H. Rao, and P. Monin
2007 “Code and conduct in French cuisine: Impact of code changes on external evaluations.” Strategic Management Journal, 28(5): 455-472.

Durand, R. and L. Paolella
2013 "Category stretching: Reorienting research on categories in strategy, entrepreneurship, and organization theory." Journal of Management Studies, 50: 1100-1123.

Elsbach, K. D. and R. M. Kramer
2003 “Assessing creativity in Hollywood pitch meetings: Evidence for a dual-process model of creativity judgments.” Academy of Management Journal, 46(3): 283-301.

Ferguson and Carnabuci
2017 “Risky Recombinations: Institutional Gatekeeping in the Innovation Process.” Organization Science 28(1): 133-151.

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

Garud, R., A. Kumaraswamy, and P. Karnøe
2010. “Path dependence or path creation?” Journal of Management Studies, 47(4): 760-774.

Gavetti, G. and D. Levinthal
2000 "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly, 45: 113-137.

2001 “Bringing cognition back in and moving forward.” Journal of Management and Governance, 5(3-4): 213-216.

Gavetti, G., H. Greve, D. Levinthal and W. Ocasio
2012 “The behavioral theory of the firm: Assessment and prospects.” The Academy of Management Annals, 6(1): 1-40.

Gittelman, M. and B. Kogut
2003 "Does good science lead to valuable knowledge? Biotechnology firms and the evolutionary logic of citation patterns." Management Science, 49: 366-382.

Goto, A. and K. Motohashi
2007 "Construction of a Japanese Patent Database and a first look at Japanese patenting activities." Research Policy, 36: 1431-1442.

Grodal, S.
2017 “Field expansion and contraction: How communities shape social and symbolic boundaries. Administrative Science Quarterly, forthcoming.

Guimmo, J.
2003 “Should all patentable inventions receive equal protection? Identifying the sources of heterogeneity in patent value.” Discussion Paper, Vol. University of Berkeley.

Haas, M., P. Criscuolo, and G. George
2015 “Which problems to solve? Online knowledge sharing and attention allocation in organizations.” Academy of Management Journal, 58(3): 680-711.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg
2001 "The NBER patent citation data file: Lessons, insights and methodological tools." National Bureau of Economic Research.

2005 "Market Value and Patent Citations." The RAND Journal of Economics, 36: 16-38.

Hansen, M. T. and M. R. Haas
2001 "Competing for attention in knowledge markets: Electronic document dissemination in a management consulting company." Administrative Science Quarterly, 46: 1-28.

Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel
1999 "Citation frequency and the value of patented inventions." Review of Economics and Statistics, 81: 511-515.

Henderson, R., A. Jaffe, and M. Trajtenberg
2005 "Patent citations and the geography of knowledge spillovers: A reassessment: Comment." The American Economic Review, 95: 461-464.

Hoffman, A. and W. Ocasio

2001 “Not all events are attended equally: Toward a middle-range theory of industry attention to external events.” Organization science, 12(4): 414-434.

Hsu, G.
2006 "Jacks of all trades and masters of none: Audiences' reactions to spanning genres in feature film production." Administrative Science Quarterly, 51: 420-450.

Hsu, G., M. T. Hannan, and Ö. Koçak
2009 "Multiple category memberships in markets: An integrative theory and two empirical tests." American Sociological Review, 74: 150-169.

Jacobs, J. A. and S. Frickel
2009 "Interdisciplinarity: A critical assessment." Annual Review of Sociology, 35: 43-65.

Jaffe, A. B., and M. Trajtenberg
1996 “Flows of knowledge from universities and federal laboratories: Modeling the flow of patent citations over time and across institutional and geographic boundaries.” proceedings of the National Academy of Sciences, 93(23):12671-12677.

1999 “International knowledge flows: evidence from patent citations.” Economics of Innovation and New Technology, 8: 105-136.

