

LEARNING FROM SMALL FAILURES: ROLE OF RELATEDNESS, FAMILIARITY,
AND STRUCTURE OF KNOWLEDGE BASE

Rajat Khanna

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in
partial fulfillment of the requirements for the degree of Doctor of Philosophy in Business
Administration in the Kenan-Flagler Business School (Strategy & Entrepreneurship
Department).

Chapel Hill
2015

Approved by:

Atul Nerkar

Rich A. Bettis

Hugh O'Neill

Isin Guler

Scott Rockart

©2015
Rajat Khanna
ALL RIGHTS RESERVED

ABSTRACT

Rajat Khanna: Learning from Small Failures: Role of Relatedness, Familiarity, and Structure of Knowledge Base

(Under the direction of Professor Atul Nerkar)

Does the approach that firms adopt in their *search* for solutions influence their ability to learn from failures? Does the *structure* of a firm's knowledge base affect learning from failures? Recent research has begun to explore the learning from small failures in experimentation. However, failures may differ in terms of learning opportunities they provide, and presence of learning opportunities may not result in better performance. In this dissertation, I focus on small failures in experimentation. I examine two aspects of search behavior. First, I explore relatedness of failures in a firm's knowledge base and examine how failures of varying degree of relatedness can lead to heterogeneous learning outcomes. Second, I investigate whether familiarity of knowledge and knowledge elements associated with failures has an effect on learning from these failures. Finally, presence of learning opportunities does not always result in increased performance, and therefore examination of factors that moderate a firm's ability to implement learning from failures is important. I argue that decomposability of a firm's knowledge base plays a crucial role in facilitating learning from failures. With the help of data on patents and their expirations for 76 pharmaceutical firms, I show that relatedness of failures has positive effect on a firm's R&D performance, but beyond a certain point increase in relatedness hurts the R&D performance. Also, failed experiments that use more familiar knowledge and knowledge elements have

negative effect on a firm's R&D performance. Finally, decomposability of a firm's knowledge base moderates these relationships such that nearly decomposable knowledge base facilitates the learning more than fully decomposable or integrated knowledge bases. By studying the role of different characteristics of small failures in learning and structure of a firm's knowledge base in incorporating that learning, this dissertation increases our understanding of mechanisms underlying R&D processes in pharmaceutical firms.

Keywords: learning, small failures, relatedness, familiarity, decomposability, knowledge base, search, experimentation, pharmaceutical industry, R&D performance.

To my parents and all my teachers

ACKNOWLEDGMENTS

With an undergraduate degree in engineering, it was a dream to pursue PhD in business. After having an initial conversation with Atul Nerkar, I was convinced that UNC is the right place for me. Soon after I joined UNC, Atul assured me that the department's highly qualified faculty and my educational and professional background will be ideal to realize my dream. This dissertation is a testament to Atul's convictions. It is an honor to thank my family, friends, and committee members who have helped me tremendously throughout the doctoral journey.

I take this opportunity to express my deep sense of gratitude to my research supervisor Professor Atul Nerkar. It has been a privilege to work under his guidance. During my first semester of PhD, Atul discussed some of the research ideas he had that led to the beginning of my first research project. He spent hours explaining me all the intricacies of research process. With an engineering mindset and lack of any sort of management training, I was often inept in my approach to solving business problems. But Atul was extremely patient throughout this time and gave me the right perspective. He taught me the art of identifying and developing research ideas and gave me enough freedom to find the research questions that interest me the most. At the same time, he made sure that I remain focused and practical in my scope of the research problem. I feel extremely proud and lucky to be able to work with him and to have him as my advisor. He has been a friend, guide, and role model to me during my time at UNC.

I was extremely fortunate to have Isin Guler on my committee. I have been working with her on different projects since my first year. She taught me how to solve small pieces of a problem separately and put them together afterwards. I have always been impressed by her tremendous intellectual capabilities and work ethics. While working on a project, at each step, she would let me take the initiative and appreciate my efforts whenever necessary. At the same time, she was openly critical of things that I can improve in future. I cannot imagine finishing this dissertation without her mentorship during the last 4 years. Thanks a lot, Isin.

I would also like to thank my other committee members including Hugh O’Neil, Rich Bettis, and Scott Rockart. They have been exceptional in terms of giving me timely feedback throughout the dissertation phase. Their feedback and suggestions have substantially improved the dissertation.

This thesis would have remained a dream had it not been for the support from my family. Aleks has been incredible both in terms of helping me emotionally and contributing intellectually to the ideas developed in this dissertation. Insights that came out of hours of discussion with her helped me develop my research questions more effectively and made this dissertation considerably richer. Being a PhD scholar herself and therefore considerably busy with her own work, it always amazed me how she never refused to dedicate her time to help me. My mother and father have been a tremendous source of strength during all these years and never lost faith in me even in worst of the times. My long vacations back home every year kept me refreshed and going.

This section on acknowledgment will be incomplete without thanking my friends including Swara, Piku, Abhishek, Shubham, Sandy, Yogesh, and Udeyan. They all provided a steady shoulder to rely on in the most desperate times. They were always welcoming to my

unexpected and spontaneous visits and proved to be the best hosts. Spending time with them made me forget about my research and helped me rejuvenate whenever necessary.

TABLE OF CONTENTS

LIST OF FIGURES	x
LIST OF TABLES	xi
1. INTRODUCTION	1
2. LITERATURE REVIEW	9
2.1 Organizational Learning and Failures	9
2.2 Knowledge Creation and Recombination.....	14
2.2.1 Experimentation in Innovation	14
2.2.2 Search	18
3. HYPOTHESES DEVELOPMENT	22
3.1 Relatedness of Failures with Non-failures	22
3.2 Relatedness among Failures	25
3.3 Familiarity of Knowledge Elements	28
3.4 Familiarity of Knowledge	32
3.5 Decomposability of Knowledge Base	35
4. RESEARCH DESIGN AND METHODOLOGY	45
4.1 Drug Discovery and Pharmaceutical Industry Context.....	45
4.2 Data and Methods.....	51
4.2.1 Dependent Variable(s).....	52
4.2.2 Independent Variables	53
4.2.3 Control Variables.....	57
4.3 Empirical Model.....	59
5. RESULTS	61
6. DISCUSSION	67
7. CONTRIBUTIONS AND LIMITATIONS	72
REFERENCES	91

LIST OF FIGURES

Figure 1 - Number of Expired Patents for Firms Included in This Study between 1985 and 2002.....	83
Figure 2 - Network of Subclasses with Red Nodes Representing Elements Underlying Failed Patents and All other Elements Denoted by Blue Nodes.....	84
Figure 3 - Network of Subclasses with Nodes Representing Elements Underlying Failed Patents.....	85
Figure 4 - Knowledge Base of Firms with Nodes of Same Colors Belonging to the Same Module.....	86
Figure 5 - Sample Data Points for Pfizer in Year 2000 and Description of Variables.....	87
Figure 6 - Nature of Relationship between Relatedness of Failures and R&D Outcome Quality at Different Levels of Decomposability.....	88

LIST OF TABLES

Table 1 - List of 76 Pharmaceutical Firms Used in This Dissertation.....	75
Table 2 - Descriptive Statistics and Correlation Matrix.....	76
Table 3 - Fixed Effects Negative Binomial Estimates for R&D Performance (DV = R&D Productivity $_{it+1}$)	77
Table 4 - Fixed Effects Negative Binomial Estimates for R&D Performance (DV = R&D Quality $_{it+1}$).....	80

1. INTRODUCTION

Failure is instructive. The person who really thinks learns quite as much from his failures as from his successes.

- John Dewey

Most firms in high technology industries have experienced failures in their research and development (R&D) activities. The process of innovation in R&D is highly uncertain, and failure is an integral part of it. One of the ways firms in these industries innovate is through the process of experimentation. There is a large body of literature that has emphasized the importance of experimentation in R&D and its role in creating innovation (e.g., Ahuja & Lampert, 2001; Lee, Edmondson, Thomke, & Worline, 2004; Nohria & Gulati, 1996; Thomke & Kuemmerle, 2002). Novel inventions are the direct result of exploration, which also occurs through experimentation and search.

Failures that occur during the course of the experimentation in R&D are considerably different from other failures explored in prior literature. Most of the research focusing on failures has examined the consequences of errors or mistakes. This stream of research includes the learning from operational failures such as acquisition problems and product defects (Hayward, 2002, Haunschild & Rhee, 2004, Henderson & Stern, 2004), or catastrophic failures such as railroad accidents and plane or orbital crashes (Haunschild & Sullivan, 2002, Madsen and Desai, 2010). Experimental nature of failures in R&D separates them from operational failures. In case of operational failures, a firm's primary purpose is to

minimize these failures, whereas failures during R&D are the only way a firm can increase its understating of ongoing R&D activities.

While the purpose of the firm is not to fail in R&D, these failures do become the only feedback mechanism through which firms receive information to improve existing R&D projects and design better future experiments (Sitkin, 1992; Fleming & Sorenson, 2004). I specifically focus on small and frequent failures in experimentation during the R&D process. These failures do not threaten a firm's existence and while not desired are expected to occur. Therefore, individuals in firms are more open to acknowledging these failures, facilitating learning as a result.

In this dissertation, I focus on firms' failed R&D efforts. Specifically, I use premature discontinuation of a patent as a sign of small failure in a firm's R&D. In 1980, United States Patent and Trademark Office (USPTO) passed a law that required firms or entities to pay maintenance fees every 4, 8, and 12 years to keep their patents active, and failure to pay the maintenance fees for a patent will lead to its discontinuation. USPTO provides data for all the patents that have been voluntarily given up by firms and individuals starting 1980. Recent research has also used this data to study small failures in experimentation and shows that these failures do provide learning opportunities for firms in the pharmaceutical industry (Khanna, Guler, & Nerkar, forthcoming, 2015).

Given that firms do learn from small failures in experimentation, a natural question to ask is: are these failures same with respect to the learning opportunities they provide? R&D in pharmaceutical firms is aimed at the creation of knowledge through recombination. The process involves search and discovery of novel combinations by using the existing pieces of knowledge (Schumpeter, 1934; Henderson & Clark, 1990; Kogut & Zander, 1992). Some of

these combinations are successful while others are unable to meet firms' expectations or are failures. Learning from failure, in the present context, occurs through generation of feedback, and it is reasonable to assume that not all failures provide identical feedback to firms.

Experimentation in R&D entails combining existing knowledge pieces in various ways, and how firms choose to combine these knowledge pieces will eventually determine their R&D outcomes. Firms' failure experience with respect to the ways they carry out the process of recombination, or search, influences their ability to learn from these failures. Such small failures represent the ways firms conduct search and failure experience with different search behaviors can have nontrivial implications for firms' learning outcomes.

In this dissertation, I explore two dimensions of search behavior associated with small failures in experimentation and their role in learning. Specifically, I ask: do relatedness and familiarity of small failures in experimentation affect the extent to which firms learn from them? As failures are natural outcome of experimentation, it is important to understand how they differ in terms of learning opportunities they provide.

Specifically, I examine how failures that characterize different search behaviors influence one of the most important outcomes of innovation activities – R&D performance, differently. A successful pharmaceutical firm is capable of producing high level of R&D performance and competent at commercializing products of R&D (e.g., Fleming, 2002; Nerkar & Roberts, 2004; Schumpeter, 1934). I focus on the former and examine how failure experience associated with different search behaviors influences a firm's ability to learn from these failures. I also distinguish between two important dimensions of R&D performance: R&D productivity (patent count) and R&D outcome quality (forward citations to patents). Previous research either uses one of the measures of R&D performance or combines the two

into one. R&D productivity and quality are distinct from each other and it is logical to study the influence of failures on these two dimensions separately to obtain a more holistic view of learning from failures (Khanna, Guler, & Nerkar, forthcoming).

R&D in general is a problem solving activity and firms search for solutions to the problem at hand which in case of the pharmaceutical industry is potential treatment for a disease. One of the ways firms search for solutions is through the process of recombination, i.e. by trying various combinations of existing knowledge pieces. In line with previous research, I refer to these knowledge pieces as knowledge elements (Fleming, 2001; Yayavaram & Ahuja, 2008).

