



Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences

Melissa A. Schilling*, Elad Green

Stern School of Business, New York University, 40 West 4th Street, New York, NY 10012, United States

ARTICLE INFO

Article history:

Received 16 February 2009
 Received in revised form 16 June 2011
 Accepted 16 June 2011
 Available online 19 July 2011

Keywords:

Innovation
 Knowledge networks
 Breakthrough
 Recombinant search
 Small-world networks
 Insight

ABSTRACT

Some ideas have dramatically more impact than others – they may overturn existing paradigms or launch new areas of scientific inquiry. Where do such high impact ideas come from? Are some search processes significantly more likely to lead to breakthrough idea generation than others? In this research, we compare “high impact” papers from the social sciences with random-but-matched articles published in the same journals in the same years. We find that search scope, search depth, and atypical connections between different research domains significantly increase a paper’s impact, even when controlling for the experience and prior publishing success of the author(s).

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Creativity and innovation have a powerful influence on organizational performance and economic growth. In many industries, innovation is now the most important driver of competitive success. As a result, a growing body of research has attempted to better understand the drivers of creativity and innovation in the workplace. A significant portion of that research has focused on either contextual factors that influence creativity and innovation such as incentives, social structure, and organizational culture (e.g., Amabile, 1983; Chatman et al., 1998; Fleming et al., 2007; Perry-Smith and Shalley, 2003) or individual attributes such as personality or skills (e.g., Baron, 1969; MacKinnon, 1965; Martindale, 1989; Oldham and Cummings, 1996; Root-Bernstein, 1989; Zhou and Oldham, 2001). A recent line of inquiry, however, focuses more explicitly on the mechanism by which individuals and organizations achieve creative outcomes: recombinant search (e.g., Fleming, 2001; Gavetti and Levinthal, 2000; Katila and Ahuja’s, 2002). The current study builds on this line of work by examining whether there are systematic differences in the search processes that lead to breakthrough ideas versus more incremental ones.

Some ideas have disproportionately more impact than others. They may launch new areas of scientific inquiry or rapidly accelerate an existing trajectory of research. These ideas, termed “high

impact ideas” here, are identifiable by the substantially higher rate at which they are cited in future work. Examples include those that fundamentally changed many areas of basic science, such as Einstein’s theory of relativity and Watson and Crick’s discovery of the double helix structure in DNA, to those whose impact was felt primarily in more well-defined scientific communities, such as Edward Said’s “Orientalism” in literary theory, or Kahneman and Tversky’s “prospect theory” in psychology, economics and management. While not all high impact ideas turn out to be correct or social-welfare enhancing, it is not hard to argue that high impact ideas have, on average, been very important for scientific and social progress.

Do the search processes that give rise to exceptionally high impact ideas differ systematically from the search processes underlying more incremental work? Some research has suggested that ideas are more likely to be high impact when they are the result of a successful connection forged between seemingly disparate bodies of knowledge (e.g., Simonton, 1995, 1999a; Schilling, 2005). Simonton (1995, 1999a,b), for example, pointed out that many of the most famous scientific breakthroughs were the result of seemingly random connections that occurred through a free associative process (what Freudians might call “primary process thinking”). In this process, an individual generates many unusual combinations between different bodies of knowledge possessed by the individual, and subjects that set to a screening process of selective retention, keeping only the best variations (much like Darwinian evolution). This view echoes the eloquent description by William James (1890:456) “Instead of thoughts of concrete things patiently following one another in a beaten track of habitual suggestion, we

* Corresponding author. Tel.: +1 212 998 0249.

E-mail addresses: mschilli@stern.nyu.edu (M.A. Schilling), elad.green@stern.nyu.edu (E. Green).

have the most abrupt cross-cuts and transitions from one idea to another, the most rarefied abstractions and discriminations, the most unheard of combination of elements, the subtlest associations of analogy; in a word, we seem suddenly introduced into a seething cauldron of ideas, where everything is fizzling and bobbling about a state of bewildering activity, where partnerships can be joined or loosened in an instant, treadmill routine is unknown, and the unexpected seems only law." The work on recombinant search lends support to this position by noting that unfamiliar or atypical combinations of knowledge yield novel outcomes with greater variance in performance. This variance can lower the average performance of outcomes, but also enables the possibility of exceptionally high performing outcomes (Fleming, 2001; Fleming and Sorenson, 2004; Katila and Ahuja's, 2002; Nelson and Winter, 1982). This line of work suggests that search scope improves the likelihood of generating a high impact idea by unleashing greater recombinatorial possibilities.

Other research, however, has emphasized the importance of search depth. Authors in this tradition typically argue that an individual's best hope of contributing a high impact idea is through attaining a deep and narrow knowledge reservoir, accrued through years of specialization. For example, some authors have observed that individuals often require at least a decade of intense study in a particular domain of knowledge prior to making a significant contribution in that domain (Gardner, 1993; Hayes, 1989; Simonton, 1999a,b). Simon and Chase even quantified this expertise by studying chess grand masters and other experts, concluding that individuals need approximately 50,000 "chunks" of richly connected information prior to making a fruitful discovery (Simon and Chase, 1973). Through intensely focused study, it is argued, these individuals may achieve an exceptional level of understanding of an area that enables breakthrough revelations.

The topic of breakthrough idea generation has garnered a significant amount of attention in both psychology and management, yet much still remains to be understood about how it occurs. In psychology, where breakthrough idea generation is studied as cognitive insight, there are a myriad of questions about what it is, why it is often accompanied by an affective response (the "Aha!") and how it might occur (e.g., Mayer, 1995; Davidson, 1995; Kaplan and Simon, 1990; Martindale, 1995; Metcalfe, 1986; Metcalfe and Weib, 1987; Simonton, 1999a,b). As Metcalfe (1995) notes: "The persistent lack of a mechanism for insight, linked with the charge that the notion of insight is somehow supernatural, has shackled researchers who would explore this most important of cognitive processes . . . We do not yet understand insight." In the management literature, breakthrough idea generation has received somewhat less attention. Instead, researchers have focused more on exploring whether search scope or search depth are more important for knowledge creation, and how individuals or firms can balance the apparent trade-off between them.

In this paper, we attempt to disentangle the seemingly competing arguments about search scope and depth by considering much more explicitly how search scope and search depth influence the structure of knowledge networks within the mind. We integrate the work in management on recombinant search with work in psychology on cognitive insight to build a set of arguments about how search shapes and refines semantic networks. This synthesis suggests that breakthrough idea generation is likely to be the result of bridging deep pools of knowledge with an atypical connection. This is an important conclusion, because when a network is sparse and highly clustered (as semantic networks have been shown to be – see Steyvers and Tenenbaum, 2005), it takes only a very small percentage of random or atypical connections to dramatically decrease the path length between nodes in the network (Watts and Strogatz, 1998). This suggests that the trade-off typically assumed to exist between search depth and search breadth

might be overstated; minor amounts of search breadth may have large payoffs.