Jaffe, A. B., M. Trajtenberg, and R. Henderson
1993 "Geographic localization of knowledge spillovers as evidenced by patent citations." Quarterly Journal of Economics, 108: 577-598.

Jaffe, A.B., M. Trajtenberg, and M. Fogarty
2000 “Knowledge spillovers and patent citations: Evidence from a survey of inventors.” American Economic Review, 90(2): 215-218.

Jones, C., M. Maoret, F. G. Massa, and S. Svejenova
2012 “Rebels with a cause: Formation, contestation, and expansion of the de novo category “modern architecture, 1870–1975.” Organization Science, 23(6): 1523-1545.

JPO
2017 “Searching for Patents, Utility Models, Designs, Trademarks.” Japan Patent Office. Accessed at jpo.go.jp/english/faqs/searching.html.

Kaplan, S. and K. Vakili
2015 "The double-edged sword of recombination in breakthrough innovation." Strategic Management Journal, 36: 1435-1457.

Katila, R.
2002 “New product search over time: past ideas in their prime?” Academy of Management Journal, 45(5): 995-1010.

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: 1183-1194.

Katila, R., E. Chen, and H. Piezunka
2012 "All the right moves: How entrepreneurial firms compete effectively." Strategic Entrepreneurship Journal, 6(2): 116-132.

Kovács B and M. Hannan
2010 "The consequences of category spanning depend on contrast." In Categories in markets: Origins and evolution (pp. 175-201). Emerald Group Publishing Limited.

Leahey, E., C. M. Beckman, and T. L. Stanko
2017 "Prominent but Less Productive The Impact of Interdisciplinarity on Scientists' Research." Administrative Science Quarterly, 62(1): 105-139.

Levinthal D. and J. March
1993 "The myopia of learning." Strategic Management Journal, 14(S2): 95-112.

Levitt B. and J. March
1988 "Organizational learning." Annual Review of Sociology, 14(1): 319-338.

Li, Q. P.G. Maggitti, K.G. Smith, P.E. Tesluk, and R. Katila
2013 "Top management attention to innovation: The role of search selection and intensity in new product introductions." Academy of Management Journal 56: 893-916.

Lo, J. Y.-C. and M. T. Kennedy
2014 "Approval in nanotechnology patents: Micro and macro factors that affect reactions to category blending." Organization Science, 26: 119-139.

Lounsbury, M. and M. A. Glynn
2001 "Cultural entrepreneurship: Stories, legitimacy, and the acquisition of resources." Strategic Management Journal, 22: 545-564.

Maggitti, P. G., K. G. Smith, and R. Katila
2013 "The complex search process of invention." Research Policy, 42(1): 90-100.

Marquis C. and M. Lounsbury
2007 "Vive la résistance: Competing logics and the consolidation of US community banking." Academy of Management Journal, 50(4): 799-820.

Martens, M., J. Jennings, and P. Jennings
2007 "Do the stories they tell get them the money they need? The role of entrepreneurial narratives in resource acquisition." Academy of Management Journal, 50(5): 1107-1132.

Monahan, G. E.
1982 "State of the art—a survey of partially observable Markov decision processes: theory, models, and algorithms." Management Science, 28: 1-16.

Mukherjee, S., B. Uzzi, B. F. Jones, and M. Stringer

2017 "How Atypical Combinations of Scientific Ideas Are Related to Impact: The General Case and the Case of the Field of Geography." Knowledge and Networks: 243-267: Springer.

Negro, G. and M. D. Leung
2013 "“Actual” and perceptual effects of category spanning." Organization Science, 24: 684-696.

Newman, M. D.
2005 “A measure of betweenness centrality based on random walks.” Social Networks, 27(1): 39-54.

Nigam A. and W. Ocasio
2010 “Event attention, environmental sensemaking, and change in institutional logics: An inductive analysis of the effects of public attention to Clinton's health care reform initiative.” Organization Science, 21(4): 823-841.