Firms can differ in terms of which knowledge elements they decide to experiment with and that can lead to different learning outcomes. For example, a firm's experience with failures can be a result of experimentation with related knowledge elements, i.e. elements that were previously combined by the firm (Breschi, Lissoni, & Malerba, 2003) or vice versa. Relatedness of failures denotes the extent to which knowledge elements underlying these failures are combined with other elements in a firm's knowledge base. Such failures share common knowledge elements with other projects in a firm and are built on common heuristics and scientific principles. Therefore, failures with higher degree of relatedness will provide information that can be used to evaluate the potential of many ongoing projects. The firm will also be able to make informed decision about how to allocate its attention and resources across different projects. However, beyond a certain point, an increase in relatedness of failures will make the causal inferences difficult and can hurt a firm's R&D performance. I show in this dissertation that firms benefit as the relatedness of failures

increases but beyond a certain point increase in relatedness begins to hurt firms' R&D performance.

Similarly, failures are also related to each other. Following the same logic as discussed earlier, relatedness among failures indicates the extent to which knowledge elements underlying these failures are combined with each other. Failures of this kind represent a firm's intention to uncover relationships among a set of elements. Failures that are highly related to each other can be a great source of information about relationships among elements and can help a firm eliminate certain combinations, reducing the search space substantially for future experiments. Though, an increase in relatedness among failures beyond a certain point indicates that the firm has tried a large number of combinations of elements and failed in all of them. Such failures suggest that the firm has experimented with elements that lack potential and it should experiment with a new set of elements, which in the long run may benefit the firm but will hurt its subsequent R&D performance. Also, beyond a point, increase in relatedness among failures implies presence of large number of common combinations across all failures, making the process of identifying causes of these failures more challenging. Empirical analysis in this study shows that firms do benefit from failures of high relatedness initially, but after a certain level, increase in relatedness among failures begins to hurt firms' R&D performance negatively.

Another dimension of search behavior studied in this dissertation is familiarity of knowledge used in failures. During experimentation, a firm can rely on familiar knowledge or use new knowledge. Firms that rely on familiar knowledge (also known as local search) use their existing stock of knowledge to search for solutions (e.g. Helfat, 1994; Rosenkopf & Nerkar, 2001; Stuart & Podolny, 1996). At the other extreme, when firms rely on less

familiar knowledge (or conduct distant search), they consciously decide not to be dependent on existing routines and stock of knowledge, but use new knowledge to find solutions (March, 1991; Miner, Bassoff, & Moorman, 2001; Rosenkopf & Nerkar, 2001). These two search behaviors are distinct and indicate a firm's intentions to conduct R&D using two considerably different approaches. A firm is likely to face fewer failures when experiments utilize more familiar knowledge (Laursen, 2012). Therefore, large number of failures as a result of experimentation with more familiar knowledge indicates that a firm's current knowledge is outdated, and it needs to rely more on less familiar knowledge. As a result, a firm's subsequent R&D performance will suffer when it faces large number of failures following experiments with familiar knowledge elements, which is supported by the empirical analysis in this study.

Similarly, firms can differ with respect to the familiarity of knowledge elements underlying failures. Repeated usage of familiar knowledge elements reduces the likelihood of errors and failures, which facilitates the process of search, making it more predictable and reliable (Eisenhardt & Tabrizi, 1995; Levinthal & March, 1981). Also, higher extent of experience with familiar knowledge elements helps firms develop deeper understanding of elements and their combinations. A firm can either experiment with more familiar knowledge elements, or it can try new elements and their combinations (Helfat, 1994; Fleming, 2001; Stuart & Podolny, 1996). Failures in experiments targeted at combining more familiar elements suggest that the firm has either exhausted or is not able to identify the potential combinations, which can hurt its R&D performance negatively. Empirical analysis supports this notion.

Finally, learning is a two-step process. First step is the generation of feedback and second step is incorporation of the feedback. Failures imply the generation of feedback. However, the presence of feedback resulting from failures does not guarantee learning. To be able to successfully learn from failures, a firm must be able to incorporate what it has learned from failures. In the present context, incorporation of learning requires making necessary changes in ongoing R&D activities. A firm's R&D activities can be represented by all knowledge elements and their combinations within the firm, and are often referred to as its knowledge base.

The last section of the dissertation examines the role of structure of a firm's knowledge base in its ability to incorporate learning from small failures in experimentation. A firm's knowledge base can display multiple structural properties, but in this dissertation, I focus on its decomposability because previous research has shown that decomposability plays an important role in a firm's ability to make changes in its knowledge base (Yayavaram & Ahuja, 2008). Decomposability implies the extent to which a firm's knowledge base can be decomposed into modules, and structures that display moderate level of decomposability or are nearly decomposable, are best positioned to adjust in response to a change (Simon, 1962, 1991).

Nearly decomposable structures lead to inventions of higher quality, primarily because these structures are flexible in nature, i.e. they are able to change more effectively (Levitan, Lobo, Kauffman, & Schuler, 1999; Levinthal & March, 1993; Yayavaram & Ahuja, 2008). Therefore, following failures, firms will find it easier to introduce necessary changes in knowledge bases characterized by near decomposability. In my analysis, I find that firms

with nearly decomposable knowledge bases are able to learn more from failures than firms with knowledge bases characterized by low or high level of decomposability.

The remainder of the dissertation is structured as follows. First, I provide detailed overview of literature on learning from failures and knowledge creation through recombination. This is followed by a section on hypotheses development. Next, I provide empirical analysis and results for the hypotheses. Discussion on overall findings, contribution of this dissertation to the literature, limitations, and directions for future research follow.

2. LITERATURE REVIEW

2.1 Organizational Learning and Failures

A large body of literature in organizational learning has established that organizations learn from experience. Recent studies have also started to examine the implications of both success and failure experiences for organizational learning (Haunschild & Sullivan, 2002; Lapre & Tsikriktsis, 2006; Baum & Dahlin, 2007). Cyert and March (1963) in their influential study clearly distinguish the ways in which firms respond to success and failure. They argue that firms have certain expectation with respect to the performance, and decision makers compare this expectation, also known as aspiration level, with actual performance. If the actual performance is above the aspiration level, decision makers infer the experience to be a success, believing that the stock of knowledge and current practices are adequate and require no changes. Also, success leads decision makers to ignore information in the outside world and use simple rules to make future decisions (Audia et al., 2000). On the other hand, if actual performance is below aspiration level, decision makers consider the experience a failure, finding it necessary to look for solutions or alternatives that can amend the problem of decreased performance.

Failures shift the focus of firms to the areas that require more attention. They inform firms that current practices may not be suited for external environment (Cyert & March, 1963; Nelson & Winter, 1982; Greve, 1998), that there may be flaws in internal strategy and structure (Chuang & Baum, 2003; Barnett & McKendrick, 2004), or that there is a need to

perform distant search to find better solutions (Levinthal & March, 1993; Audia et al., 2000; Denrell, 2003). Failures not only provide a firm with motivation to begin distant search to identify superior alternatives, but also provide information on the direction of non-local search (Wildavsky, 1988).

Research in management has explored the implications of various types of failures for organizational learning. Some of these studies analyze failures that are catastrophic in nature in terms of their consequences for firms, e.g. accidents in railroad and orbital launch vehicle industries (Haunschild & Sullivan, 2002; Madsen & Desai, 2010); while other studies focus on operational failures such as automobile recalls, unsuccessful acquisitions, and technological obsolescence (Haunschild & Rhee, 2004; Hayward, 2002; Henderson & Stern, 2004).

Even though most firms experience failure at some point, learning from it is not straightforward. Strangely, even though the idea that firms should learn from failures seems obvious, we rarely observe it taking place (Cannon & Edmondson, 2005). Firms frequently struggle to learn from failures, because individuals in firms are averse to associate with failure and tend to distance themselves from it (Cannon & Edmondson, 2005). The alienation of organizational members from a failure makes it difficult to uncover the underlying causes of the failure, impeding the process of learning from it as a result. Firms often have policies and procedures in place that reward success and penalize failure. Such routines provide no incentives for organizational members to identify with a failure and learn from it (Argyris, 1990).

In addition to firm-specific policies, psychological factors such as desire of an individual to maintain high self-esteem and be in control of the situation also obstruct the

process of learning by attributing success to his/her own actions and failure to actions of others or environmental factors (Weiner, 1971, 1985). In the presence of success-driven incentive structure in firms and psychological barriers associated with admitting a failure, individuals either ignore or hold somebody else responsible for the failure, thus undermining learning from it (Cannon & Edmondson, 2005). Recent research also highlights that firms may learn more from failures of other firms than their own failures (Baum & Dahlin, 2007). Therefore, availability of learning opportunities following failures is not enough for learning to take place, and factors such as a firm's openness in acknowledging failures, and ability of individuals to overcome the aversion to associate with failures play a crucial role in deriving valuable insights. Bill Gates, the former chief executive of Microsoft Corp., expresses this perspective succinctly (Perman, 2006):

In the corporate world, when someone makes a mistake, everyone runs for cover. At Microsoft, I try to put an end to that kind of thinking. It's fine to celebrate success, but it's more important to heed the lessons of failure. How a company deals with mistakes suggests how well it will bring out the best ideas and talents of its people, and how effectively it will respond to change

Although current studies have greatly increased our understanding of learning from failure, research on small failures that occur in R&D is limited. R&D process is incremental and path dependent, requires substantial time and investment, and extremely uncertain in nature (Nelson & Winter, 1982). Therefore, failure is a natural outcome in the course of innovation in high technology industries, e.g. pharmaceutical, semiconductor, optical disk, etc. One of the ways firms in these industries innovate is through the process of experimentation. As mentioned earlier, R&D process is inherently uncertain, and

experimentation helps in evaluating potential alternatives and finding the optimal solution (Ahuja & Lampert, 2001; Cannon & Edmondson, 2005; Thomke & Kuemmerle, 2002).

Most of these experiments fail, and even though firms do not particularly welcome such outcomes, they are aware of them. However, by no means I suggest that failures are desirable. In fact, quite the opposite: failures can deplete scarce resources, discourage individuals, disrupt ongoing activities within a firm, and even be tragic (McGrath, 2011). Inevitability of failures during experimentation though helps firms overcome many obstacles to learning (Thomke, 2003; Thomke & Kuemmerle, 2002).

Most studies on learning from failure examined failures that are catastrophic, and therefore require significant organizational attention, or operational in nature. However, relatively fewer studies have explored the implications of small failures for organizational learning. Small failures often are an excellent source of learning and prevent future large failures (Cannon & Edmondson, 2005; Sitkin, 1992). They are also an important precursor to learning as they can effectively direct a firm's focus to areas that require immediate attention. As these failures are not threatening to the survival of the firm, it is possible to both attend to such failures and engage in search for solutions. Small failures, in this respect, bring a unique opportunity for learning.

However, small failures are often ignored by organizational members. Small failures may not receive adequate attention from organizations, and feedback resulting from these failures can also be forgotten (being less in amount) in short periods (Sitkin, 1992; Levinthal & March, 1993; Baumard & Starbuck, 2005). Cannon and Edmondson (2005) argue that one of the primary reasons behind organizations not learning is the neglect of small failures. As small failures, at the time of their occurrence, seem insignificant and unimportant, firms tend

to ignore these failures and miss out on an opportunity to learn. As a result, in the absence of learning from small failures, it is difficult to prevent large failures later in organizational life (Tucker & Edmondson, 2003).

As discussed, small failures provide feedback that can help a firm in the long run, but at the same time they are at the risk of being ignored by individuals within the firm. Given opposing views on whether firms learn from small failures, scholars have studied factors that can explain this discrepancy (Cannon & Edmondson, 2005; Haunschild & Rhee, 2004; Henderson & Stern, 2004; Thomke, 2003).

One such factor that can facilitate learning and is also applicable to small failures examined in this dissertation is experimentation (Campbell, 1991; Cannon & Edmondson, 2005; Hedberg, Nystrom, & Starbuck, 1976). Experiments are expected to have uncertain outcomes and intended for learning (Thomke, 2003). Therefore, failures during experimentation are unavoidable and not perceived in a negative light within a firm (McGrath, 2011; Sitkin, 1992). Inevitability of failures in experimentation helps individuals in firms to overcome the tendency to detach themselves from failures and promotes learning as a result (Goleman, 1985). Firms that engage in experimentation experience more failures, but they tend to be more innovative and successful compared to those firms that do not encourage experimentation (Thomke, 2003).