We examine these ideas empirically in the context of social science articles. Using a novel coding and analysis method, we assess whether search scope, depth, or atypical connections in a social science article are related to its relative impact, while controlling for such factors as the field, year, and journal of publication, and the prior experience and publishing success of the authors. To do this, we compare a set of "high impact" articles and a control group of articles published in the same journals and same years. Both sets of articles are published in the same top tier social science journals, thus this is not an examination of how search processes differ between winning ideas and losing ideas, but rather how search processes differ between good ideas and *exceptionally high impact* ideas.

In the next section, we first explain why knowledge creation is fundamentally a process of building and refining associational networks. We then integrate research on search from the management literature with work on creativity in the psychology literature to generate arguments about the roles of search depth, search scope, atypical connections, and prior knowledge. We then describe the data and methods for testing these ideas, followed by the results. In the final section, we summarize the findings and discuss their implications for research and practice.

2. Knowledge networks and recombinant search

Knowledge is more than mere information. Information refers to a signal (Dretske, 1981) or a flow of messages (Nonaka, 1994) that only becomes *knowledge* when and if it is integrated within a pattern of associations that give it meaning (Bartlett, 1932; Mayer and Greeno, 1972). The network of patterns within which the information is embedded structures how that information is understood, and how that information relates to what is believed to be true (Nonaka, 1994). In absence of such connections between bits of data, the data would be meaningless – it would not be knowledge per se. Knowledge is thus literally a network.

2.1. Knowledge networks

Knowledge networks exist, and knowledge creation occurs, at many different levels of analysis. At the group and inter-organizational network levels, the term "knowledge networks" has come to refer to the distribution of knowledge in a network of nodes and links wherein nodes may be individuals, groups, or organizations, and links are the methods by which they share or exchange information and knowledge (Hansen, 2002; Monge and Contractor, 2003). We focus here on the individual cognition level, though many of the fundamental processes have analogs at other levels of analysis.

At the individual level, the knowledge network refers to the pattern of associations between concepts in the mind. Early models of cognition proposed that concepts were organized in the memory via chains or hierarchies of association (Medin et al., 2002). More recent models of cognition have taken an explicit network approach; some of the most well known include semantic memory models and parallel distributed processing (PDP) models. Most semantic memory models graphically depict the organization of memory by showing concepts as *nodes*, and relations as *links* (Anderson and Bower, 1973; Collins and Quillian, 1969). Concepts are connected by some relationship, and the connections provide inference about the nature of that relationship. PDP models (or "connectionist models") also take a network approach, but attempt to achieve a closer structural coherence with the physiology of the brain (Medin et al., 2002). To more closely approximate how cognition occurs in the mind, PDP models suggest that concepts

are represented in patterns of activity over the same set of nodes, rather than in particular nodes. As an individual learns, the strength of connection between nodes is adjusted, causing a stimulus to prompt the activation of a modified set of nodes. The associations between concepts give them meaning; thus the knowledge contained in the network is defined by the network's structure (Bates and Elman, 2002).

Knowledge creation occurs when new information is integrated within the network, or when the existing information within the network is recombined in new ways (Schilling and Phelps, 2007). A long line of research emphasizes the latter method, suggesting that the creation of new knowledge is most often the result of novel recombinations of known elements of knowledge, problems, or solutions (Gilfillan, 1935; Nelson and Winter, 1982; Usher, 1954) or the reconfiguration of the ways in which knowledge elements are linked (Henderson and Clark, 1990). Both the amount and diversity of new information integrated into the network (or degree of change in the way existing knowledge is combined) will influence how novel the new knowledge is (Fleming, 2001; Katila and Ahuja's, 2002; March, 1991).

Knowledge creation is also influenced by the existing composition and structure of the knowledge network (Schilling and Phelps, 2007). The knowledge elements that already reside in the network, and the pattern of associations among them, strongly condition how new information is integrated into the network and/or how new combinations among existing information occur and are understood. For instance, psychologists have found that varying the sequence of instruction can induce individuals to assimilate new knowledge to different schemata: "...new learning involves the development of cognitive structure that results from relating new ideas and accommodating existing structures... since the outcome of learning is jointly determined by the new material and the structure to which it is assimilated, the use of different procedures could lead to the development of markedly different structures during the learning of the same new concept" (Mayer and Greeno, 1972). Thus the network both enables information to become knowledge, and determines the nature of that knowledge.

In sum, knowledge creation is the integration of new information into the knowledge network, and/or the recombination of existing knowledge within the network. Additionally, knowledge created through either process is shaped by the structure and content of the existing knowledge network. Thus knowledge creation is a function of new information, existing information, and the existing network.

2.2. Recombinant search

The process of exploring different potential solutions to a problem, including the identification of new knowledge elements or new relationships between knowledge elements, may be termed recombinant search. "Search" refers to "an act of scrutiny, inquiry or examination in an attempt to find something, gain knowledge, etc." (Webster's). In the process of deliberate search, individuals often seek external information to augment or refine their existing knowledge networks. However, at the same time, the individual may consciously or subconsciously seek new associations between information or ideas already possessed. The ability for the mind to subconsciously explore new combinations between known information or ideas may be precisely why the "Aha!" moment of insight often occurs after a period of incubation during which time the individual may not perceive themselves to be working on the problem (Dorfman et al., 1996; Wallas, 1926).

Much of the research on recombinant search implicitly or explicitly invokes a spatial metaphor. Reference is made to searching in the "neighborhood" of past solutions (Cyert and March, 1963) or practicing "local" or "distant" search (Nelson and Winter, 1982).

The term local search is typically meant to convey when an actor searches deeply but with low scope. The inverse – low depth and high scope – is indicative of distant, or exploratory, search. Though much of the research on local and distant search poses these as choices on two different ends of a continuum, recent work described below suggests that breakthrough idea generation may be most likely when individuals are able to combine the benefits of both the deep knowledge reservoirs achieved through search depth and the unexpected connections of exploratory search (Katila, 2000; Schilling, 2005).

2.2.1. Search depth

Search *depth* refers to the extensiveness of search within a given knowledge area. The deeper the search effort, the more cumulative experience an actor accrues in a domain of knowledge and the greater their competence in that area (Katila, 2000). The research on local search suggests that most individuals are predisposed to search for solutions in areas in which they have existing knowledge (Dosi, 1988; Helfat, 1994; March, 1991; Martin and Mitchell, 1998; Nelson and Winter, 1982; Patel and Pavitt, 1997; Stuart and Podolny, 1996). Bounded rationality of individuals limits their ability to search all possible domains of knowledge (Simon, 1978) and biases them toward more salient areas of their own prior experience (Cyert and March, 1963). The more knowledge an individual has in a particular domain, the more likely they are to understand the nature of the relationships between different ideas. As associations are challenged or reinforced over time, the more accurate the pattern of associations should become, and the more efficient the individual should be in searching for a solution among them (Dosi, 1988; Harlow, 1959). As noted previously, a well-established line of research in psychology maintains that only through developing a deep reservoir of knowledge in an area is an individual likely to make a meaningful impact.