Ocasio, W.
1997 "Towards an attention-based view of the firm." Strategic Management Journal, 18: 187-206.

2011 "Attention to attention." Organization Science, 22: 1286-1296.

Pahnke, E., R. Katila, and K. Eisenhardt
2015 “Who takes you to the dance? How partners’ institutional logics influence innovation in young firms.” Administrative Science Quarterly, 60(4): 596-633.

Paolella, L. and R. Durand
2016 "Category spanning, evaluation, and performance: Revised theory and test on the corporate law market." Academy of Management Journal, 59: 330-351.

Payne, J., J.Bettman, and M. Luce
1996 “When time is money: Decision behavior under opportunity-cost time pressure.” Organizational Behavior and Human Decision Processes, 66(2): 131-152.

Pfeffer, J.
1993 "Barriers to the advance of organizational science: Paradigm development as a dependent variable." Academy of Management Review, 18: 599-620.

Piezunka, H. and L. Dahlander
2015 "Distant search, narrow attention: How crowding alters organizations’ filtering of suggestions in crowdsourcing." Academy of Management Journal, 58: 856-880.

Podolny, J. M. and T. E. Stuart
1995 "A role-based ecology of technological change." American Journal of Sociology, 100: 1224-1260.

Podolny, J., T. Stuart, and M. Hannan
1996 “Networks, knowledge, and niches: Competition in the worldwide semi-conductor industry, 1984-1991.” American Journal of Sociology, 102(3): 659-689.

Porac, J. and H. Thomas
1990 “Taxonomic mental models in competitor definition.” Academy of Management Review, 15(2): 224-240.

Porter, A. L., A. S. Cohen, J. D. Roessner, and M. Perreault
2007 "Measuring researcher interdisciplinarity." Scientometrics, 72: 117-147.

Rao, H. and S. Giorgi
2006 “Code breaking: How entrepreneurs exploit cultural logics to generate institutional change.” Research in Organizational Behavior, 27: 269-304.

Rao, H., P. Monin, and R. Durand
2003 "Institutional Change in Toque Ville: Nouvelle Cuisine as an Identity Movement in French Gastronomy 1." American Journal of Sociology, 108: 795-843.

Rips, L., J.
1975 “Inductive judgments about natural categories.” Journal of Memory and Language, 14(6): 665.

Rosch, E.
1975 “Cognitive reference points.” Cognitive Pychology, 7(4): 532-547.

Rosch, E. and B. B. Lloyd
1978 Cognition and Categorization. Hillsdale, NJ: Lawrence Erlbaum Associates.

Rosch, E., C. Simpson, and R. Miller
1976 “Structural bases of typicality effects.” Journal of Experimental Psychology: Human Perception and Performance, 2(4): 491.

Rosenkopf, L. and A. Nerkar
2001 "Beyond local search: boundary-spanning, exploration, and impact in the optical disk industry." Strategic Management Journal, 22: 287-306.

Ruef, M. and K. Patterson
2009 "Credit and classification: The impact of industry boundaries in nineteenth-century America." Administrative Science Quarterly, 54: 486-520.

Singh, J. and L. Fleming
2010 "Lone inventors as sources of breakthroughs: Myth or reality?" Management Science, 56: 41-56.

Stuart, T. E. and J. M. Podolny
1996 "Local search and the evolution of technological capabilities." Strategic Management Journal, 17: 21-38.

Sullivan, 2010
2010 “Competition and beyond: Problems and attention allocation in the organizational rulemaking process.” Organization Science, *21*(2): 432-450.

Thornton, P., W. Ocasio, and M. Lounsbury
2012. The institutional logics perspective: A new approach to culture, structure, and process. Oxford University Press on Demand.

Trajtenberg, M.
1990 “A penny for your quotes: patent citations and the value of innovations.” *The* Rand Journal of Economics, 172-187.