To sum up, previous studies on learning have either focused on catastrophic failures or operation failures. These failures are not a result of experimentation, and firms would rather avoid them. In contrast, small failures studied here are a result of experimentation in a firm's R&D, and objective of the firm is not to minimize these failures but to use the information from these failures to improve its understanding of R&D activities. Recent

research supports the notion that firms in pharmaceutical industry learn from small failures in experimentation (Khanna, Guler, & Nerkar, forthcoming 2015). However, not all small failures are same in terms of the information they provide. Systematic examination of how small failures differ with respect to their relationships with other projects, including both failures and non-failures, and implications of such heterogeneity for firm performance is essential to develop better understanding of the innovation process.

2.2 Knowledge Creation and Recombination

Innovation in pharmaceutical industry involves finding a combination of knowledge elements that can be used to treat a disease of interest. In that respect, pharmaceutical firms create new knowledge through the process of recombination, i.e. by combining existing knowledge elements in different ways (Fleming, 2001). Firms are not aware of outcomes of these combinations beforehand and experimentation is the only way to find out which combinations work and which do not. Over time firms learn which elements and their combinations are more or less useful in different contexts. This knowledge can help firms eliminate elements and their combinations from future experiments, thus reducing the search space substantially in which the firms have to find solutions (Vincenti, 1990).

The underlying processes behind knowledge creation and recombination are experimentation and search. Faced with a problem, firms are searching for its solution, and because of the high uncertainty and lack of information about the outcome of different combinations, firms engage in experimentation.

2.2.1 Experimentation in Innovation

Several studies have examined the influence of breakthrough innovations, and innovations in general, on firms and industries (Belderbos, Carree, & Lokshin, 2004;

Brouwer & Kleinknecht, 1999; Capon, Farley, Lehman, & Hulbert, 1992). However, processes such as problem solving and experimentation behind innovation activities have not attracted the attention of scholars as much. The products of R&D, in most part, are a manifestation of problem solving embedded in innovation process, and yet literature on technological change has largely ignored it. With an increased understating of complexity of underlying technological landscape, experimentation and problem solving become indispensable and represent a major proportion of economic activity (Carter, 1995).

When outcomes of various scenarios, which are part of a problem, cannot be predicted or logically derived, experimentation is the only way to solve the problem. Experimentation involves trial-and-error learning, and insights generated from the learning are used to obtain an idea about the direction in which the solution to the problem might lie (Barron, 1988).

The idea behind learning from experimentation, in the present context, is as follows: a firm is presented with a number of alternatives, but the outcome of these alternatives is neither known nor can be estimated computationally. ‘Learning by doing’ remains the logical approach to determine the superiority of these alternatives. The alternatives may not even contain the best possible solution, as the alternatives do not originate based on the true relationships among elements but the firm’s understanding of these relationships. These alternatives are then tested within the constraints of experimental design (Duncker, 1945; Simon, 1969). Every experiment or a set of experiments provides information on the outcomes that may or may not be expected. Unexpected results are used to correct the ‘error’ on the part of the firm, and future experiments are refined and modified based on this information. This process is iterated until the firm reaches satisfactory solution, all possible

alternatives are exhausted, or the firm decides to discontinue the process for some other reason. Owing to the uncertainty in experimentation, a firm may even find solution to the problem it never intended to solve. One of the examples is the discovery of Viagra. Scientists, working on a drug for a cardiovascular disease, ended up discovering the use of the drug for a different therapeutic indication (Foye, 2008).

Experiments are often conducted on simple models, i.e. simplified version of the real world situation. Such simplification helps firms test multitude of scenarios in a short amount of time, which is not possible in the real world. Since these simple models do not incorporate all attributes of reality, the application of successful experiments can still lead to failure in real world settings (Thomke et al., 1998; Tyre & von Hippel, 1997). Failures during experimentation therefore are part of the process and expected to occur, and resulting information can be used to find better solutions.

The fundamental factor that differentiates failures studied in this dissertation from failures studied in previous research is their experimental nature. As failures are inevitable in experimentation, individuals are less averse to acknowledging these failures and are also able to overcome psychological biases associated with them (Thomke, Hippel, & Franke, 1998). As a result, instead of distancing themselves from failures, individuals in firms analyze these failures more objectively to find potential causes. In the absence of severe consequences following failures, individuals are motivated to identify reasons of failures and update their beliefs as the new information from failures presents itself, as opposed to defending their flawed information to avoid punishment. Nystrom and Starbuck (1984) also discuss the importance of experiments in learning:

People who see themselves as experimenting are willing to deviate temporarily from practices they consider optimal in order to test the validity of their assumptions.

Small failures studied here occur during the drug discovery process in the pharmaceutical industry. Pharmaceutical firms combine different pieces of knowledge during the discovery phase to develop a potential drug for the disease of interest. These pieces of knowledge are often referred to as knowledge elements and essentially are chemical compounds owned by firms. Thus, small failures in experimentation are the result of a firm's unsuccessful attempts to use a single knowledge element or combine more than one knowledge elements during the drug discovery process, with latter occurring more often. In most cases, firms have little knowledge about outcomes of these experiments. As a result, many experiments do not lead to desired outcomes, or in other words fail. Small failure in experimentation in this regard refers to a R&D project that involved one or more knowledge elements and failed to meet the firm's expectations, leading to its discontinuation.

In the absence of complete knowledge of underlying science, the most effective way to move forward is to experiment with different combinations of knowledge elements. Previous studies have demonstrated the role of problem-solving in product and process development (Allen, 1966; Clark & Fujimoto, 1991; Marples, 1961; Smith & Eppinger, 1997; Thomke, 1997; Clark & Wheelwright, 1992). All these studies provide evidence of trial-and-error learning via the process of experimentation. However, the role of failed experimentation in the recombinant R&D efforts and its impact on performance outcomes has not been examined in detail. This dissertation not only examines the link between failed experiments and innovation outcomes, but also explores how this relationship varies depending on the different characteristics of experiments.

2.2.2 Search

Invention or innovation is the process of combining (or recombining) individual components (Schumpeter, 1939, p. 88), or in this case, knowledge elements. These elements can themselves comprise other elements and so on. Invention often involves finding desirable outcome either by recombining old elements or combining new elements. Henderson and Clark (1990) provide clear distinction between the different ways firms engage in the process of innovation. They emphasize the importance of architectural innovation in addition to the already prevalent notions of incremental (Nelson & Winter, 1982; Ettlie, Bridges, & O'Keefe, 1984; Tushman & Anderson, 1986) and radical innovation (Dess & Beard, 1984; Dewar & Dutton, 1986). Architectural innovation involves reconfiguration of existing knowledge elements. It is unique in terms of its effect on firms' decision making in that it does not require firms to acquire knowledge about new elements but understand more deeply how these elements are related to each other (Henderson & Clark, 1990).

How firms choose to innovate can be viewed as a search activity over a technology landscape. Consider an area that has 'hills', 'peaks', and 'valleys' with smooth climbing on a hill representing incremental innovation and jumping directly to a peak as radical innovation. There is no 'one-size-fit-all' strategy for all firms with respect to how firms should search over the landscape, but rather the strategy depends on a firm's set of existing capabilities and understanding of the landscape. For example, Dell, which has competitive advantage in efficient manufacturing and supply chain management relies more on incremental innovation, as introducing radical innovation will lead to substantial changes in other operations, leading to less efficient manufacturing and distribution. Apple, on the other hand, is known for breakthrough innovations and engaging only in incremental search can hurt its competitive advantage (Fleming, 2012).

By viewing knowledge elements and their combinations part of a technology landscape, firms can achieve better understanding of their existing position and potential options that can be leveraged going forward. Following Kauffman's NK model (Kauffman, 1993), a technology landscape can be portrayed as a collection of N knowledge elements and interdependency among these elements is denoted by K (Fleming & Sorenson, 2001, 2004; Sorenson, Rivkin, & Fleming, 2006). There has been a tradition of using NK models to explore different aspects of a firm's behavior including organizational adaptation, technological evolution, decision making, and competitive behavior (e.g. Ethiraj & Levinthal, 2004; Frenken, 2000; Gavetti & Levinthal, 2000; Levinthal, 1997; Rivkin & Siggelkow, 2003). I use empirical data and rely on descriptive characteristics of the NK model to explain the heterogeneity in learning outcomes as a result of failure experience in different search behaviors.

Given presence of true relationships among knowledge elements in a landscape, firms attempt to uncover these relationships through experimentation, which is essentially a manifestation of the search process. At first, inventors have no way of knowing the topology of landscape and rely on small steps, i.e. local search to gain some understanding of existing relationships among elements (March & Simon 1958; Cyert & March 1963). This strategy may work well in landscapes that have one smooth peak, as eventually all firms will reach the peak regardless of where they start in the landscape. In cases where landscape is rugged, local search may not be the best strategy. Firms may trap themselves on a local peak that is much lower than the other local peaks and the global peak. Any local search made at this inferior local peak will lead to decrease in fitness value and firms will tend to stay on this peak. Under these circumstances, firms may be better off exploring distant regions in the

landscape. Although, such undertaking is risky in nature, it has the potential to provide huge pay-offs to firms (Rosenkopf & Nerkar, 2001).

Search in the technology landscape, in this sense, can be interpreted as a problem-solving activity (Nelson & Winter, 1982; Katila & Ahuja, 2002). Firms are faced with a certain problem and to solve the problem they are looking for an optimal solution in the landscape. In the absence of complete understanding of underlying science, searching and trying different configurations of knowledge elements is the only way boundedly rational decision makers can learn about the causal relationship among elements and discover the optimal solution (Fleming & Sorenson, 2004; March & Simon, 1958; Sitkin, 1992).

Two primary reasons for firms to engage in search are: (1) decision makers are cognitively limited (Cyert & March, 1993; Carter, 1971; Anderson, 1983; Pinfield, 1986). Even when decision makers have all the information regarding the individual elements in the landscape, it is impossible to predict the outcomes of various configurations of these elements, and (2) even if we assume that decision makers are equipped with extraordinary skills to solve complex problems, they still cannot find optimal solution because of the nature of the problem, as most of these problems do not have a closed-form solution (Simon, 1996).

As discussed, even for small number of elements in the landscape, the potential number of possible combinations can be exponential in nature, and because of the bounded rationality of decision makers, there is no way of finding the location of ‘peaks’ and ‘valleys’ in the landscape ex ante. One of the ways firms increase their understanding of the landscape is through the process of search. Search is evolutionary in nature, i.e. it is path dependent and often involves recombination of existing knowledge elements (Fleming, 2001; Gilfillan, 1935; Mokyr, 2002; Nelson & Winter, 1982).

Prior research has shown that many novel inventions result from the synthesis and recombination of a firm's existing technologies (Basalla, 1988; Gilfillan, 1935; Hargadon & Sutton, 1997; Henderson & Clark, 1990; Schumpeter, 1939; Usher, 1954; Weitzman, 1996). Firms often search in the neighborhood, and most of the progress is due to the small changes, unless there is a technological discontinuity (Dosi, 1988; Sahal, 1981; Tushman & Anderson, 1986). On these lines, invention is the result of process of combining new or existing knowledge elements, and primary goal of the search is to identify superior elements and their combinations (Fleming, 2001; Henderson & Clark, 1990).

Literature on search has broadly discussed the importance of different search behaviors in innovation (Gavetti & Levinthal 2000; Rivkin, 2000). Studies in this stream have explored the role of a firm's overall experience (including both success and failure experience) with different search behaviors in its innovative performance. For example, Fleming (2001) and Katila and Ahuja (2002) explore how a firm's experience with new or familiar elements can lead to differential innovation outcomes. Similarly, a firm's experience with local or distant search can also have an effect on its subsequent innovative performance (Rosenkopf & Nerkar, 2001).

However, literature has not yet examined how search behaviors characterized mainly by failure experiences influence innovation outcomes. As firms respond differently to success and failure (Cyert & March, 1963), they are also likely to learn differently from the two experiences. By examining the search behavior corresponding to failure experience and its effect on innovation outcomes, this study helps identify mechanisms specific to failures and therefore uncover some of the intricacies of learning in innovation process.

3. HYPOTHESES DEVELOPMENT

The aim of science is not things themselves, ..., but the relations between things: outside those relations there is no reality knowable.