On the other hand, some authors have also argued that repeatedly searching within the same domain of knowledge is likely to result only in incremental advancements, while decreasing the likelihood of highly novel or "radical" solutions (Fleming and Sorenson, 2004; Katila, 2000; Mezias and Glynn, 1993; Schilling, 2005). The continuous exploitation of familiar bodies of knowledge can exhaust their potential as sources of novel solutions as all possible combinations of knowledge elements are eventually achieved (Fleming, 2001; Kim and Kogut, 1996). Furthermore, as an individual becomes highly specialized in a knowledge domain, they can become prone to "Einstellung," whereby learners who have previously solved a problem a particular way will form a problem-solving set that mechanizes their problem solving, constraining them from developing creative solutions (Luchins, 1942; Mayer, 1995). Many forms of learning can become automatized such that when faced with a particular situation, the learner automatically recalls a representation, and it is difficult not to do so (Gick and Lockart, 1995). When an individual has well-reinforced expectations about the direction a search path should take, it constrains their ability to explore different possibilities, and may prevent them from generating "preinventive forms" with a more natural or universal structure (Finke, 1995: 262).

2.2.2. Search scope and the role of atypical connections

Search *scope* refers to the number or breadth of knowledge domains searched, and is often invoked to refer specifically to knowledge domains in which the actor lacks prior experience or competence (Fleming, 2001; Katila, 2000). A growing body of literature suggests that searching more distant domains of knowledge can help an individual to avoid becoming trapped in an inefficient local optima (Ahuja and Lampert, 2001; Gavetti and Levinthal, 2000; Levinthal and March, 1993; Perkins, 1995). Recombinatory search that is high in scope can enhance creativity and innovation

through at least two mechanisms. First, an increase in search scope increases the number of knowledge elements available for recombination (Fleming, 2001; Simonton, 1995). All else being equal, the larger the set of knowledge elements searched, the greater are their combinatorial possibilities (Fleming and Sorenson, 2001). Second, search scope increases the variance in the outcomes of search (Fleming, 2001; March, 1991). The “value of variance” (Mezias and Glynn, 1993) in search is that while it increases the number of failures, it also increases the number of highly novel or radical solutions to be realized (Levinthal and March, 1981; March, 1991). In essence, it is because a search of broad scope is more likely to lead to atypical connections that it enables breakthrough idea generation (Schilling, 2005; Simonton, 1999a).

Search scope can also influence knowledge creation in other ways. First, searching diverse knowledge domains provides actors with multiple, varied interpretations of an extant problem. Such heterogeneity of meaning can lead to useful reconceptualizations of problems (Huber, 1991; Kaplan and Simon, 1990) and increase the actor's ability to integrate the novel knowledge into their existing knowledge base (Cohen and Levinthal, 1990; Nonaka, 1994). Access to alternative perspectives regarding problems and solutions can help actors to apply solutions from one domain to problems in another, a process known as analogical transfer (Gentner and Gentner, 1983; Holyoak, 1984; Loewenstein et al., 1999; Schilling et al., 2003). Second, the effort required to make sense of and understand these diverse knowledge domains challenges the stability of existing cognitive structures and cause-effect relationships (cf. Ahuja and Lampert, 2001). This can stimulate a critical self-examination of one's own cognitive structures leading to “second-loop” learning (Argyris and Schon, 1978). Such efforts at incorporating diverse knowledge elements into one's current knowledge structures can facilitate the construction of novel linkages and associations among them, resulting in highly novel insights and solutions (Cohen and Levinthal, 1990; Simonton, 1999a,b).

As suggested previously, search that is high in scope can be costly (Cyert and March, 1963; Kauffman et al., 2000; Nelson and Winter, 1982). In contrast to the efficiencies gained from experience in local search, high scope search is more uncertain, difficult, and less successful on average (Fleming, 2001; Nelson and Winter, 1982). When actors look for solutions in knowledge domains in which they lack prior experience and expertise, they must expend greater effort and resources to understand and integrate this diverse knowledge. Furthermore, as the number and variety of knowledge elements searched increases, the ability of actors to attend to and comprehend their interactions is diminished due to limited cognitive capacity (Fleming and Sorenson, 2001; Simon, 1978). The uncertainty associated with the outcomes of search also tends to rise with the scope of search. The recombination of well-understood knowledge components is more certain than that of relatively novel knowledge elements due to the benefits of previous experience – through experience, actors gain better understanding of which elements to recombine and which to avoid and what combinations are better than others for certain problems and contexts (Nelson and Winter, 1982).

2.2.3. Combining deep knowledge reservoirs with atypical connections

The depth and scope of search are often posed as different ends of the same continuum, as if investing in or utilizing more of one comes at the expense of the other. This unidimensional conception presupposes, however, that the amount of search individuals can conduct is fixed, and thus more units allocated to depth (i.e., within a field) leads to less units allocated to scope (i.e., across fields). This assumption is, of course, incorrect. Individuals vary in their abilities, experience, motivation and other factors that influ-

ence the amount of effort they will exert toward search, and the scope of knowledge they will attempt to apply to a given problem. It is altogether possible, for example, for one individual to search both deeper *and* broader than another. Depth and scope are thus separate, though interdependent, dimensions (Katila and Ahuja's, 2002).

A small body of research suggests that search depth and scope may positively interact, enabling increases in knowledge creation that are not possible when search efforts focus on one dimension or the other. Scope provides for the identification of potentially promising domains of knowledge that can then be further investigated, fine-tuned, and understood through search depth (Gavetti and Levinthal, 2000). Such fine-tuning increases the individual's understanding of the new knowledge domain and the ease to which it can be integrated and combined with the actor's established knowledge stock. Thus, the iterative use of depth and scope might lead to more solutions, and solutions of greater novelty, than the focused use of either mode of search (Katila, 2000; Katila and Ahuja, 2002).

A closer examination of the micro-mechanisms underlying knowledge creation, however, suggests an even more nuanced understanding of breakthrough idea generation. As noted previously, knowledge creation is the process of integrating new information into one's knowledge network or recombining existing information in the knowledge network in new ways. The knowledge elements in the mind (including both ideas and concepts) are not randomly connected to one another, but rather are highly structured (Anderson and Hinton, 1989; Steyvers and Tenenbaum, 2005). The likelihood of two ideas being associated together is a probability of some function of their similarity on one or more dimensions (sometimes termed “semantic distance” or “semantic relatedness”) (Collins and Loftus, 1975; Rips et al., 1973). Association based on similarity results in significant clustering. Further, such networks are likely to be sparse. Forging and maintaining links between concepts in the mind has a cost in terms of time and effort (Simon, 1955), and links that are not reinforced over time can diminish (Martindale, 1995). These costs make it difficult (if not impossible) to densely connect every possible node in the network to every other node; instead cognitive networks are likely to be characterized by dense connectivity among closely related nodes, and much sparser connectivity (if any) between nodes that are only distantly related. Though such order and clustering is extremely valuable in terms of giving structure and meaning to individual knowledge nodes and sets of knowledge nodes (Bartlett, 1932; Mayer and Greeno, 1972), it also results in relatively long path lengths in the network. Long path lengths make it more difficult and time consuming for an individual to search their cognitive network, and may make the individual less likely to find a solution that is not in the immediate domain of the problem.