Trajtenberg, M., R. Henderson, and A. Jaffe
1997 "University versus corporate patents: A window on the basicness of invention." Economics of Innovation and new technology, 5: 19-50.

Trapido, D.
2015 “How novelty in knowledge earns recognition: The role of consistent identities.” Research Policy, 44(8): 1488-1500.

Tversky, A.
1977 “Features of similarity.” Psychological review, *84*(4): 327.

USPTO
2012. Overview of the U.S. Patent Classification System (USPC). Accessed at https://www.uspto.gov/sites/default/files/patents/resources/classification/overview.pdf

2017. 7-Step U.S. Patent Search Strategy Guide. Accessed at https://www.uspto.gov/sites/default/files/documents/7%20Step%20US%20Patent%20Search%20Strategy%20Guide%20%282016%29%20Long%20Version.pdf

Uzzi, B., S. Mukherjee, M. Stringer and B. Jones
2013 “Atypical combinations and scientific impact.” Science, 342(6157): 468-472.

Vergne, J. P. and T. Wry
2014 "Categorizing categorization research: Review, integration, and future directions." Journal of Management Studies, 51: 56-94.

Wry, T. and M. Lounsbury
2013 "Contextualizing the categorical imperative: Category linkages, technology focus, and resource acquisition in nanotechnology entrepreneurship." Journal of Business Venturing, 28: 117-133.

Wry, T., M. Lounsbury, and M. A. Glynn,
2011 “Legitimating nascent collective identities: Coordinating cultural entrepreneurship. Organization Science, 22(2): 449-463.

Wry, T., M. Lounsbury, and P. D. Jennings
2014 "Hybrid vigor: Securing venture capital by spanning categories in nanotechnology." Academy of Management Journal, 57: 1309-1333.

Yanagisawa, T.

2016 “Outline of the Patent Examination Process at the Japan Patent Office (JPO)” Accessed at http://www.wipo.int/edocs/mdocs/aspac/en/wipo_ip_kul_2_16/wipo_ip_kul_2_16_1.pdf.

Yasukawa, S. and S. Kano
2015 "Comparison of examiners’ forward citations in the United States and Japan with pairs of equivalent patent applications." Scientometrics, 102: 1189-1205.

Zuckerman, E. W.
1999 "The categorical imperative: Securities analysts and the illegitimacy discount." American Journal of Sociology, 104: 1398-1438.

Table 1. Descriptive Statistics (N = 21,864)

Variable	Mean	S.D.	Min	Max	Source
1. 3-year forward citation z-score (all)	0.82	1.91	0.00	47.00	PATSTAT
2. 3-year forward citation z-score (within-class)	0.47	1.31	0.00	29.00	PATSTAT
3. 3-year forward citation z-score (cross-class)	0.34	1.04	0.00	32.00	PATSTAT
4. Secondary classes	-0.01	1.01	-0.91	8.37	PATSTAT
5. Backward citations	0.03	1.11	-1.05	26.30	PATSTAT
6. Inventor prior patents	-0.05	0.87	-0.61	14.54	PATSTAT
7. Inventor prominence	0.76	0.42	0.00	1.00	CONSTRUCTED
8. Inventor tech HHI (Focus)	-0.21	0.97	-6.70	1.06	CONSTRUCTED
9. Applicant prior patents	-0.05	0.87	-0.60	15.12	PATSTAT
10. Applicant prominence	0.76	0.42	0.00	1.00	CONSTRUCTED
11. Applicant tech HHI (Focus)	-0.34	0.89	-55.82	1.75	CONSTRUCTED
12. Claims	0.24	1.20	-1.22	19.39	PATSTAT
13. Time to approval	0.16	0.98	-2.70	3.00	PATSTAT
14. Art unit change in applications	0.14	0.53	-0.93	14.00	PATSTAT
15. Grant year	2002.17	3.70	1995.00	2010.00	PATSTAT
16. German Patent	0.33	0.47	0.00	1.00	PATSTAT
17. Japanese Patent	0.33	0.47	0.00	1.00	PATSTAT
18. Proximity	0.00	1.00	-1.53	4.69	CONSTRUCTED
19. Density	0.00	1.00	-5.12	2.31	PATSTAT
20. Atypicality	0.00	1.00	-1.08	2.31	CONSTRUCTED