- Henri Poincare - Science and Hypothesis (1905; p. xxiv)

3.1 Relatedness of Failures with Non-failures

Firms learn more from failures that arise from heterogeneous causes interacting in complex ways, than from failures resulting from homogenous causes (Haunschild & Sullivan, 2002). Faced with a failure, firms often look for the cause on the surface and ignore the latent conditions responsible for the failure. When failures seem to have originated from heterogeneous factors, and these factors interact with each other in complex ways, firms are forced to think harder and dig deeper to discover the causes of failures (Reason, 1997). Also, instead of blaming each other, individuals faced with complex problems tend to work together and try to find the solution collectively (Jehn, Northcraft, & Neale, 1999). Such in-depth analysis is helpful in finding elements and combinations of elements responsible for failures and therefore facilitates more learning.

In the context of this study, failures vary in terms of which knowledge elements they combine. Some of these failures have higher degree of relatedness, i.e. contain elements that are part of other projects as well, whereas others may be isolated in nature, i.e. they comprise relatively independent knowledge elements. Expanding on the conceptual definition in earlier research, two elements are said to be related to each other when they appear often in the same invention (Teece, Rumelt, Dosi, & Winter, 1994).

In this respect, higher relatedness implies extensive combination of elements underlying failures with other elements, which will result in firms perceiving the source of these failures heterogeneous in nature, and therefore will invoke greater commitment in terms of time and resources. In addition, failures with high degree of relatedness with non-failures indicate that knowledge elements underlying these failures are often combined with other elements in projects that did not fail yet. Presence of these elements in non-failures will provoke deeper analysis and provide greater motivation for a firm to find the elements and their combinations responsible for failures. Such analysis will reveal vital information, which can be used to improve projects that did not fail yet. Therefore, firms will learn more from failures with higher degree of relatedness because of both in-depth analysis and availability of more information.

However, failures consisting of knowledge elements highly related to other elements in a firm's knowledge base will not always result in greater learning for three reasons. First, presence of knowledge elements underlying failures in many other projects will make a firm question the potential of all of these projects, and even though some of these projects may prove valuable in future, the firm will find it difficult to identify them. Therefore, too many common knowledge elements across failures and non-failures will make the process of separating superior elements and their combinations from the inferior elements difficult, resulting in causal ambiguity and inhibiting learning as a result (Reed & Defillippi, 1990). Second, higher relatedness of failures with other projects suggests that the firm experimented with same knowledge elements it has combined previously, which will hinder the firm's ability to learn about new knowledge elements and their combinations. Third, high relatedness of failures with non-failures is also an indication of fundamental flaws in a firm's

R&D activities, because of the presence of same knowledge elements in both failed projects and projects that did not fail yet.

Thus, higher relatedness of failures with non-failures will enable a firm to identify non-failures that contain combinations in which the firm experienced failure. By comparing the combinations of elements across failures and non-failures, a firm can prioritize its ongoing projects and commit its time and investment to more productive areas. For example, if a firm experiences failures in a combination of elements x, y, and z, it can deprioritize non-failures that also use the same combination and focus more on other projects. Such comparison allows firms to identify projects that have less likelihood of failures and allocate more attention to these projects. However, beyond a point, an increase in relatedness of failures with non-failures indicates the presence of large number of common combinations across failures and non-failures and will make the process of prioritizing ongoing projects challenging. In this situation, a firm has no option of choosing some projects over others as most of the non-failures contain combinations in which the firm experienced failures.

Therefore, I hypothesize:

Hypothesis 1: The relatedness of failures with non-failures in a firm's R&D has a curvilinear association with a firm's subsequent R&D performance.

3.2 Relatedness among Failures

Failures are part of a firm's overall knowledge base, but they also represent a separate knowledge base in which elements underlying failures are connected to each other.

Relatedness of failures in a knowledge base made only of failures can also have important implication for learning.

Higher *relatedness among failures* suggests that a firm attempted to combine elements, associated with failures, often with each other. By trying and failing in various combinations of same set of elements, firms attain deeper knowledge of elements and their relationships. Firms can leverage this knowledge by combining the elements in a novel way, which can be a source of architectural innovation. To quote from Henderson and Clark (1990):

The essence of an architectural innovation is the reconfiguration of an established system to link together existing components in a new way.

On the other hand, lack of relatedness among elements implies either the firm did not experiment with different combinations, or it recombined the elements rather successfully. It is also possible that the firm experimented with wider set of elements, i.e. instead of recombining a smaller set of elements extensively; it tried fewer combinations of a larger set of elements. Therefore, lower relatedness of failures indicates that either the firm was successful in experiments it conducted, or chose to experiment with large number of elements but only superficially. In any case, there is not much to learn from such failures as in case of failures of higher relatedness.

Too much relatedness among failures, however, can obstruct the process of learning. As the relatedness among failures increases, a firm can use this information and avoid using same combinations in future experiments. However, beyond a certain point, increase in relatedness indicates that the firm has combined the elements extensively and failed in all of them. Failures of this kind do not reduce search space for future experiments, but suggest that firm should look for solutions elsewhere. Therefore, following failures of this kind, a firm is required to experiment with different set of elements that can take a long time before the firm

will begin to identify productive combinations and will hurt its subsequent R&D performance. Higher extent of relatedness among elements underlying failures also suggests that the firm combined most of the elements across all failures. Presence of large number of combinations across most of the failures will make it difficult for firms to identify combination(s) responsible for these failures.

Let us consider a case in which a firm experiments with 10 knowledge elements e_1, e_2, \dots, e_{10} in 10 different projects, and all these projects did not result in desired outcomes, or in other words, failed. Now, let us say that all these 10 projects combine 9 of the 10 knowledge elements. These projects can be visualized as:

$p_1 = \{e_1, e_2, e_3, e_4, e_5, e_6, e_7, e_8, e_9\}$
 $p_2 = \{e_1, e_2, e_3, e_4, e_5, e_6, e_7, e_8, e_{10}\}$
 $p_3 = \{e_1, e_2, e_3, e_4, e_5, e_6, e_7, e_9, e_{10}\}$
 \cdot
 \cdot
and so on.

All these projects contain most of the elements, and therefore it would be impossible for a firm to understand which knowledge elements and their combinations are responsible for failures. For example, let us consider that the presence of any one of the combinations of e_1, e_2 , and e_3 is in fact the cause of all failures. However, as all projects comprise many other common elements and their combinations, a firm has no way to find that the elements e_1, e_2 , and e_3 are the sources of failures. Alternatively, such failures can also result from a rather big failure in a firm's R&D. For example, a pharmaceutical firm may be engaged in a research program that focuses on a set of related diseases. Therefore, potential treatment for these diseases will involve higher degree of overlap between knowledge elements, which will lead to high value of relatedness among failures.

As these failures are not equivalent to several small failures, but instead indicate a big failure, i.e. failure of the research program, the firm will experience severe negative consequences, overshadowing any potential benefits of learning from these failures. Following the discussion, I argue that *relatedness among failures* will provide a firm with vital information about the relationships among knowledge elements, but at the same time too much relatedness will either result in the firm's inability to identify elements and their combinations responsible for failures, or is an indication of a relatively large failure that will hurt the firm's subsequent R&D performance.

Initially, higher value of relatedness among failures will provide a firm with information about the combinations that do not work, which can help the firm eliminate these combinations from future experiments, reducing the search space considerably. However, increase in relatedness among failures beyond a certain point signals failures in most of the combinations of elements involved, and a firm will need to experiment with different set of elements. Working with new set of elements will take time before a firm can develop some understanding of how these elements and their combinations work, and even though in the long run it can be beneficial, the firm will experience a drop in its subsequent R&D performance. These arguments lead to the second hypothesis:

Hypothesis 2: The relatedness among failures in a firm's R&D has a curvilinear association with a firm's subsequent R&D performance.

3.3 Familiarity of Knowledge Elements

Earlier sections discussed relatedness of failures with non-failures and with each other. I also discussed the implications of relatedness for learning and subsequent R&D performance. The focus of this section is on familiarity of knowledge elements that are part

of the failures, i.e. the extent to which firms have experience with these elements.

Conditional on the search behavior, a firm can combine familiar elements, combine new elements with familiar ones, or combine new elements with each other. Failures representing experimentation with familiar elements will strengthen a firm's understating of linkage between existing elements, whereas failures characterized by combination of less familiar elements will present the firm with new information about potential combinations.

Previous research has explored the role of firm's overall experience in producing novel inventions. Experience with less familiar knowledge elements broadens the scope of search and increases firms' overall understanding of search landscape. Experiment with unfamiliar technology is similar to the process of exploration and does not guarantee success. Firms, unfamiliar with new elements, are aware of negative outcomes associated with such experiments. At the same time, experiments with less familiar knowledge elements suggest firms' awareness and need for change. Experimentation with less familiar knowledge elements increases the likelihood of failures, but such failures provide important information about underlying relationships that were previously unknown. Feedback from experience that combines unfamiliar knowledge elements also helps firms avoid competency traps (Leonard-Barton, 1992; Levitt & March, 1988) and facilitates boundary spanning explorations (Rosenkopf & Nerkar, 2001). Experiments that incorporate less familiar knowledge elements may seem detrimental but feedback from these failures exposes other dimensions of technological landscape. By experimenting with knowledge elements previously unknown, a firm can considerably change its position in the landscape and explore other parts of the landscape. A firm's experience with exploration indicates its inclination to break through the

current technological paradigm and incorporate necessary changes in the existing knowledge base.

In this study, however, I specifically emphasize the role of search behavior in failure experience. Failures can result from experiments involving familiar elements or new elements. Familiarity of knowledge elements underlying failures is not discreet but continuous in nature and signifies the extent of a firm's experience with these elements, i.e. how frequently and recently elements that are part of failures have been used by the firm (Fleming, 2001; Vincenti, 1990).

Experimentation with new elements is full of surprises and considerably more uncertain than experimentation with old or familiar elements (Fleming, 2001). Firms usually engage in experimentation with new knowledge elements for various reasons such as (1) to solve a problem for which solution does not exist in present domain and (2) to gain or maintain competitive advantage in a volatile market. Experimentation with familiar knowledge elements, on the other hand, is more certain as firm is already aware of at least some of the relationships among these elements.

Owing to the prior experience with familiar elements, a firm has some knowledge about the outcomes of experiments, and this knowledge can help a firm design experiments in a more effective manner. To provide an example relevant to the present context, consider a pharmaceutical firm that is trying to find a treatment for a certain disease. Let us also assume that the firm already has some knowledge about the disease, i.e. it is aware of some knowledge elements that can interact with the disease-specific receptors or genes. The firm can rely on already known knowledge elements and recombine them to find the potential treatment. Such experiments are often easier to design and in most cases their outcomes can

be predicted with certain degree of certainty. However, it is possible that the best treatment of the disease requires knowledge elements that are currently not used by the firm or are used relatively less, i.e. they are less familiar to the firm. In such circumstances, if a firm does not search outside its knowledge base and continues relying on familiar elements, it will not be able to find an effective treatment of the disease.

Experimentation involving familiar elements is more certain in nature and less likely to lead to futile inventions (Fleming, 2001). High familiarity of elements is also an indication of their recent use and hence ensures less forgetting and effective subsequent usage (Argote, Beckman, & Epple, 1990). Experimentation with unfamiliar elements, however, is uncertain in nature and scholars have emphasized the value of recombining familiar elements to avoid the uncertainty (Altschuler, 1998; Mead & Conway, 1980). Strategy to rely on familiar elements to create inventions sounds promising, but unfortunately advantages of working with familiar elements do not last forever. Firms will eventually exhaust the useful combinations or at least will believe that they did either because of their inability to find further useful combinations or presence of social bias (Henderson, 1995; Kim & Kogut, 1996; Sahal, 1985)

In summary, large number of failures in familiar elements indicates that the firm attempted to find solutions by recombining old elements but did not succeed. Such failures indicate a firm's incompetence in combining familiar elements and reluctance in experimenting with new elements. However, failures that comprise new elements indicate a firm's attempt to explore new parts of the landscape and can often present novel solutions. Since these failures represent experiments with new elements, they are expected to occur and are an excellent source of new information (Fleming & Sorenson, 2004). Even though a firm

will experience more failures when experimenting with new knowledge elements, the quality and type of feedback generated by these failures will lead to more learning opportunities and therefore better subsequent R&D performance. This leads to the next hypothesis:

Hypothesis 3: A firm's experience with failures in more familiar knowledge elements is negatively associated with its subsequent R&D performance.