Atypical connections can create a shortcut in the knowledge network that results in a dramatic shift in the individual's knowledge network. Formerly distant ideas may be brought into close proximity, simultaneously reorienting the individual's perception of distance between other elements that are associated with these ideas. The dramatic decrease in path length between formerly distant representations may prompt the individual to search for and note other similarities. Relationships that had never been previously considered might suddenly seem obvious, causing the rapid formation of new links between the representations without any prompting from external input, resulting in a cascade of node and link changes.¹ Notably, the perceived significance of the shift in the

¹ This may be the dynamic underlying the affective response that characterizes cognitive insight – the “Aha!” feeling when sudden clarity emerges (Dominowski and Dallob, 1995; Gick and Lockart, 1995; Schilling, 2005).

individual's knowledge network is a function of both the unexpectedness of the connection, and the magnitude of change it creates in the network of representations. The latter suggests that when atypical connections create shortcuts between deep or large knowledge reservoirs, the resulting shift in the individual's cognitive network is likely to be especially dramatic.

An intriguing implication of this argument is that it may only require one (or a few) atypical connections to dramatically reorient a network of representations. As demonstrated in the work on small-world networks, when a network is sparse and highly clustered, it only takes a very small percentage of random or atypical links to cause a phase transition in connectivity of the network, bringing the nodes of the network significantly closer (Watts and Strogatz, 1998), and potentially leading to the cascade of associations noted to characterize cognitive insight. This has important implications for the tension between exploration and exploitation. The trade-off that is typically assumed to exist between search depth and search breadth might be overstated – only minor amounts of search breadth might be required to achieve the large payoffs of exploration if the individual has good reasoning or intuition about areas to explore.²

2.3. Current search versus prior search

When an individual searches for a solution to a particular problem, their search will be guided by a number of factors that include (but are not limited to) what others have done in the area, what the individual perceives to be most relevant to the problem at hand, what knowledge they possess as a result of previous searches, and their natural search tendencies. The search depth, search scope, or atypical connections manifest in a given paper is thus to some degree due to the current search effort, and is to some degree an artifact of the searches the individual has done before. If an individual already has deep expertise in a particular domain that is relevant to the current search problem, they are very likely to draw from those knowledge repositories for the current search effort. Consistent with this, one observes that authors use many of the same cites across their papers – some cites become almost reflexive in nature. However, it is also possible for individuals to draw from deeper or broader knowledge reservoirs than what is reflected in the paper's references, which introduces some noise into the measures used here (as discussed at greater length in Section 5).

Other things being equal, an author with extensive experience in the field (i.e., years of participating in the field and/or success in publishing) might have advantages in generating a high impact article. First they are likely to both have a larger knowledge stock upon which to draw, and be better able to assimilate and utilize the fruits of their searches because of learning transfer efficiencies (Ellis, 1965). Second, an experienced author might be better at selecting research topics that have high impact potential (thus better allocating their search efforts), be better at framing the articles in such a way that their potential is realized by their readership, or in proactively spreading and legitimizing their ideas (Birkinshaw and Mol, 2006; Hargadon, 2002). Third, prior experience and publishing success may act as a signaling and legitimization device that serves to increase the likelihood of others reading and citing the work. Finally, the fourth mechanism is the result of a selection effect: Only individuals with some degree of publishing success are likely to be retained in a field that emphasizes publishing as a key outcome – thus individuals with prior publishing success get more tries at producing a high impact idea. Though one can bring to

mind examples of “one hit wonders” (an expression borrowed from the music industry, where it connotes a musical artist known for only one hit single and short-lived fame), and in mathematics it is often argued that success comes early or not at all, in the social sciences it is uncommon for inexperienced authors to produce heavily cited works. Notably, however, there is recent evidence to suggest that an individual's prior success can also diminish their likelihood of generating a breakthrough idea. Like the arguments made previously about the tendency for individuals to exploit existing competencies, prior success in an area can induce individuals to rely on familiar knowledge and routines, leading to increasingly incremental ideas (Audia and Goncalo, 2007). We thus control for the author's (s') longevity in the field, and prior publishing success.

3. Methods

To test the hypotheses above, we used data from Thomson Scientific's High Impact Papers database, novel measures of search scope and depth based on the Dewey decimal system, and bibliometric data gathered on the authors, as described below.

3.1. Data

Thomson Scientific's High Impact Paper's database ranks the most influential papers in specific fields of science and social science based on their citation counts. The database includes the 200 most cited papers of each year since 1981, and additionally provides metrics for the “expected” number of citations for an article published in the same journal of the same year, as well as bibliographic information. For the current study, the ten highest impact papers for each of the following four disciplines were chosen, for a total of 40 papers: Economics, Management, Psychology, and Sociology. To provide a control group, for each of these articles, two other articles were randomly chosen from the same journal, published in the same year (termed “random-but-matched” articles from this point forward), resulting in a final sample of 120 articles. Because the bibliographic data collection and coding for each article is quite burdensome, this sample size was chosen to economize on data handling while still providing enough degrees of freedom for statistical testing.

For each article, we created measures of search scope, search depth, and atypical connections based on coding of the references used in the articles, and we also collected bibliographic data on the authors, as described below. Before describing the measures, however, it is useful to first present an overview of how the Dewey decimal system works.

3.2. The Dewey decimal system

The Dewey decimal system was created by Melvil Dewey in 1876 as a structure for the organization of library collections, and it has been continuously revised to meet the evolving needs of libraries and electronic information access environments. The Online Computer Library Center (OCLC) has owned and maintained the Dewey decimal system since 1988. It is the world's most widely used system for the classification of library materials by topic.

The system is constructed as a hierarchical classification scheme with three primary levels. The highest level divides library materials into ten main classes representing the major disciplines: 000 Generalities; 100 Philosophy & psychology; 200 Religion; 300 Social sciences; 400 Language; 500 Natural sciences & mathematics; 600 Technology (Applied sciences); 700 The arts; 800 Literature & rhetoric; and 900 Geography & history. At the next level, each of these main classes is divided into finer-grained categories known as the “Hundred Divisions.” For example, the “Hundred Divisions” for Social Sciences is divided into the following: 300 Social Science,

² For an example of what we mean by “reasoning” here, Fleming and Sorenson note that scientific theories may give inventors foresight about potentially fruitful areas of exploration, improving their likelihood of benefiting from distant search.

310 Collections of general statistics; 320 Political science; 330 Economics; 340 Law; 350 Public administration & military science; 360 Social problems & service, association; 370 Education; 380 Commerce, communications, transportation; and 390 Customs, etiquette, folklore. Each of the Hundred Divisions, is in turn, divided into even finer categories. For example, the “Thousand Sections” for Economics includes 330 Economics, 331 Labor economics, 332 Financial economics, 333 Economics of land & energy, and so forth.

Individual library materials, such as academic journals, are assigned a Dewey decimal classification number that is at least three digits long per the divisions discussed above, but may also be given a classification number that is even longer by adding digits after a decimal, enabling the material to be assigned to an even more narrowly defined concept. For example, the American Economic Review has a Dewey decimal classification number of 330.973. The thousand division of 330 refers to economics in general, while the 9 after the decimal point indicates the topic “Economic situations and conditions,” and the 73 following is a place number that indicates the US.

The Dewey decimal system is the most widely used foundation for development of knowledge taxonomies and retrieval tools in the world (Saeed and Chaudhry, 2001, 2002; Thompson et al., 1997; Weigand, 1998). Research that has examined how well the Dewey system partitions information into discrete classes has concluded that the Dewey concept definitions have a high degree of class integrity (i.e., the concepts are well-defined and disjoint) (Thompson et al., 1997).