Table 2. Correlations (N = 21,864)

Variable	1	2	3	4	5	6	7	8	9	10	11	12	13
1. 3-year forward citation z-score (all)													
2. 3-year forward citation z-score (within-class)	0.85												
3. 3-year forward citation z-score (cross-class)	0.76	0.30											
4. Secondary classes	0.02	-0.07	0.13										
5. Backward citations	0.10	0.07	0.09	0.05									
6. Inventor prior patents	0.10	0.09	0.07	0.02	0.01								
7. Inventor prominence	0.05	0.05	0.04	0.01	0.00	0.25							
8. Inventor tech HHI (Focus)	-0.04	0.01	-0.08	-0.12	-0.09	-0.11	-0.13						
9. Applicant prior patents	0.10	0.09	0.07	0.02	0.01	0.93	0.25	-0.10					
10. Applicant prominence	0.05	0.04	0.04	0.00	0.00	0.24	0.93	-0.13	0.25				
11. Applicant tech HHI (Focus)	-0.08	-0.06	-0.07	-0.04	-0.06	-0.19	-0.37	0.36	-0.20	-0.37			
12. Claims	-0.06	-0.06	-0.04	0.04	-0.03	-0.08	-0.13	0.06	-0.08	-0.12	0.08		
13. Time to approval	-0.09	-0.09	-0.04	0.08	0.06	0.06	0.06	-0.12	0.05	0.05	0.02	0.01	
14. Art unit change in applications	0.00	-0.01	0.01	0.01	-0.01	0.01	0.00	-0.03	0.01	0.00	-0.01	0.02	0.04
15. Grant year	-0.27	-0.24	-0.20	0.06	0.03	-0.06	0.04	-0.03	-0.06	0.04	0.12	0.03	0.56
16. German Patent	-0.21	-0.19	-0.15	0.05	0.01	0.05	0.04	-0.23	0.05	0.04	-0.22	0.05	0.38
17. Japanese Patent	-0.25	-0.21	-0.19	-0.02	0.03	-0.14	-0.07	0.21	-0.14	-0.05	0.28	0.19	-0.17
18. Proximity	0.04	0.00	0.06	0.10	0.00	0.06	0.00	-0.03	0.06	0.00	-0.05	0.04	0.06
19. Density	0.03	0.05	0.00	-0.04	0.02	0.11	0.07	-0.01	0.11	0.06	-0.07	0.03	0.08
20. Atypicality	0.03	-0.07	0.14	0.85	0.04	0.03	0.01	-0.11	0.02	0.01	-0.05	0.03	0.07

*** The correlations are for the full pooled sample and include all observations for the US, Germany, and Japan.**

Table 2. Correlations (continued)

Variable	14	15	16	17	18	19
15. Grant year	-0.02					
16. German Patent	0.06	0.12				
17. Japanese Patent	-0.04	0.34	-0.50			
18. Proximity	-0.10	-0.04	0.10	-0.12		
19. Density	-0.20	-0.02	0.18	-0.12	0.65	
20. Atypicality	0.01	0.04	0.05	-0.04	0.09	-0.04