3.4 Familiarity of Knowledge

In the previous section, I examined whether elements associated with failures are familiar to the firm and what can be the implications of such differences for learning. In this section, I argue that familiarity of knowledge will also be instrumental in learning from these failures, that is, while using more familiar knowledge a firm relies on knowledge that resides within its domain. Failures characterized by use of less familiar knowledge, however, are built on the knowledge that lies outside a firm's domain. In this study, I consider use of more familiar knowledge a result of local search; whereas use of less familiar knowledge indicates a firm's attempt to rely more on distant search (Rosenkopf & Nerkar, 2001).

Knowledge utilized in experimentation can be either close to firms' existing knowledge or distant in nature. Firms can perform local search in which they try to build on something they already know. Local search allegedly is the most common method in the innovation process and is also referred to as exploitation by researchers (Cyert & March, 1963; Nelson & Winter, 1982; Stuart & Podolny, 1996). Essentially, in local search, firms rely on the knowledge they already possess. This process is incremental in nature, i.e. in each step; firms explore new knowledge that is similar or local to the existing knowledge in firms. Local search is a widespread phenomenon in drug discovery process, in which pharmaceutical firms scan through thousands of molecules in high-throughput screening, and

after selecting few molecules that match properties of the potential drug, firms make small changes in these molecules to produce commercially viable drugs (Drews, 2000). In the presence of uncertainty and lack of knowledge of elements and their combinations, inventors often rely on local search to develop greater understating of technology landscape (March & Simon, 1958; Vincenti, 1990).

However, too much reliance on local search prevents inventors from discovering novel solutions. Especially, when interdependency among knowledge elements is high, the technology landscape consists of many peaks, and local search can stall the process of innovation by exhausting all possible combinations in the neighborhood (March, 1991; Fleming, 2001). Distant search, on the other hand, involves use of knowledge that is outside of a firm's boundaries. Experimentation with knowledge that is beyond a firm's expertise is risky in nature because of the uncertainty and ambiguity involved in the process. At the same time, such experiments expose the firms to new knowledge that can be used to combine elements in more innovative ways. Searching beyond the knowledge boundaries of a firm leads to better innovation outcomes (Rosenkopf & Nerkar, 2001).

Local search and distant search can be viewed as exploitation and exploration, and the firm should maintain a balance between these two activities to remain competitive in the market (March, 1991). Experiment with knowledge distant to a firm's core competencies will often result in failures but will also improve its overall understanding of technology landscape. Feedback from such failures can help the firm combine knowledge elements in a more creative way, leading to increase in R&D performance. Firms can greatly benefit from such information, especially when they are trapped on the local peak in the landscape, and local search can only hurt their performance.

It does not imply that distant search will always result in better performance. In fact, more often it will lead to failure, but feedback generated in the process is invaluable to the process of innovation, and can help firms greatly in finding the optimal solution. Ingenuity of distant search, however, cannot compensate for steadiness and certainty of local search in innovation. Both are crucial to the firm, and while excessive exploitation can lock out novel technologies and render the firm in suboptimal equilibria, too much exploration will prevent firms from reaping benefits of experimentation by generating too many ideas without any distinct competence (March, 1991).

Prior research has examined how local vs. distant search can have differential effect on firms' innovation outcomes. However, in this study, I specifically focus on failure experience of a firm and not its overall experience. Following the arguments made in previous section, and considering the nature of local and distant search discussed in this section, firms will experience more failures when experimenting with distant knowledge. Such experience generates valuable feedback and can help the firm locate itself on a higher peak in the landscape, especially when landscape is rugged in nature. When landscape is rugged, the likelihood of getting trapped on a local peak is high and one of the ways to breakthrough from such a situation is to rely on distant knowledge.

What do the failures in local vs. distant knowledge imply for learning from them? I make a case in this study that the knowledge a firm uses to create inventions is different from type of elements (familiar or unfamiliar) discussed in previous section. For example, if knowledge elements are ingredients to an invention, then knowledge is its recipe. Though it is possible that use of distant knowledge can introduce firms with new elements, this is not

always true. A firm can use distant knowledge to combine familiar elements and rely on local knowledge to experiment with new elements.

Following the nature of local search discussed in this section, firms should experience fewer failures as the experiments will utilize more familiar knowledge. Therefore, large number of failures as a result of local search indicates that a firm's current knowledge is outdated and it needs to rely more on distant knowledge. Such failures do not lead to valuable information and therefore are not a great source of learning. However, distant search involves use of new knowledge and is risk prone. Failures following distant search are expected to occur and increase a firm's understanding of the technological landscape. Such experiences generate valuable feedback and can help a firm locate itself on a higher peak in the landscape. They also indicate a firm's willingness and motivation to look beyond its knowledge domain to find better solutions. Overall, higher number of failures resulting from experimentation targeted at distant search as compared to local search characterizes a firm's willingness to experiment with distant knowledge and ability to leverage local knowledge effectively.

Hypothesis 4: A firm's experience with failures utilizing more familiar knowledge is negatively associated with its subsequent R&D performance.

3.5 Decomposability of Knowledge Base

Within an industry, firms search over the same technology landscape. Based on the initial choices firms make with respect to the knowledge elements and their combinations, they locate themselves at different positions in the landscape. The set of knowledge elements and how they are combined are referred to as the knowledge base of a firm. In other words, a knowledge base constitutes of knowledge elements linked with each other in a certain way.

Structure of a knowledge base, therefore is the nature of linkages among knowledge elements. The search behavior and structure of a knowledge base are inherently interrelated. Structure of a knowledge base is a result of the search behavior of the firm. In the absence of knowledge about true relationships between elements, firms attempt to combine these elements based on their respective understanding of the landscape, and such attempts are responsible for the structure of knowledge base. This is one of the reasons for observing considerable heterogeneity in firms' knowledge base even when they are searching in the same technology landscape.

Earlier research has largely overlooked the structural aspects of the knowledge base and primarily focused on the size of knowledge base and its effect on innovation related outcomes (Ahuja & Katila, 2001; Fleming, 2001). Scholars have also explored the implications of overlap in knowledge bases of firms for their ability to absorb knowledge from their neighbors (Lane & Lubatkin, 1998; Mowery, Oxley, & Silverman, 1996). However, how elements in knowledge bases are related to each other and how difference in structure of the knowledge base can impact innovation activities are rather less studied (Yayavaram & Ahuja, 2008). How firms decide to combine elements in the knowledge base contributes to the structure of knowledge base. Thus, structure of the knowledge base serves as a proxy for firms' understanding of technological landscape. For example, scientists in firm A may choose to combine elements x and y always in addition to other elements, whereas scientists in firm B may decide not to do so and instead combine elements x and z or some other combination. A firm's decision to combine certain elements and not others represents its belief regarding which elements work well together and vice versa (Yayavaram & Ahuja, 2008).

How firms choose to combine elements in the knowledge base can differ in a meaningful way. Decision to use certain elements together in an invention should not be confused with the true relationship among these elements. There is an underlying relationship among elements, known as interdependence, in each landscape, and firms' choice of combining elements does not necessarily match with true relationships (Yayavaram & Ahuja, 2008; Yayavaram & Chen, 2013). Firms' understanding of the landscape contributes to their decision to prefer some combinations over other (Simon, 1996), which is the primary source of variation among knowledge bases of different firms. There may be an inherent relationship between elements x and y such that any changes made in x causes the change in contribution of y to the overall performance, but not necessarily all firms are aware of this relationship, and therefore we may observe only some firms using this particular combination instead of all firms.

Firms can also combine certain elements without consciously acknowledging the actual interdependencies. Considering extremely large number of ways in which firms can combine elements in their knowledge base, it is reasonable to assume that firms will have substantially different knowledge bases with respect to the structure. As a result, structures of some of the knowledge bases are more decomposable or modular in nature than structures of other knowledge bases. On the other hand, some knowledge bases may look similar to a large connected cluster, i.e. highly integrated or non-decomposable. Structural properties of a knowledge base can play an important role in incorporating the information that firms receive following failures and therefore can have serious implications for learning from failures.

One of the key aspects of learning from failures is firms' ability to not just identify necessary changes, but successfully integrate them in the existing knowledge base.

Therefore, it is critical for knowledge bases to be able to undergo the required changes.

Knowledge bases of firms are not static in nature and evolve over time, and depending on the information from the failures, firms may need to dissolve old ties because these ties may not be as productive as the firm thought previously. Also, information from failures can indicate potential for new ties and motivate the firm to incorporate them in its knowledge base.

Structure of a firm's knowledge base can be instrumental in its ability to incorporate changes, be it dissolving old ties or introducing new ties, and could be one of the reasons for some firms being able to learn from failures more effectively than other firms.

Knowledge base of a firm can exhibit several patterns including decomposability, hierarchy, scale-free, and small-world, which can influence the firm's ability to make required changes in its knowledge base (Rivkinn & Siggelkow, 2007; Yayavaram & Chen, 2013). In this dissertation, I examine one of the structural properties of a knowledge base, specifically its decomposability, for two reasons. First, previous studies have looked at the effect of decomposability on product architecture and organization structure (Baldwin & Clark, 2000; Ethiraj & Levinthal, 2004; Sanchez & Mahoney, 1996; Schilling, 2000), but how decomposability can be a limiting factor in learning from failures has not been studied. Second, prior research has also explored the role of decomposability in generating more innovative products, specifically in turbulent environment (Baldwin & Clark, 2000; Ulrich & Eppinger, 1999). Yayavaram and Ahuja (2008) show that decomposability of knowledge base affects usefulness of inventions such that nearly decomposable knowledge base

increases the usefulness of firms' inventions. They also find that near-decomposability increases the knowledge base's capacity to change.

These studies have certainly increased our understanding of the role of decomposability in product and R&D innovation but lacks in-depth understanding of mechanisms involved behind such effect. Why do firms decide to change some combinations among elements in the knowledge base? How do firms learn about new combinations, and what is the motivation behind such changes? In this dissertation, I intend to investigate one such mechanism, i.e. learning from failures. There can potentially be several reasons for firms to dissolve some ties and introduce others, and one of the reasons is that firms receive information from failed experiments and use this information to improve ongoing projects.

For example, if a firm has previously combined elements x and y successfully and, in recent experimentation, found that combination of x and y is not valuable anymore, it will identify inventions that use this particular combination and see whether it is possible to improve those inventions or it is best to discontinue them. Firm can also experiment with completely new elements u and v instead. Failures provide important information and necessary impetus to make changes in the old knowledge but structure of existing knowledge base can facilitate or inhibit such initiatives. Studying the role of structure of knowledge base in learning from failures will further increase our understanding of innovation process and uncover one of the sources of changes in a firm's knowledge base.

Decisions to create new ties and dissolve old ties among knowledge elements are the primary sources of variation in structure of knowledge bases across firms. The resulting structure is rarely entirely decomposable or completely integrated but rather lies somewhere in between. In other words, the decomposability of a structure is a continuous property of

knowledge base and can lie anywhere between 0 and 1 with 0 being non-decomposable and 1 being entirely decomposable structure. Knowledge base that has all of its elements highly clustered or densely connected to each other is non-decomposable in nature. A highly decomposable knowledge base, on the other hand, has its elements divided into modules such that within modules elements will be densely connected whereas elements across modules will be seldom connected.

Knowledge bases that have discernible modules, and at the same these modules are connected through ties between elements across modules are nearly decomposable (Simon, 1962). Near-decomposability is one of the key properties of complex systems and provides various systems including biological, physical, and social an adaptive advantage. Near-decomposability plays a crucial role in emergence of complex systems by providing stable intermediate forms in which modules are not affected by other modules in the short run, but in the long run, the system as a whole can adapt because of the cross-module linkages (Simon, 1962). Following Simon's conceptualization of near-decomposability, research on modularity (Baldwin & Clark, 2000; Schilling, 2000) and loose coupling (Weick, 1976) also demonstrated that nearly decomposable systems are adaptive, i.e. able to undergo necessary changes and persistent in nature.

Nearly decomposable knowledge base enables firms to balance exploitation and exploration by facilitating both deep knowledge of modules and exposure to variations in existing knowledge through ties between modules. Literature on product innovation has shown that firms often develop the most novel innovations when they are specialists in their core knowledge, engage in exploration of new knowledge, and have an arrangement in place that can integrate existing knowledge and new knowledge attained through exploration

(Gupta, Smith, & Shalley, 2006; Katila & Ahuja, 2002). Nearly decomposable knowledge base provides unique advantages to firms by generating superior knowledge and deeper understanding of knowledge elements within modules, exposing new combinations among elements and ideas through exploration, and enabling firms to assimilate new knowledge into existing knowledge base. Such mechanisms facilitate independent development of modules without experiencing too much disturbance from other modules, but at the same time cross-modules linkages ensure the spread of innovation across modules in the long run.