We were able to find Dewey decimal numbers in the OCLC Worldcat system for 5318 (i.e., 79%) of the references cited by the articles in the sample. The remaining references were mostly unpublished works or papers in publications such as conference proceedings that were not identifiable in Worldcat.

3.3. Dependent variables

Two different dependent variables are compared here. The first is a dummy variable (0,1), *high impact*, indicating whether the article was drawn from the “high impact” paper database. This dependent variable is analyzed using logistic analysis. The second is a continuous variable, *relative impact*, that is created by dividing the actual citation count of the paper by the “expected citation” of the paper based on the journal and year in which it was published. The measure of “expected citation” is provided in the Thomson High Impact Papers database, and because the papers in the control group are drawn from the same journals and years as the high impact papers, this expected citation score can be used for the random-but-matched papers as well. The advantage of this measure is that it takes account of the fact that there is substantial variation in impact even among the high impact papers. Relative impact was natural log transformed to improve its normality.

3.4. Independent variables

3.4.1. Search depth

The depth of any hole is a function of both how many shovelfuls of earth are removed, and the area over which one’s effort is applied. Similarly, to measure depth of search, we multiply the number of references by the concentration of those references within (a) bibliographic field(s). First we calculate the concentration of references. The concentration measure is based on a Herfindahl-Hirschman Index (HHI) measure. It is calculated as the sum of the squared percentage of references that are in each Dewey decimal hundred division represented in the references. Thus if an article cites ten references, and seven of those ten (70%) are in one hundred division and three of those ten (30%) are in another, the article’s reference concentration would be $.70^2 + .30^2$ which equals .58. This measure

achieves its highest value when all of the articles cited by the paper are in the same hundred division, or 1^2 which equals 1.³ The concentration measure is multiplied by a count of the total number of references in the article to yield a measure of depth. Thus if two articles in a discipline have the same degree of concentration of references but one has many more references, it will score higher on depth. Similarly, if two articles in a discipline have the same number of references but one is significantly more concentrated, it will score higher on depth. The depth measure was log transformed to improve its normality.

3.4.2. Search scope

To capture the scope of literary terrain covered in an article, we created a concentric weighted count of Dewey decimal classifications represented in the references. This measure takes into account that journals in different main classes, for example, are more different from each other than journals that are in the same main class. The scope score is calculated as follows: each unique main class (the highest level) represented in the references is worth three points, each unique hundred division is worth two points, and each unique thousand section is worth one point. For example, if an article cited five articles whose journals had the following Dewey decimal classifications, 658.05, 305, 338.7, 616.8, and 658.5, the article would receive a scope score of 18 (6 points for having two different main classes, 600 and 300; 8 points for having four different hundred divisions, 650, 300, 330, and 610; and 4 points for having four different thousand sections, 658, 305, 338, 616). This measure was log transformed to improve its normality.

Referencing norms vary significantly across disciplines and journals. Psychology articles, for example, tend to cite far more references than economics articles. Furthermore, the tendency to cite across multiple literary areas may vary both with discipline norms, and with where a discipline falls within the Dewey decimal classification system. For example, whereas economics, sociology, and psychology each have their own hundred division (and thus an article citing only economics papers could still cite publications from a number of different categories at the thousand section level), management is a thousand section category (658), so an article that cites only management articles would only be citing publications from the same thousand section category, making it appear to have far less scope than an economics article that cited only other economics articles.⁴ These differences (and others) across fields are managed by controlling for the journal in which each article is published with the journal-year dummy variables.

3.4.3. Atypical connections

To measure atypical connections, we first calculated the number of times each possible pair of Dewey Decimal codes appeared together in an article. We then normalized this count by the maximum number of articles in which either code appeared, as Dewey Decimal codes that appear in more articles are at more risk of co-appearing with other Dewey codes. This gave us a probability index of any pair of Dewey codes being co-cited. Then, for each article,

³ Though intuitively concentration would seem to be the opposite of scope, the measures are not directly related in this way. It is altogether possible for two articles to have the same scope, but have very different concentrations and vice versa. For a simple example, consider one article that has ten references, eight of which are in the classification 658, and the remaining two are in the classifications 305 and 338, versus another article that also has ten references, but five of the references are in 658, three are in 305, and two are in 338. Both articles have the same scope score of 15, but the first has a concentration score of .66 and the second has a concentration score of .38. The first article’s references are thus significantly more concentrated than the second article’s references.

⁴ This example is an extreme; there were no articles in the sample that cited only management articles.

Table 1
Descriptives and correlations.

| Variable | Mean | St Dev | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
|----------------------------------|-------|--------|------------------|------------------|-------|------------------|-----|-------|-------|-------|-------|
| 1. High impact | .33 | .47 | | | | | | | | | |
| 2. Relative impact | 1.13 | 1.09 | .92** | | | | | | | | |
| 3. Years | 12.48 | 9.48 | .39** | .41** | | | | | | | |
| 4. Prior cites | 4.61 | 2.33 | .41** | .43** | .62** | | | | | | |
| 5. Author count | 1.57 | .69 | .11 | .13 | .25** | .38** | | | | | |
| 6. Depth | 2.87 | .72 | .19 [†] | .23 [†] | .15 | .01 | .06 | | | | |
| 7. Scope | 3.33 | .57 | .32** | .31** | .29** | .21 [†] | .03 | .34** | | | |
| 8. Depth × scope | 9.69 | 3.29 | .31** | .32** | .25** | .12 | .07 | .89** | .71** | | |
| 9. Atypical connections | .62 | .49 | .34** | .32** | .30** | .22 [†] | .05 | .34** | .61** | .53** | |
| 10. Atypical connections × depth | 1.88 | 1.57 | .35** | .32** | .29** | .21 [†] | .05 | .54** | .62** | .70** | .68** |

[†] $p < .05$.

** $p < .01$ (two-tailed test).

we created a list of every pair of Dewey codes that co-appear in the article, and their corresponding co-citation probability. We then identified the 5% least probable connections made in each of the four social science fields examined, and used a dummy variable (0,1) to indicate if a given article included one of the improbable co-citations (“atypical connections”). Articles with a “1” have one or more atypical connections; articles with a “0” do not.⁵

3.4.4. Prior experience and publishing success of authors

As argued previously, there are a number of ways that an author's prior experience and visibility might influence the likelihood of their article being high impact. We thus include the following measures: (1) *Years*, which captures the experience of the author, as the number of years between completing the terminal degree and publication of the article. We use the experience of the most experienced author when there are multiple authors. (2) *Prior cites*, which captures the author's publishing success by counting the number of citations to the author's prior work, accumulated up to the year of publication of the target article. We use the cites of the most highly cited author when there are multiple authors.⁶ The latter measure was natural log transformed to improve its normality. (3) *Author count*, controls for the number of authors on the paper.

3.4.5. Journal and year of publication

To control for significant differences in referencing norms across fields and journals, and for differences in length of time for which the article is at risk of being cited, dummy variables were included for $n - 1$ of the unique journal-year combinations represented in the sample.