Table 3. Panel Models of Within-Class Forward Citations

	Model 1		Model 2		Model 3	
Constant	62.530	***	58.953	***	59.463	***
	(16.172)		(16.150)		(16.149)	
Secondary classes	-0.055	***	0.009		0.008	
	(0.011)		(0.018)		(0.018)	
Backward citations	0.089	***	0.087	***	0.087	***
	(0.009)		(0.009)		(0.009)	
Inventor prior patents	-0.013		-0.012		-0.012	
	(0.029)		(0.029)		(0.029)	
Inventor prominence	-0.022		-0.025		-0.025	
	(0.062)		(0.062)		(0.062)	
Inventor tech focus	0.030	**	0.026	**	0.026	**
	(0.012)		(0.012)		(0.012)	
Applicant prior patents	0.086	***	0.080	***	0.078	***
	(0.029)		(0.029)		(0.029)	
Applicant prominence	0.045		0.046		0.045	
	(0.063)		(0.063)		(0.063)	
Applicant tech focus	-0.015		-0.017		-0.018	
	(0.013)		(0.013)		(0.013)	
Claims	0.016		0.015		0.015	
	(0.011)		(0.011)		(0.011)	
Time to approval	0.064	***	0.061	***	0.062	***
	(0.023)		(0.023)		(0.023)	
Art unit change in applications	-0.021		0.003		0.003	
	(0.019)		(0.019)		(0.019)	
Grant year	-0.031	***	-0.029	***	-0.030	***
	(0.008)		(0.008)		(0.008)	
German patent	-0.297	***	-0.337	***	-0.339	***
	(0.024)		(0.025)		(0.025)	
Japanese patent	-0.062		-0.072	*	-0.071	*
	(0.039)		(0.039)		(0.039)	
Proximity			-0.055	***	-0.055	***
			(0.016)		(0.016)	
Density			0.125	***	0.127	***
			(0.017)		(0.017)	
Atypicality			-0.067	***	-0.068	***
			(0.018)		(0.018)	
Proximity X Atypicality					0.011	
					(0.014)	
Density X Atypicality					-0.032	**
					(0.014)	
R^2	0.034		0.039		0.039	
N	21,864		21,864		21,864	
Patent Families	7,288		7,288		7,288	
N per Patent Family	3.000		3.000		3.000	
rho	0.274		0.271		0.271	
r2_b	0.014		0.029		0.030	
r2_o	0.026		0.035		0.036	
F	36.552		34.295		30.999	

*** p < 0.10; ** p < 0.05; *** p < 0.01; two-tailed tests.**
*** Clustered standard errors (by patent family) are in parentheses**
*** Patent family fixed effects are included in all models**

Table 4. Panel Models of Across-Class Forward Citations

	Model 4		Model 5		Model 6	
Constant	142.995	***	146.085	***	145.389	***
	(14.860)		(14.835)		(14.827)	
Secondary classes	0.082	***	0.011		0.010	
	(0.010)		(0.017)		(0.017)	
Backward citations	0.060	***	0.063	***	0.063	***
	(0.008)		(0.008)		(0.008)	
Inventor prior patents	0.088	***	0.087	***	0.087	***
	(0.027)		(0.027)		(0.027)	
Inventor prominence	-0.118	**	-0.113	**	-0.113	**
	(0.057)		(0.057)		(0.057)	
Inventor tech focus	-0.062	***	-0.058	***	-0.058	***
	(0.011)		(0.011)		(0.011)	
Applicant prior patents	-0.017		-0.012		-0.011	
	(0.027)		(0.027)		(0.027)	
Applicant prominence	0.175	***	0.173	***	0.174	***
	(0.058)		(0.058)		(0.058)	
Applicant tech focus	-0.006		-0.003		-0.003	
	(0.012)		(0.012)		(0.012)	
Claims	-0.004		-0.004		-0.003	
	(0.010)		(0.010)		(0.010)	
Time to approval	0.149	***	0.152	***	0.150	***
	(0.021)		(0.021)		(0.021)	
Art unit change in applications	0.001		-0.014		-0.013	
	(0.017)		(0.017)		(0.017)	
Grant year	-0.072	***	-0.073	***	-0.073	***
	(0.007)		(0.007)		(0.007)	
German patent	-0.044	**	-0.016		-0.007	
	(0.022)		(0.023)		(0.023)	
Japanese patent	0.206	***	0.225	***	0.230	***
	(0.036)		(0.036)		(0.036)	
Proximity			0.094	***	0.086	***
			(0.015)		(0.015)	
Density			-0.104	***	-0.106	***
			(0.016)		(0.016)	
Atypicality			0.077	***	0.077	***
			(0.017)		(0.017)	
Proximity X Atypicality					0.035	***
					(0.013)	
Density X Atypicality					0.012	
					(0.013)	
R^2	0.028		0.033		0.034	
N	21,864		21,864		21,864	
Patent Families	7,288		7,288		7,288	
N per Patent Family	3.000		3.000		3.000	
rho	0.278		0.280		0.280	
r2_b	0.012		0.013		0.013	
r2_o	0.017		0.020		0.021	
F	29.771		29.225		27.230	