In this dissertation, I argue that firms receive the information from failures and incorporate it into existing projects to improve them in any way possible. I further argue that availability of information from failures by itself is necessary but not a sufficient condition for learning from them. Ability of firms to make necessary changes in the existing knowledge base can be a critical limiting factor in such learning. Near-decomposability of knowledge base will not only allow firms to make required changes in respective modules, but also provides a unique opportunity to spread such learning across modules. Therefore, I believe that firms with near decomposable knowledge bases will leverage the learning from failures most and that will be reflected in their future innovative performance.

Previous research has assumed that near decomposable knowledge base results in better innovations because of its ability to both adapt and persist at the same time, but by specifically looking at its role in firm's ability to learn from failures, current dissertation provides empirical evidences of the above belief and uncovers one of the ways firms can enhance their learning capabilities. In addition to the prior knowledge that nearly decomposable knowledge base is promising to better inventions, this study offers insights into how firms can leverage the structure of their knowledge base. By understanding the

mechanisms involved in the process of innovation and role of structure of the knowledge base, firms can better manage their innovations and take greater advantage of their knowledge base.

One of the aspects of learning from failures that has largely gone unnoticed is a firm's ability to effectively incorporate learning from failures. Failures in experimentation helps firms uncover underlying relationships among knowledge elements, but given inertia and path dependency, firms may not be able to incorporate what has been learned from failures. This is particularly true and therefore interesting to study in high technology industries such as pharmaceutical, where number of knowledge elements is high and extent of interaction among these elements can be large. As firms begin to explore true relationships among knowledge elements in the landscape, next step is to build and apply what has been learned effectively. The introduction of new architecture based on the learning from failures requires investment of time and resources (Louis & Sutton, 1991). In addition, to assimilate the learning from failures, a firm's knowledge base must be capable of going through the required changes (Yayavaram & Ahuja, 2008).

Knowledge base of a firm is equivalent to complex adaptive system that evolves over time (Fleming & Sorenson, 2001). Elements in the knowledge base interact with each other to varying degrees, and these interactions change as firms' understanding of relationships among elements in technology landscape improves. In the context of this study, firms observe the failures of patents, and each patent represents a specific combination of certain knowledge elements. To incorporate information from these failures, firms must be able to make necessary changes in the existing knowledge base. One of the desirable properties of complex systems, knowledge base in this case, is the near-decomposability (Simon, 1962).

Consider an example of a knowledge base that is not decomposable, i.e. completely integrated in nature. In such knowledge base, each element interacts with all the other elements. Now let us say a firm learns from previous failures that a certain combination of elements does not provide desirable outcomes, and it decides to discontinue the projects that encompass this particular combination of elements. This is impossible to achieve in an entirely integrated knowledge base because each element is connected with all other elements, and removing a particular combination implies discontinuing all ongoing projects in the firm. On the other extreme, when knowledge base is completely decomposable, it is easy to incorporate learning from failures and make necessary changes, but such learning will result in limited performance increment as it is constrained to a single module in the knowledge base, and even though it is possible that the changes made in this module can potentially benefit other modules, there is no way to find this out in a completely decomposable knowledge base.

In order to implement learning in the most effective way, a firm's knowledge base should have distinct modules, and these modules should be connected with each other through elements across modules, a property of nearly-decomposable structures (Simon, 1961). Configuration of elements in a nearly-decomposable knowledge base facilitates the implementation of learning from failures. Separate modules in the knowledge base enables the changes within modules without disturbing other elements too much, and links between modules allows the transfer of learning across modules.

By its very nature, knowledge base that is entirely decomposable or non-decomposable indicates that a firm is either missing critical linkages or considering some irrelevant linkages among elements, respectively (Yayavaram & Ahuja, 2008), which makes

the process of incorporating learning and making subsequent changes ineffective. For example, in case of an entirely decomposable knowledge base, though a firm will find it easier to incorporate learning within modules, such learning will be limited in nature in the absence of critical cross-module linkages. However, when knowledge base is non-decomposable (or integrated), a firm will find it difficult to incorporate the learning due to the presence of too many linkages among elements. A firm with knowledge base that exhibits moderate levels of decomposability will not only find it easier to incorporate learning within modules, but cross-module linkages will also lead to greater learning benefits in a long run. Therefore, I hypothesize:

Hypothesis 5a: At moderate (as compared to low and high) level of decomposability of knowledge base, benefits of relatedness of (and among) failures for R&D performance will be more, i.e. the positive slope of proposed inverted-U shape relationship will be steeper.

Hypothesis 5b: At moderate (as compared to low and high) level of decomposability of knowledge base, negative effect of too much relatedness of (and among) failures on R&D performance will be less, i.e. negative slope of proposed inverted-U shape relationship will be flatter.

4. RESEARCH DESIGN AND METHODOLOGY

The dissertation explores three aspects of learning from small failures in experimentation. First, implications of relatedness of failures in a firm's knowledge base for learning are investigated. Second, role of search behavior associated with failures in learning is examined. Third, given failures do provide learning opportunities, the role of structure of a firm's knowledge base in incorporating learning is studied. I focus on patented R&D efforts of firms in the pharmaceutical industry to study various aspects of failures and their role in learning. Drug discovery process in the pharmaceutical industry involves high degree of experimentation, and patents represent outcome of the experimentation process (Thomke, 2003). The only way firms in pharmaceutical industry can improve their understanding of underlying science is through the process of experimentation in which different knowledge elements and their combinations are tested time and again, and the ones that show potential are developed further.

4.1 Drug Discovery and Pharmaceutical Industry Context

The dissertation aims to explore the three aspects of learning from small failures in experimentation. First, implications of relatedness of failures in a firm's knowledge base for learning are investigated. Second, effect of search behavior associated with failures on learning is examined. Third, given failures do provide learning opportunities, the role of structure of a firm's knowledge base in incorporating such learning is studied. The ideal set-up to answer these questions calls for a scenario in which: (1) firms are bound to experiment in order to gain competitive advantage, and (2) failures resulting from these experiments are

small in terms of their consequences.

Drug discovery process in the pharmaceutical industry satisfies this criterion better than any other process to my knowledge (Thomke, 2003). The only way firms in the pharmaceutical industry can improve their understanding of underlying science is through the process of experimentation. During experimentation, different knowledge elements and their combinations are tested time and again, and many of these experiments are considered failures, primarily because of their perceived lack of value by the firm. As these failures are frequent, and firms experience large number of these failures on a regular basis, it is reasonable to assume that they are small in terms of their consequences. Also, failures during experimentation do not threaten firms' existence and do not require firms to make major changes, further suggesting that these failures are small in nature.

Firms in pharmaceutical industry rely on their respective proprietary chemical libraries, which consist of chemical compounds that are tested against disease targets. Except rarely, when a chemical compound represents a single chemical group, each chemical compound is a combination of a set of chemical groups. These chemical groups characterize knowledge elements that are combined by firms during experimentation.

A chemical library defines a firm's experimental search space in which the firm looks for solutions (Thomke & Kuenmmmerle, 2002). Though, it is possible that there exists an overlap between the search spaces of different firms, it is unlikely for firms to have same search space and search trajectory. It is believed, and reasonably so, that there exists a true relationship between various knowledge elements, and firms engage in experimentation to uncover these relationships (Yayavaram & Chen, 2013). Firms differ in terms of knowledge elements they choose to experiment with initially. In case there is an overlap between

knowledge elements of two firms, they can still differ in the ways they choose to combine these elements. In this sense, firms not only search in different parts of the experimental space but often search along different trajectories as well. The only way for firms to uncover true relationships among different knowledge elements is to experiment. Even though most of these experiments fail, because true relationships among elements will always be much less than all the possible relationships, firms have no other way to learn about the search space and must experience these failures.

I discuss drug discovery process briefly to support the choice of the pharmaceutical industry to answer questions in this dissertation. As mentioned, firms experiment with chemical compounds that are part of their respective proprietary chemical libraries. Chemical compounds in these libraries are tested against the disease target of interest. The process is similar to trying all the keys to unlock a newly discovered lock, but much more complex in nature (Thomke & Kuemmerle, 2002). The capability to successfully experiment with different knowledge elements, and access to diverse chemical libraries are known to provide significant competitive advantage to firms in the pharmaceutical industry (Iansiti, 2000; Thomke, 1998, 2001; Thornke et al., 1998). Chemical libraries represent years of commitment and tacit knowledge on the part of the firms. GSK, for example, is one of the top players in respiratory diseases, and in addition to scientific expertise and understanding of respiratory diseases; it has the ability to identify chemical compounds that can act on disease targets effectively. Similarly, Eli Lilly has developed a chemical library focused on diseases related to central nervous system (Thomke & Kuemmerle, 2002).

Although, nearly all large pharmaceutical firms focus on all diseases to some extent, they tend to be experts in some disease areas, and such expertise is a result of years of

experimentation and resource commitment. Pharmaceutical firms differ in terms of which elements they experiment with, and how they combine these elements. Therefore, even when firms experiment with same elements, they can differ in terms of combination of elements used.

As discussed, before pharmaceutical firms patent or test their drugs in clinical trials, there is long process of pre-discovery and pre-clinical research to identify suitable drugs (PhRMA, 2007). There are four major steps in this process. First, firms try to understand the disease and its underlying causes. This involves study of changes in genes and how these changes can subsequently lead to the disease. Second, scientists look for a ‘target’, one of the altered genes or molecules, which can interact with the potential drug. Third, the identified target is validated for its role in disease and its successful interaction with the drug molecule. Fourth, once scientists have some understanding of the disease and potential drugs, they begin the process of finding lead compounds or drugs that can be used to treat the disease. After conducting preliminary safety tests and optimization studies, firms select few compounds that are further tested in preclinical trials, which establish the safety of drugs in animals before they can be tested in humans.

On an average, of the 5,000-10,000 compounds tested in the fourth step, only around 250 are selected for preclinical testing (lead compounds). After testing lead compounds extensively in preclinical studies, firms select 1 to 5 compounds that enter clinical trials. These compounds are called ‘drug candidates’. After identifying handful of ‘drug candidates’, firms file investigational new drug (IND) application with Food and Drug Administration (FDA) to begin clinical trials on humans. Eventually, one of the five drug candidates successfully clears all four stages of clinical trials and is commercialized in the

market for the treatment of the disease. Broadly, of the 5,000 to 10,000 chemical compounds in the fourth step, pharmaceutical firms file patents for around 250 lead compounds. Of these 250 lead compounds, only up to 5 are tested in clinical trials, and out of these five, only one eventually becomes a product (PhRMA, 2007; Thomke & Kuemmerle, 2002).

I study small failures in the form of premature expirations of patents in the pharmaceutical industry. Premature expiration of a patent is defined as a firm's decision to not pay the maintenance fees and give up its right on the patent before 20 years. Data on patent expirations provides ideal settings to study small failures in experimentation. In the first stage of discovery, thousands of chemical compounds are tested using high-throughput screening and most of them are discarded. At this stage, firms have no expectations from any particular compound, and the sole purpose is to select lead compounds that show some potential. Since firms expect to select only several lead compounds from this stage, rejection of most of the compounds is not considered a failure but part of the process. Pharmaceutical firms tend to patent lead compounds before or during the pre-clinical trials. Although, firms are aware that most of the lead compounds (~98%) that go through pre-clinical trials will be discarded, they have similar expectations from all these compounds and also commit considerable amount of time and resources during the preclinical trials of each compound.

Failures at this stage mirror closely with the idea of small failures in experimentation studied in this dissertation and are perfectly captured in patent expiration data. The failure of a lead compound in R&D process suggests that certain elements or their combination did not result in desirable outcomes, and the firm needs to further experiment in order to develop a successful drug except that the firm now has more information resulting from failures. Data

on patents and their expirations allow us to study a firm's R&D efforts into projects that once showed some potential and lived up to the firm's expectations.