4. Results

Descriptive statistics and correlations for the variables are shown in Table 1. As shown, there is a significant positive correlation between the impact measures, the years in the field, and the cites prior to publishing the articles examined here. Not surprisingly, prior experience and prior publishing success have a strong positive correlation with both a paper's likelihood of being in the high impact set, and its relative impact score. The impact measures are also significantly and positively related to the scope measure, the atypical connections measures, and interaction terms between depth and scope, and depth and atypical connections. Rel-

ative impact was also significantly and positively related to depth, though the dummy variable for high impact was not. It is worth noting that there is a positive and significant correlation between the depth and scope measures for the data examined here, consistent with arguments that scope and depth should not be considered two ends of a continuum. There is also a significant and positive correlation between scope and atypical connections, consistent with the arguments that scope can enable the formation of atypical connections.

Table 2 reports the logistic regression results for the binary dependent variable, “high impact.” We first test the more traditional depth versus scope arguments, and then turn to the “atypical connections” arguments made here. In Model 1, only the control variables for the journal-year and author attributes (*years*, *prior cites*, and *author count*) are entered. This model is able to correctly classify 79.2% percent of the articles. The first two author attribute variables, *years* and *prior cites*, are positive and significant, suggesting that authors' prior experience and publishing success significantly increases the likelihood of writing a high impact paper. The third author control, *author count*, is not significant, indicating that having more authors on a paper does not directly impact the likelihood of a paper being high impact. In Model 2, the depth and scope variables are added. Adding these variables significantly increased the Chi-squared statistic ($p < .01$ significance of the change), and increases the likelihood of the model correctly classifying the articles to 87.5%. The depth variable has a significant ($p < .05$) and positive coefficient, and its estimated odds ratio indicates that for every additional unit of search scope, the odds of an article being high impact increases by a factor of 4.95. The scope variable also has a significant and positive coefficient ($p < .01$), and its odds ratio indicates that for each additional unit of depth, the likelihood of an article being “high impact” is increased by a factor of 168.9. In Model 3, the interaction term *depth × scope* is entered. Adding this variable leads to a modest decrease in the log likelihood ratio, and the variable would be significant only with a two-tailed test, indicating evidence for a modest interaction effect. Model 4 includes squared terms for the depth and scope variables to explore the possibility of curvilinear effects. Only the squared scope variable comes close to achieving significance, and its coefficient is positive. This is surprising as it suggests an increasing, rather than diminishing, effect of scope.⁷

⁵ All three of the measures of “atypical connections” yield nearly identical results; we thus report here only the results with the dummy variable as this operationalization is most consistent with our arguments.

⁶ We also ran our models using a sum of the authors' years of experience when there are multiple authors, and the sum of the authors' prior cites. The results were nearly identical.

⁷ Though it was suggested that we center the variables prior to creating interaction and squared effects, research indicates that mean centering does not correct collinearity problems in moderated regression models (Echambadi and Hess, 2007) and is likely unnecessary given that modern statistical packages use double precision for all of the calculations, minimizing the potential for collinearity to cause problems. We did, however, explore the possibility of curvilinear effects through graphical analysis and SPSS curvefit estimation. We found no evidence suggestive of significant curvilinear effects.

Table 3
OLS regression of relative impact.^a

| Variables | Model 1 B (SE) | Model 2 B (SE) | Model 3 B (SE) | Model 4 B (SE) | Model 5 B (SE) | Model 6 B (SE) |
|------------------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
| Constant | .63 (.66) | −3.25** (.99) | −.12 (2.38) | −1.96 (2.68) | −.78 (.71) | −.79 (.88) |
| Years | .04** (.02) | .03† (.02) | .03† (.02) | .03* (.02) | .03 (.02) | .03 (.02) |
| Prior cites | .18** (.07) | .18** (.06) | .17** (.06) | .17** (.06) | .20** (.06) | .20** (.06) |
| Author count | −.06 (.18) | −.00 (.16) | −.02 (.16) | −.02 (.16) | −.04 (.17) | −.04 (.17) |
| Depth | | .55** (.16) | −.59 (.81) | 1.70† (.95) | .53* (.17) | .53* (.26) |
| Scope | | .82** (.24) | −.12 (.69) | −.99 (1.57) | | |
| Depth × scope | | | .34 (.24) | | | |
| Depth squared | | | | −.20 (.16) | | |
| Scope squared | | | | .27 (.24) | | |
| Atypical connections | | | | | .57* (.24) | .59 (.95) |
| Depth × atypical connections | | | | | | −.01 (.34) |
| Adj. R squared | .07 | .29 | .30 | .29 | .24 | .23 |
| F of change | 1.35 | 13.30** | 2.08 | 1.16 | 9.75** | .00 |

Coefficients for dummies for journal-year sets omitted.

† $p < .10$.

* $p < .05$.

** $p < .01$ (two-tailed test).

or finding a solution in a “homing space” wherein many clues lead to its almost inevitable discovery (Perkins, 1995), may not result in a solution that is perceived as novel. The associations may conform to the reader’s expectations, and may not significantly redirect inquiry in a way that causes the article to be seen as a seminal work. By contrast, when connections are made between ideas that had seemed unrelated or incongruent, the connection may be unexpected, and it may prompt readers to consider numerous other associations between, or applications of, ideas articulated in the paper (Schilling, 2005). One of the ways that authors may identify and forge relationships between disparate ideas is through search scope, though not all search scope will result in atypical connections.

The results indicated no evidence for interaction effects between search scope and depth, or between search scope and atypical connections. While at first blush this would seem to imply that an author could focus on search scope or atypical connections alone to increase the impact of their papers, it is important to remember that the results are comparing published works with exceptional published works, and that to be published, the articles probably must meet threshold levels of depth in their citations. Thus the failure to find a significant interaction effect should not be interpreted as indicating that an article could rely on high scope or atypical search processes alone.

The results also indicated no evidence for diminishing effects of search scope and search depth, suggesting that there is no penalty for searching extremely deeply or broadly (though again, journal norms may truncate the distribution of what is acceptable in terms of depth and breadth of referencing). Furthermore, the results indicated a positive effect for an author’s experience in the field and prior publishing success. Taken together, these results imply that for the social sciences at least, drawing from more knowledge, both deeply and broadly, improves the impact of a published work.

This study offers several contributions to the extant research. In addition to the main findings described above, it helps to clarify the relationships between search depth, search scope, and breakthrough idea generation. The results here indicate that search depth and search scope are positively correlated with one another, consistent with Katila and Ahuja, 2002 arguments that these two dimensions should not be considered two ends of a continuum: some individuals search both more deeply and more broadly. Furthermore, both have positive direct effects on breakthrough idea generation. This result obviates the debate about whether it is search depth or scope that facilitates breakthrough idea generation – it is both.

Another contribution of the study is the development of a novel coding system for measuring search scope, search depth, and atypical connections, that may be of interest to those who do research in this area. This coding system enables researchers to characterize the search process evidenced in a paper’s references using readily available public data. The Dewey decimal codes are based on well-defined and non-overlapping classes, helping to avoid some of the problems that emerge in using patent citations (another common measure of search scope and search depth).¹⁰ Future research may wish to compare the results that emerge from using both systems of measurement. For example, when an invention leads to both the publishing of a scientific article and a patent, do measures of search scope and search depth based on both article references and patent citations yield congruent patterns? The coding system could also be used to identify degrees of difference and commonality between articles, which could be useful for a number of research directions. This would, for example, enable us to identify periods of convergence and divergence in research streams, and may help us to identify articles that significantly redirected research in an area (akin to discontinuous innovation).