*** p < 0.10; ** p < 0.05; *** p < 0.01; two-tailed tests.**

*** Clustered standard errors (by patent family) are in parentheses**

*** Patent family fixed effects are included in all models**

Table 5. Panel Models of All Forward Citations

	Model 7		Model 8		Model 9	
Constant	134.185	***	133.543	***	133.478	***
	(15.515)		(15.525)		(15.525)	
Secondary classes	0.015		0.011		0.010	
	(0.010)		(0.018)		(0.018)	
Backward citations	0.100	***	0.100	***	0.100	***
	(0.009)		(0.009)		(0.009)	
Inventor prior patents	0.045		0.045		0.045	
	(0.028)		(0.028)		(0.028)	
Inventor prominence	-0.088		-0.087		-0.087	
	(0.060)		(0.060)		(0.060)	
Inventor tech focus	-0.018		-0.018		-0.018	
	(0.011)		(0.011)		(0.011)	
Applicant prior patents	0.053	*	0.052	*	0.051	*
	(0.028)		(0.028)		(0.028)	
Applicant prominence	0.141	**	0.140	**	0.141	**
	(0.061)		(0.061)		(0.061)	
Applicant tech focus	-0.011		-0.011		-0.011	
	(0.013)		(0.013)		(0.013)	
Claims	0.008		0.007		0.008	
	(0.011)		(0.011)		(0.011)	
Time to approval	0.137	***	0.136	***	0.136	***
	(0.022)		(0.022)		(0.022)	
Art unit change in applications	-0.017		-0.009		-0.009	
	(0.018)		(0.018)		(0.018)	
Grant year	-0.067	***	-0.067	***	-0.067	***
	(0.008)		(0.008)		(0.008)	
German patent	-0.225	***	-0.236	***	-0.232	***
	(0.023)		(0.024)		(0.024)	
Japanese patent	0.101	***	0.106	***	0.109	***
	(0.038)		(0.038)		(0.038)	
Proximity			0.021		0.017	
			(0.015)		(0.016)	
Density			0.023		0.024	
			(0.017)		(0.017)	
Atypicality			0.005		0.005	
			(0.017)		(0.017)	
Proximity X Atypicality					0.029	**
					(0.013)	
Density X Atypicality					-0.016	
					(0.014)	
R^2	0.040		0.041		0.041	
N	21,864		21,864		21,864	
Patent Families	7,288		7,288		7,288	
N per Patent Family	3.000		3.000		3.000	
rho	0.285		0.284		0.284	
r2_b	0.005		0.007		0.007	
r2_o	0.020		0.022		0.022	
F	43.400		36.298		32.745	

*** p < 0.10; ** p < 0.05; *** p < 0.01; two-tailed tests.**
*** Clustered standard errors (by patent family) are in parentheses**
*** Patent family fixed effects are included in all models**