To summarize, premature expirations of patents, i.e. voluntary decision of firms to give up rights on their intellectual property provide ideal settings to study learning from small failures in experimentation for following reasons: First, patents are clear indication of a firm's R&D efforts and provide understanding of experimentation carried out by firms. Second, since patents are not products themselves; in fact, very few of them are ever commercialized, patent expirations do not threaten a firm's existence in a significant way. Also, firms studied in this dissertation are mid-to-large size and discontinue patents frequently and in large numbers. Therefore, patent expiration can be ruled out as a large failure. Third, the decision to discontinue patents is voluntary in nature. Though, it is possible that in some cases such decisions are influenced by the feedback from the market, the decision is not forced upon a firm by any external authority. The voluntary nature of decision making behind patent expirations further makes a case to study organizational learning in this context because of two conflicting theories around this topic: one theory suggests that volition is positive for learning because autonomous decision making promotes commitment and problem solving, while the other theory argues that externally imposed failures are better for learning because they come as a shock and help firms overcome inertia, facilitating deeper analysis (Haunschild & Rhee, 2004). Fourth, unlike the case of large failures that can threaten a firm's existence, or operational failures that happen after product launch where individuals are averse to associating themselves with failures, small failures in experimentation are result of search carried out by firms, and these failures are expected outcomes of experimentation. The purpose of the firms is not to minimize these failures but

to improve their understanding of underlying science based on the information from these failures. Experimental nature of small failures helps individuals in firms overcome biases associated with failures, facilitating learning from the failures. Finally, patents provide information on the underlying knowledge elements and their combinations. Small failures in the form of patent expirations reveal information about the elements and combinations that did not result in desired outcomes. Such detailed information is necessary to study relational characteristics of failures, search behavior associated with failures, and structural characteristics of knowledge base that are the focus of this dissertation. In short, not only patent expirations provide an opportunity to study learning from small failures in experimentation, but the information contained in patents also facilitates the detailed analysis of mechanisms involved in such learning.

4.2 Data and Methods

I test the hypotheses with data on patent expirations in the pharmaceutical industry. I define failure as the expiration of a patent before its natural life of 20 years. In 1980, The United States Patent and Trademark Office (USPTO) made it necessary for firms to pay maintenance fees every 4, 8, and 12 years to keep their patents active. A firm or entity failing to pay the fees at the scheduled time will have its patent expired. The most likely reason of early discontinuation of a patent is the lack of value as perceived by a firm (Serrano, 2010).

The sample consists of all patents filed by 76 pharmaceutical firms between 1975 and 2005. In line with previous research, I considered 3-digit USPTO classes 514 and 424 to identify patents in the pharmaceutical industry (Anand, Oriani, & Vassolo, 2010; Guler & Nerkar, 2012). These patents belonged to more than 200 firms. Some of the topics that I explore in this study include relatedness of failures and structure of firms' knowledge bases,

and therefore require reasonable amount of patenting activity from firms. For example, study of relationship among failures for firms that have only one failure is irrelevant. Similarly, knowledge base of a firm with very few patents will not exhibit any distinct structure. Therefore, questions explored in this dissertation are only applicable to firms that are reasonably active in patenting and terminations. For this reason, I removed the firms that had a total of less than 50 patents, or did not patent consecutively for 10 years between 1980 and 2002, leading to the final sample consisting of 76 pharmaceutical firms.

Table 1 lists all pharmaceutical firms used in this study. USPTO provides information on all the expired patents, including the stages at which they were discontinued, i.e. first stage (after 4 years of grant date), second stage (after 8 years of grant date), and third stage (after 12 years of grant date). After the introduction of maintenance fees regime in 1980 by the USPTO, the first year in which patents expired was 1985 (law was applicable to patents granted in year starting 1981), so the dataset captures all expired patents starting year 1985. Figure 1 shows the premature expiration of patents between 1985 and 2002 for 76 pharmaceutical firms analyzed in the current study. Of 47,625 patents granted to the 76 firms in the sample, 22,626 patents (48%) expired as of 2005. Since I measure learning at the firm level, I aggregated the number of expired patents at the firm-year level, leading to a final panel that contains 1731 firm-year observations¹.

4.2.1 Dependent Variable(s)

The dependent variable used in this study is the R&D performance of a firm. I use both the productivity and the quality of R&D outcomes to measure R&D performance of the

¹ Pharmaceutical industry experienced big wave of mergers and acquisitions (M&As) during the time period of this study. In case of M&A, I kept the firms as separate entities before M&A occurred and combined them into a single entity afterwards.

firm. I measure productivity as the number of patents granted to a firm in a given year. Patents are considered a source of competitive advantage in many industries as they provide the patent holder the right to exclude other firms and individuals from using the knowledge covered by patents (Cockburn, Henderson, & Stern, 2000). For firms in industries such as pharmaceuticals, patents are direct results of innovation.

I measure the quality of the R&D performance as the total number of citations to the patents of a firm in a given year. Numerous studies provided evidence of correlations between importance of patents by technical specialists and citations to patents, and have established the use of citations as legitimate proxy for quality of innovative or inventive performance (Pavitt, 1988; Jaffe, Trajtenberg, & Henderson, 1993; Trajtenberg, 1990).

4.2.2 Independent Variables

Five key independent variables are used to test the hypotheses developed in this study. Below, I provide list of all independent variables and how I measure them.

4.2.2.1 Relatedness of failures and relatedness among failures

Hypotheses 1 and 2 use *relatedness of failures* and *relatedness among failures*, respectively as independent variables. Hypothesis 1 uses relatedness of failed patents with patents that did not fail yet, whereas Hypothesis 2 uses relatedness of failed patents with each other. That said, the methodology used to estimate the relatedness is same in both cases. I borrow from the previous studies (Breschi, Lissoni, & Malerba, 2003; Makri, Hitt, & Lane, 2010; Nesta & Paolo Saviotti, 2005; Teece et al., 1994) and calculate the relatedness of patents based on their subclasses and their linkages with other subclasses. Specifically, relatedness of a failure is the extent to which its subclasses are combined with other subclasses in a firm. Higher relatedness of a patent indicates that its subclasses are combined

with large number of subclasses in other patents. Scholars in strategy have also used the same measure to capture interdependency of a patent but with the assumption that higher extent of relatedness among subclasses implies the interdependence of these subclasses with each other (Fleming & Sorenson, 2004, Ganco, 2013). Subclasses used to measure relatedness represent knowledge elements discussed throughout the dissertation. The measure is estimated by first calculating relatedness of subclasses underlying a failed patent, and then normalizing it by number of subclasses in the patent. *Relatedness of failures* at the firm level is constructed by averaging relatedness of all failed patents for each firm in a given year.

$$Relatedness\ of\ patent\ a = \frac{\sum_{i \in a} \sum_{j \notin i} \frac{\text{count of patents in subclasses } i \text{ and } j}{\text{count of patents in subclass } i}}{\text{count of subclasses of patent } a},$$

where i belongs to subclasses in patent a , and j belongs to all other subclasses, except i , that are combined with subclass i and present in non-failures, i.e. patents that did not expire yet. Figure 2 provides an illustration of the *relatedness of failures* by showing the relationships of knowledge elements underlying failures (colored red) with other elements (colored blue) in the knowledge base for two pharmaceutical firms in 1995. The graph depicts a network of subclasses in which a firm is patented, with nodes representing subclasses and edges denoting common patents. Two nodes (or subclasses in this case) are connected to each other if they occur or are combined in the same patent. Higher relatedness of a patent indicates greater number of edges between subclasses of that patent and other subclasses. As shown in Figure 2a, for Fisons Pharmaceuticals, subclasses underlying failures are moderately related to other subclasses in knowledge base; whereas in case of Rhône-Poulenc, the relatedness of failures is quite high as the subclasses underlying failures are combined with many other subclasses (Figure 2b).

Relatedness among failures is also measured using the same methodology, except that now j belongs to all subclasses, excluding i , that are combined with i and present in failed patents. Figure 3 illustrates the same graph but only considers subclasses that are part of failures in a given year, representing *relatedness among failures*. Figure 3a shows the relatedness among elements underlying failures for Bristol-Myers Squibb (BMS) in 1991, and suggests that the firm did not combine the subclasses with each other often, thus indicating lower degree of *relatedness among failures*, compared to the failures in 1995 (Figure 3b), when the firm combined the subclasses extensively with each other, indicating higher degree of *relatedness among failures*.

4.2.2.2 Familiarity of knowledge and knowledge elements

Hypotheses 3 and 4 use search behaviors as independent variables in the form of familiarity of elements (old vs. new) and type of knowledge (local vs. distant) underlying failures, respectively. To estimate familiarity of a failed patent, I rely on a measurement used by Fleming (2001) that takes into account both frequency and recency of use of subclasses in the patent. After calculating familiarity of each patent, I take average of familiarity of all failed patents for each firm in a given year to construct the independent variable for hypothesis 3. Also, following Fleming (2001), I use the time constant of knowledge base as 5 years that signifies loss of two-thirds of knowledge in 5 years or 18% per year.

Familiarity of patent a 's subclass $i \equiv I_{ai} = \sum_{\substack{\text{all patents } b \\ \text{granted before patent } a}} 1 \{ \text{patent } b \text{ uses subclass } i \}$

$$\times e^{-\left(\frac{\text{application date of patent } a - \text{application date of patent } b}{\text{time constant of knowledge base}}\right)}$$

$$\text{Familiarity of patent } a \equiv F_a = \frac{\sum_{\substack{\text{all subclasses } j \\ \text{of patent } a}} I_{aj}}{\sum_{\substack{\text{all subclasses } j \\ \text{of patent } a}} 1}.$$

To calculate the independent variable used in hypotheses 4, which is the type of knowledge used in failed patents, I rely on the measure used in Katila and Ahuja (2002) that captures the extent to which a firm uses new knowledge in its inventions. As before, I estimate the distance of knowledge used in each failed patent and then take average for all failed patents for each firm in a given year.

$$\text{Familiarity of knowledge used in patent } a \equiv S_a = \frac{\text{Number of new citations}_t}{\text{Number of old citations}_t}$$

4.2.2.3 Decomposability of a knowledge base

Hypotheses 5a and 5b use interaction of decomposability of knowledge base and relatedness of failures as an independent variable. Knowledge base can be visualized as a network in which nodes are subclasses, and two nodes have an edge if they have at least one patent in common. To estimate decomposability of a knowledge base, I borrow from Newman (2004) that uses random network as base model to identify presence of modules in a focal network, which is the knowledge base of a firm in a given year in present context. Scholars in management have started to take advantage of above method to analyze the structure of various networks (Gulati, Sytch, & Tatarynowicz, 2012; Zhou, 2013). The algorithm functions by dividing the network in modules based on the most ‘between’ edges, i.e. edges with highest betweenness centrality. However, depending on the cut-off value of betweenness centrality, one will have the network divided into different number of modules. Therefore, next step is to measure the decomposability of all possible divisions of the network to find the optimum number of modules in which the network should be divided (Newman & Girvan, 2004). Decomposability, therefore, can be defined as:

$$D = \sum_{i=1}^n [e_i - e_{ir}],$$

where n is the number of modules in which the network is divided, e_i is the fraction of edges within module i , and e_{ir} is the fraction of edges in the module if edges were placed at random. Essentially, this formula estimates the difference between observed number of edges between nodes within modules and expected number of edges. Expected number of edges is estimated on the basis of random network in which edges within and between modules are equally likely. The algorithm is iterated for all values of n , and the highest value of D is the decomposability of network or knowledge base.

Figure 4 provides graphical representation of difference between the knowledge bases of Rhône-Poulenc and Fison Pharmaceuticals in 1991. Knowledge base of Rhône-Poulenc (Figure 4a) is composed of independent modules with little to none interactions between elements across modules. Firms with such knowledge bases will find it easier to incorporate learning and make changes in separate modules, but will be limited in terms of leveraging advantages of such learning across entire knowledge base because of the absence of cross-module linkages. Knowledge base of Fison Pharmaceuticals (Figure 4b), on the other hand, also displays presence of modules, but unlike the case of Rhône-Poulenc, these modules are connected with each other through linkages between elements across modules. Such structures enjoy the transfer of learning insights from one module to other modules.

4.2.3 Control Variables

I use several control variables that can have confounding effects on the dependent variable, R&D performance. One of the control variables is the size of a firm's R&D units, as large R&D teams in an organization are likely to have higher productivity. I calculate the size of the R&D team by counting the number of inventors that have patents for each firm in a given year from the USPTO database (McFadyen & Cannella, 2004).