The study also has implications for other areas of future research. First, while the fundamental arguments of the paper were based on network dynamics of cognition, our measures only capture adjacencies within the sample (i.e., whether two Dewey codes are co-cited in papers). It would be fascinating (and valuable) to construct a method of more fully capturing the knowledge network dynamics of search and breakthrough idea generation. Second, though using the Dewey codes in references provides an interesting and novel way to trace an author’s search process, it is undoubtedly incomplete. Individuals are likely to draw from a deeper or broader knowledge reservoir than what is reflected in the paper’s references; we can only hope that the patterns we capture here are reflective of the patterns that exist in the true knowledge stocks used. This would also be an interesting area to develop new methods for more fully capturing an individual’s search process. Third, though the results indicated that search scope, search depth, and atypical connections can help to differentiate the exceptional published articles from more ordinary published articles, they do not address how these search processes impact the likelihood of an article being published at all. Future research could tackle this issue by comparing published articles with those that fail to achieve publica-

¹⁰ References share a disadvantage with patent citations, however, in that both might be partially driven by the strategic interests of the author(s).

tion. It would be interesting to see if Fleming's 2001 results could be replicated, i.e., while search scope can enable the breakthrough idea generation that leads to high impact papers, does it also increase the likelihood of failing to publish a paper?

Acknowledgements

This project would not have been possible without the ideas and assistance of Jonathan Weitz, who first proposed the idea of using the Dewey Decimal system to code article content, and helped gather the first wave of data. The authors would also like to gratefully acknowledge the suggestions and assistance of Heather Clark, Kathy Eisenhardt, Nile Hatch, and Joe Porac. An earlier version of the paper was presented at the 2010 Winter Strategy Conference in Utah. This research was supported by the National Science Foundation under Grant No. SES-0234075.

References

- Ahuja, G., Lampert, C.M., 2001. Entrepreneurship in the large corporation: a longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal* 22, 521–543.
- Amabile, T.M., 1983. *The Social Psychology of Creativity*. Springer Verlag, New York.
- Anderson, J., Bower, G.H., 1973. *Human Associative Memory*. Halstead, New York.
- Anderson, J.A., Hinton, G.E., 1989. Models of information processing in the brain. In: Hinton, G.E., Anderson, J.A. (Eds.), *Parallel Models of Associative Memory*. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 23–64.
- Argyris, C., Schon, D., 1978. *Organisational Learning: A Theory of Action Perspective*. Addison Wesley, Reading, Mass.
- Audia, P.G., Goncalo, J.A., 2007. Past success and creativity over time: a study of inventors in the hard disk drive industry. *Management Science* 53, 1–15.
- Baron, F., 1969. *Creative Person and Creative Process*. Holt, Rinehart, Winston, New York.
- Bartlett, F.C., 1932. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, London.
- Bates, E.A., Elman, J.L., 2002. Connectionism and the study of change. In: Mark Johnson (Ed.), *Brain Development and Cognition: A Reader*, second ed. Blackwell Publishers, Oxford.
- Birkinshaw, J., Mol, M., 2006. How management innovation happens. *Sloan Management Review* 47 (4), 81–88.
- Chatman, J.A., Polzer, J.T., Barsade, S.G., Neale, M.A., 1998. Being different yet feeling similar: the influence of demographic composition and organizational culture on work processes and outcomes. *Administrative Science Quarterly* 43, 749–780.
- Cohen, W.M., Levinthal, D.A., 1990. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* 35, 128–152.
- Collins, A.M., Loftus, E.F., 1975. A spreading-activation theory of semantic processing. *Psychological Review* 82, 407–428.
- Collins, A.M., Quillian, M.R., 1969. Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior* 8, 240–247.
- Cyert, R., March, J., 1963. *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.
- Davidson, J.E., 1995. The suddenness of insight. In: Sternberg, R., Davidson, J. (Eds.), *The Nature of Insight*. MIT Press, Bradford, pp. 125–155.
- Dominowski, R.L., Dallob, P., 1995. Insight and problem solving. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. 33–62.
- Dorfman, J., Shames, V.A., Kihlstrom, J.F., 1996. Intuition, incubation, and insight: implicit cognition in problem solving. In: Underwood, G. (Ed.), *Implicit Cognition*. Oxford University Press, Oxford, pp. 257–285.
- Dosi, G., 1988. Sources, procedures, and microeconomic effects of innovation. *Journal of Economic Literature* 26, 1120–1171.
- Dretske, E., 1981. *Knowledge and the Flow of Information*. MIT Press, Cambridge, MA.
- Echambadi, R., Hess, J.D., 2007. Mean-centering does not alleviate collinearity in moderated regression models. *Marketing Science* 26, 438–445.
- Ellis, H., 1965. *The Transfer of Learning*. MacMillan Company, New York.
- Finke, R.A., 1995. Creative insight and preinventive forms. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. 255–280.
- Fleming, L., 2001. Recombinant uncertainty in technological search. *Management Science* 47, 117–132.
- Fleming, L., Mingo, S., Chen, D., 2007. Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly* 52, 443–475.
- Fleming, L., Sorenson, O., 2004. Science as a map in technological search. *Strategic Management Journal* 25, 909–928.
- Gardner, H., 1993. *Creating minds: An Anatomy of Creativity seen Through the Lives of Freud, Einstein, Picasso, Stravinsky, Eliot, Graham, and Gandhi*. Basic Books, New York.
- Gavetti, G., Levinthal, D., 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly* 45, 113–137.
- Gentner, D., Gentner, D.R., 1983. Flowing waters or teeming crowds: mental models of electricity. In: Gentner, D., Stevens, A. (Eds.), *Mental Models*. Lawrence Erlbaum, Hillsdale, NJ, pp. 99–130.
- Gick, M.L., Lockart, R.S., 1995. Cognitive and affective components of insight. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. 197–228.
- Gilfillan, S.C., 1935. *The Sociology of Invention*. Follett, Chicago.
- Hansen, M.T., 2002. Knowledge networks: explaining effective knowledge sharing in multiunit companies. *Organization Science* 13, 232–250.
- Hargadon, A., 2002. Brokering knowledge: linking learning and innovation. *Research in Organizational Behaviour* 24, 41–85.
- Harlow, H.E., 1959. Learning set and error factor theory. In: Koch, S. (Ed.), *Psychology: A Study of a Science*, vol. 2. McGraw-Hill, New York, pp. 492–537.
- Hayes, J.R., 1989. *The Complete Problem Solver*, second ed. Lawrence Erlbaum, Hillsdale, NJ.
- Helfat, C., 1994. Evolutionary trajectories in petroleum firm R&D. *Management Science* 40, 1720–1747.
- Henderson, R., Clark, K., 1990. Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly* 35, 9–30.
- Holyoak, K.J., 1984. Analogical thinking and human intelligence. In: Sternberg, R.J. (Ed.), *Advances in the Psychology of Human Intelligence*, vol. 2. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 191–230.
- Huber, G.P., 1991. Organizational learning: the contributing processes and the literatures. *Organization Science* 2, 88–115.
- James, W., 1890. *Principles of Psychology*, vol. 1. Henry Holt, New York.
- Kaplan, C.A., Simon, H.A., 1990. In search of insight. *Cognitive Psychology* 22, 374–419.
- Kauffman, S., Lobo, J., Macready, W.G., 2000. Optimal search on a technology landscape. *Journal of Economic Behavior & Organization* 43 (2), 141–166.
- Katila, R., 2000. Innovation search determinants of new product introduction and their radicality: The case of industrial robotics. Unpublished doctoral dissertation. The University of Texas at Austin.
- Katila, R., Ahuja, G., 2002. Something old, something new: a longitudinal study of search behavior and new product introduction. *Academy of Management Journal* 45, 1183–1194.
- Kim, D.J., Kogut, B., 1996. Technological platforms and diversification. *Organization Science* 7, 283–301.
- Levinthal, D., March, J.G., 1981. A model of adaptive organizational search. *Journal of Economic Behavior and Organization* 2, 307–333.
- Levinthal, D., March, J.G., 1993. The myopia of learning. *Strategic Management Journal* 14, 95–112.
- Loewenstein, J., Thompson, L., Gentner, D., 1999. Analogical encoding facilitates transfer in negotiation. *Psychonomic Bulletin Review* 6, 586–597.
- Luchins, A., 1942. Mechanization in problem solving: the effect of Einstellung. *Psychological Monographs* 54 (6), whole no. 248.
- MacKinnon, D.W., 1965. Personality and the realization of creative potential. *American Psychologist* 17, 484–495.
- March, J.G., 1991. Exploration and exploitation in organizational learning. *Organization Science* 2, 71–87.
- Martin, X., Mitchell, W., 1998. The influence of local search and performance heuristics. *Research Policy* 26 (7/8), 753–772.
- Martindale, C., 1989. Personality, situation, and creativity. In: Glover, J., Ronning, R., Reynolds, C. (Eds.), *Handbook of Creativity*. Lawrence Erlbaum, Mahwah, NJ, pp. 279–301.
- Martindale, C., 1995. Creativity and connectionism. In: Smith, S.M., Ward, T.B., Finke, R.A. (Eds.), *The Creative Cognition Approach*. MIT Press, Cambridge, MA, pp. 249–268.
- Mayer, R.E., 1995. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. 3–32.
- Mayer, R.E., Greeno, J.G., 1972. Structural differences between learning outcomes produced by different instructional methods. *Journal of Educational Psychology* 63, 165–173.
- Medin, D.L., Ross, B.H., Markman, A.B., 2002. *Cognitive Psychology*. Wiley, John, & Sons, New York.
- Metcalfe, J., 1986. Feeling of knowing in memory and problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12, 288–294.
- Metcalfe, J., Weib, D., 1987. Intuition in insight and noninsight problem solving. *Memory & Cognition* 15, 238–246.
- Metcalfe, J., 1995. Foreword. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. ix–xiv.
- Mezias, S.J., Glynn, M.A., 1993. The three faces of corporate renewal: institution, revolution, and evolution. *Strategic Management Journal* 14, 77–101.
- Monge, P.R., Contractor, N., 2003. *Theories of Communications Networks*. Oxford University Press, New York.
- Nelson, R.R., Winter, S., 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Nonaka, I., 1994. A dynamic theory of organizational knowledge creation. *Organization Science* 5, 477–501.
- Oldham, G.R., Cummings, A., 1996. Employee creativity: personal and contextual factors at work. *Academy of Management Journal* 39, 607–634.
- Patel, H., Pavitt, K., 1997. The technological competency of the world's largest firms: complex path-dependent, but not much variety. *Research Policy* 26, 141–156.
- Perkins, D.N., 1995. Insight in minds and genes. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. 495–534.