Figure 1. Within-Class Citations – Density and Atypicality Interaction

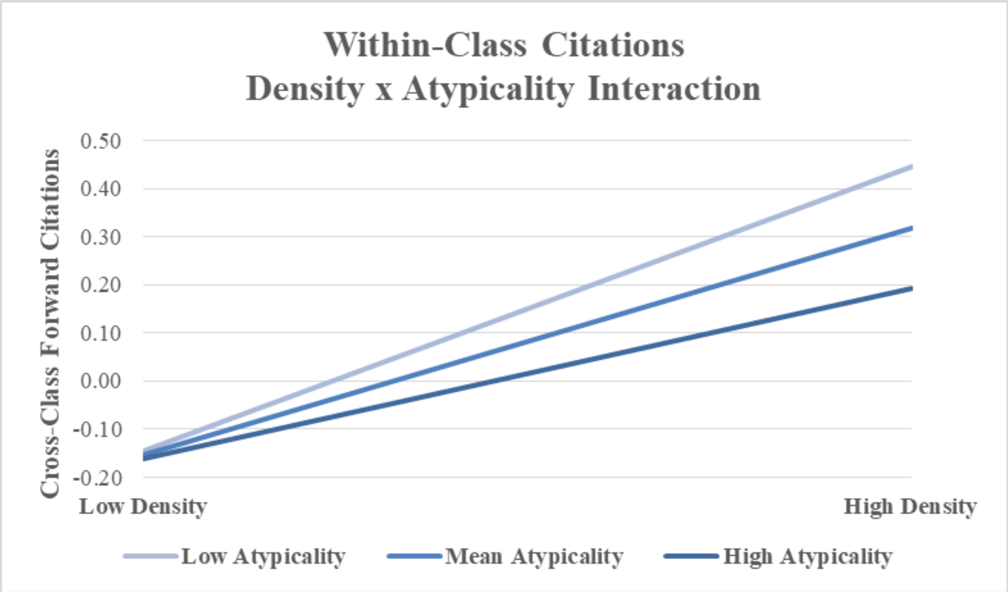

Figure 2. Across-Class Citations – Proximity and Atypicality Interaction

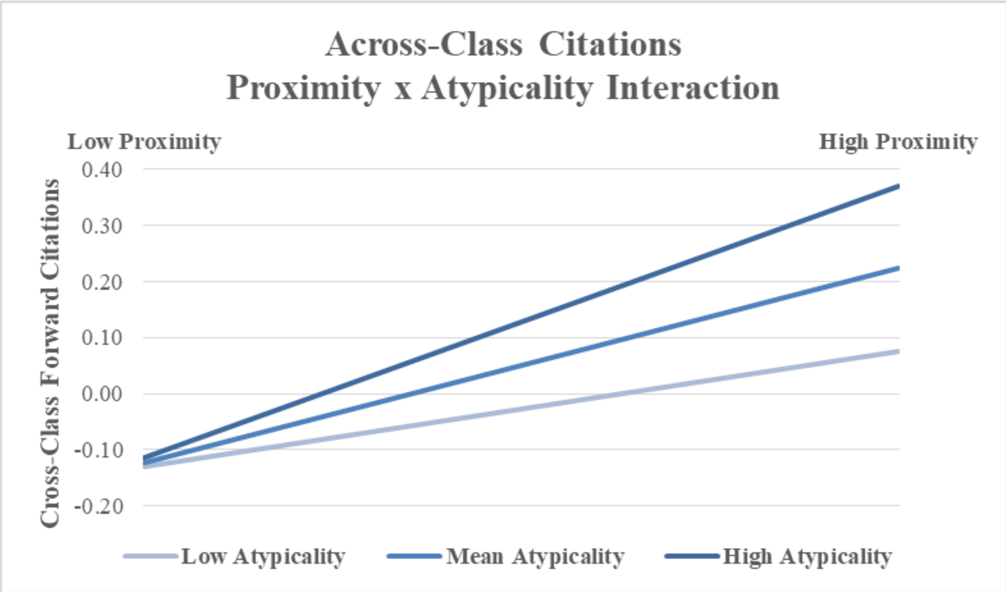

Figure 3. All Citations – Proximity and Atypicality Interaction