- Perry-Smith, J.E., Shalley, C.E., 2003. The social side of creativity: a static and dynamic social network perspective. *Academy of Management Review* 28, 89–106.
- Rips, L.J., Shoben, E.J., Smith, E.E., 1973. Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior* 12, 1–20.
- Root-Bernstein, R.S., 1989. Who discovers and who invents. *Research Technology Management* 32 (1), 43–51.
- Saeed, H., Chaudhry, A.S., 2001. Potential of bibliographic tools to organize knowledge on the internet: the use of Dewey decimal classification scheme for organizing web-based information resources. *Knowledge Organization* 28 (1), 17–26.
- Saeed, H., Chaudhry, A.S., 2002. Using Dewey decimal classification scheme (DDC) for building taxonomies for knowledge organization. *Journal of Documentation* 58, 575–583.
- Schilling, M.A., 2005. A small-world network model of cognitive insight. *Creativity Research Journal* 17 (2&3), 131–154.
- Schilling, M.A., Phelps, C., 2007. Interfirm collaboration networks and knowledge creation: the impact of large scale network structure on firm innovation. *Management Science* 53, 1113–1126.
- Schilling, M.A., Vidal, P., Ployhart, R.E., Marangoni, A., 2003. Learning by doing something else: variation, relatedness, and the learning curve. *Management Science* 49, 39–56.
- Simon, H.A., 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* 69 (1), 99–118.
- Simon, H.A., 1978. Rationality as process and product of thought. *American Economic Review: Proceedings* 68, 1–16.
- Simon, H.A., Chase, W.G., 1973. Skill in chess. *American Scientist* 61, 394–403.
- Simonton, D.K., 1995. Foresight in insight? A Darwinian answer. In: Sternberg, R.J., Davidson, J.E. (Eds.), *The Nature of Insight*. MIT Press, Cambridge, MA, pp. 465–494.
- Simonton, D.K., 1999a. Creativity as blind variation and selective retention: is the creative process Darwinian? *Psychological Inquiry* 10 (4), 309–328.
- Simonton, D.K., 1999b. *Origins of Genius*. Oxford University Press, New York.
- Steyvers, M., Tenenbaum, J.B., 2005. The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. *Cognitive Science* 29, 41–78.
- Stuart, T., Podolny, J., 1996. Local search and the evolution of technological capabilities. *Strategic Management Journal* 17 (summer special issue), 21–38.
- Thompson, R., Shafer, K., Vizine-Goetz, D., 1997. Evaluating Dewey concepts as a knowledge base for automatic subject assignment. In: *Proceedings of the second ACM international Conference on digital libraries*, pp. 37–46.
- Usher, A., 1954. *A History of Mechanical Inventions*. Harvard University Press, Cambridge, MA.
- Wallas, G., 1926. *Art of Thought*. Harcourt, Brace and Company, New York.
- Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of 'small-world' networks. *Nature* 393, 440–442.
- Weigand, W.A., 1998. The "Amherst Method": The origins of the Dewey Decimal Classification Scheme. *Libraries & Culture* 33 (2), 175–194.
- Zhou, J., Oldham, G.R., 2001. Enhancing creative performance: effects of expected development assessment strategies and creative personality. *Journal of Creative Behavior* 35, 151–167.