Another control variable that can explain a firm's R&D performance is the number of alliances that firms make with other firms. Previous literature has established that alliances are a source of information and capabilities, and there exists a positive link between the extent of a firm's alliance activity and its R&D performance (Ahuja, 2000; Ahuja & Katila, 2001). I control for the number of alliances made by each firm between 1975 and 2005.

I also control for technological diversity. A large body of literature has shown that firms that diversify in multiple technologies not only benefit from the technological spillovers, but also buffer themselves against the risk of investments in a single technology (Jaffe, 1986; Vega, 2006). Also, diversification exposes a firm to new opportunities and promotes innovation (Nelson, 1959). I include the Herfindhal Index for subclasses in which a firm patents in a given year to control for the effect of technological diversity on R&D performance.

Geographic diversification can also affect innovation in multinational firms. Presence in countries other than the home country increases the availability of necessary resources and capabilities for innovation, and previous research has recognized the positive role of diversification in R&D performance (Kobrin, 1991). In order to account for the effect of geographic diversification on innovative performance, I include the count of the number of countries in which a firm patents in a given year.

I also control for lagged value of dependent variable and total number of failures to account for firm-specific strategies with respect to the way they conduct R&D. Figure 5 provides a brief description of all variables and includes a sample data point for Pfizer in 2000.

4.3 Empirical Model

The following models are used to test the relationships between relatedness of (and among) failures (H1 & H2), search behavior (H3 & H4), and structure of knowledge base (H5a & H5b), and R&D performance in a panel dataset of 76 firms for a period of 30 years:

$$P_{i,t+1} = \beta_1 P_{i,t} + \beta_2 F_{i,t} + \beta_3 C_{i,t} + u_{i,t} \quad (1)$$

$$Q_{i,t+1} = \beta_1 Q_{i,t} + \beta_2 F_{i,t} + \beta_3 C_{i,t} + v_{i,t} \quad (2)$$

In equation (1), P_{it+1} is R&D productivity, which is measured as number of patents filed by firm i in year $t+1$, $P_{i,t}$ is the lagged value of productivity for firm i , and $F_{i,t}$ are lagged expired patent characteristics, including relatedness of failures, relatedness among failures, familiarity of elements, type of knowledge, and interaction terms including relatedness of failures and decomposability of knowledge base, for firm i in period t . $C_{i,t}$ are lagged control variables, and include total number of expired patents, size of the R&D team, number of alliances, technological diversity, and geographical diversity. Equation (2) has same set of variables except that dependent variable Q_{it+1} is R&D outcome quality, measured as citations to the patents of firm i in period $t+1$, and $Q_{i,t}$ is R&D outcome quality for firm i in period t . I conducted Hausman test to choose between random effects and fixed effects model, and the test rejected the null hypothesis that difference in coefficients is not systematic, suggesting the use of fixed effects model over random effects.

The dependent variables used to measure R&D performance are count in nature, and Poisson model is an appropriate choice, but because of the presence of higher dispersion in dependent variable (DV), I instead use negative binomial model that accounts for over dispersion in DV. All five hypotheses are tested with fixed effects negative binomial model. Though, fixed effect model reduces degree of freedom substantially, especially when the

number of groups is large and the number of observations is small, it is preferred over population average and random effects model when time invariant characteristics of firms can lead to biased estimates, which is true in the present study, as validated by the Hausman test. I also include time dummies in the model to account for time-specific trends that can explain the R&D performance of the firms in the sample.

5. RESULTS

Table 2 presents the descriptive statistics and correlation matrix for all the variables used in the analysis. Number of observations for all variables is 1731, except for variables that are calculated using failures and alliances. The first time expired patents start to appear in the dataset is in the year 1985, leading to total number of observations for those variables fewer than other variables. Similarly, alliance data was not available for all the firms in the dataset. *Relatedness of failures* (with non-failures) varies between 0 and 35 with a mean value of 3; though, most of the values (95%) are below 8.3. *Relatedness among failures* is between 0 and 20 with a mean value of 1.7. Most of the values (95%) are below 4.2. *Familiarity of knowledge elements* ranges from 0 to 79 with a mean value of 4.5. *Type of knowledge* measures ratio of new citations to old citations, which ranges between 0 and 1 with a mean value of 0.32. Decomposability ranges from 0 to 0.9 with a mean value of 0.5.

Tables 3 and 4 contain models that are used to test the hypotheses 1–5 for productivity and quality measures of R&D performance, respectively. Model 1 in both tables is the test for control variables without any hypothesized independent variables. Model 2 comprises *relatedness of failures* and control variables (Hypothesis 1). Model 3 includes *relatedness among failures* along with control variables (Hypothesis 2). Models 4 and 5 incorporate independent variables *familiarity of knowledge elements* and *type of knowledge*, respectively

along with the control variables (Hypotheses 3 and 4). Models 6 and 7 test interactions of decomposability with *relatedness of failures* and *relatedness among failures*, respectively (Hypotheses 5a and 5b). Model 8 is the full model with all independent variables, interactions terms, and control variables, and it is also used to interpret coefficients of all independent variables.

Hypothesis 1 argues for a curvilinear relationship between relatedness of failures with non-failures and R&D performance and is supported for both measures of R&D performance. Model 2 in Table 3 indicates positive and significant coefficient for linear term, and negative and significant coefficient for squared term of relatedness of failures, suggesting that relatedness of failures has curvilinear effect on productivity measure of R&D performance, such that an increase in relatedness leads to an increase in R&D productivity, but beyond a certain point increase in relatedness begins to affect R&D productivity negatively.

Similarly, relatedness of failures has positive effect on R&D quality, but affects the quality negatively at large values (Model 2 in Table 4). The inflection point beyond which the relatedness of failures begins to impact R&D performance negatively corresponds to the values of 10 and 9 for R&D productivity and R&D outcome quality, respectively. Most of the failures in the sample have relatedness value of less than 8, except in some instances when the value is higher than 8 and can be as high as 35. Even though curvilinear effect is not relevant for most firms, some firms in the sample did experience negative effect of relatedness of failures on R&D performance. Interpretations of coefficients in Model 2 in Table 3 revealed that for a unit increase in relatedness, on average, firms experienced 6% increase in R&D productivity. Similar analysis with respect to Model 2 in Table 4 led to the

finding that a unit increase in relatedness corresponds to 8% increase in R&D outcome quality. Beyond the inflection points, a unit increase in relatedness is associated with 0.3% and 0.4% decrease in R&D productivity and R&D outcome quality, respectively.

Hypothesis 2 predicts a curvilinear relationship between relatedness among failures and R&D performance and is also supported as evident in Model 3 of both Tables 3 and 4. Both R&D productivity and quality are positively associated with relatedness among failures, but only to a certain point, after which an increase in relatedness among failures leads to decrease in R&D productivity and quality. The inflection point for both productivity and quality measure of R&D performance is at the value of 5, which is within the range of data used to test hypotheses in this study, but is above most of the values of relatedness among failures in the sample. This is similar in nature to the effect of relatedness of failures examined in Hypothesis 1, where few firms experienced negative impact. As before, analysis of coefficients in Model 3 of Table 3 suggests that a unit increase in relatedness among failures by one unit leads to an increase in R&D productivity by 14%. However, beyond inflection point, increase of one unit in relatedness among failures decreases the R&D productivity by 0.8%. Also, results are similar in nature for R&D outcome quality with relatedness among failures associated with 17% increase before inflection point and 1% decrease afterwards (Model 3 in Table 4).

Hypothesis 3 argues that a firm's experience with failures in more familiar knowledge elements will provide less learning opportunities and therefore will result in negative R&D outcomes, as compared to the firm's experience with less familiar elements. In the present context, experience with familiar elements implies the extent to which a firm has used these elements in the past. Failure in familiar elements indicates a firm's inability to make use of

prior knowledge about these elements. However, failure in experiments with less familiar elements is expected and provides the firm with new information about elements and their combinations. Model 4 in Tables 3 and 4 tests the effect of familiarity of elements on R&D productivity and R&D outcome quality, respectively. As evident from Model 4, the familiarity of elements has negative and significant ($p\text{-value} = 0.000$) effect on both R&D productivity and R&D quality, which supports Hypothesis 3. In other words, as a firm experiences failures in experiments that use more familiar knowledge elements, its subsequent R&D performance decreases. For the sample of firms in this study, for a unit increase in familiarity of elements, R&D productivity of firms decreased by approximately 5%. Similarly, firms experienced 2% decrease in R&D quality for a unit increase in familiarity of knowledge elements. Overall, results suggest that firms are able to learn more and therefore improve their R&D performance when they experience failures in less familiar elements as compared to more familiar ones.

In hypothesis 4, I discuss familiarity of knowledge used in failures experienced by firms. Specifically, I explore whether there are any implications of familiarity of knowledge used in failures for firms' R&D performance. Model 5 in Tables 3 and 4 reports results for above relationship for R&D productivity and quality, respectively. The sign of the coefficient for familiarity of knowledge is negative and significant, suggesting that firms benefit more from experience with failures that use less familiar knowledge. Simple calculations show that for a unit increase in familiarity of the knowledge used in failures, firms' R&D productivity decreased by 0.14% and R&D outcome quality decreased by 0.17%.

Hypotheses 5a and 5b predict that firms learn the most when decomposability of their knowledge base is moderate, i.e. neither too low nor too high. In other words,

decomposability of knowledge base moderates the relationships tested in both Hypotheses 1 and 2 such that, at moderate level of decomposability, positive slope of inverted-U relationship will be steeper (5a), and negative slope will be flatter (5b), as compared to when decomposability is low or high. Model 6 in Table 3 tests this relationship with dependent variable as R&D productivity and independent variable as *relatedness of failures*, same as the one used in Hypothesis 1. I created three dummies, each for a level of decomposability, with ‘low’ for values of decomposability below 0.4, ‘moderate’ for values between 0.4 and 0.6, and ‘high’ when values are above 0.6. In model 6 of Table 3, coefficient of interaction term containing linear term of relatedness of failures and moderate decomposability is higher than coefficient of interaction term that contains linear term of relatedness of failures and low decomposability, suggesting that positive slope of relationship between relatedness of failures and R&D productivity is higher when decomposability is moderate as compared to when decomposability is low. Same is true for the high level of decomposability. I plotted the above relationship at all three levels of decomposability to show whether firms learn the most when decomposability is moderate.

As shown in Figure 6, positive slope of the relationship between relatedness of failures and R&D productivity is indeed the highest at moderate level of decomposability, supporting hypothesis 5a. However, negative slope of above relationship is not flatter for moderate level of decomposability as compared to low and high levels, in contrast to hypothesis 5b. Overall, the results support Hypothesis 5a but not 5b, i.e. positive impact of relatedness of failures on R&D productivity is highest when decomposability is moderate, as hypothesized in 5a, but negative impact is not the lowest, in fact it is the highest, when decomposability is moderate, in contrast to hypothesis 5b. Similar results hold when

dependent variable is R&D quality and independent variable is *relatedness of failures*. Also, results are of same nature when independent variable is *relatedness among failures*, and dependent variables are R&D productivity and R&D quality. I will discuss the implications of these results in the section on discussion.

6. DISCUSSION

The dissertation broadly considers three aspects of learning from small failures in experimentation. First, I explore the link between relational characteristics of small failures in experimentation and R&D performance of a firm. Second, implications of search behavior associated with failures for firms' R&D performance are studied. Finally, I examine how structure of a firm's knowledge base moderates the discussed relationships.

With the help of data on patents and their expirations for firms in the pharmaceutical industry, I show that a firm learns most from failures that are result of experimentation with knowledge elements that are highly related to other knowledge elements in the firm's knowledge base. Failures that rank low on relatedness are not connected to other projects in a firm's knowledge base. Such failures provide information to firms that is relatively easier to interpret but limited in nature with respect to learning and improving subsequent R&D activities. However, too much relatedness of knowledge elements with other elements can lead to causal ambiguity, which prevents firms from discerning causal structure of relationships and identifying combination of elements liable for failures. In some cases, too much relatedness is also a signal of experimentation with inconsequential knowledge elements, and a firm's unwillingness to experiment with new knowledge elements, limiting the scope of learning. For example, if a failure combines elements $e1$, $e2$, and $e3$, and these elements are not connected to any other element in the knowledge base, the firm does not benefit from this information in terms of improving ongoing projects. However, if all or some of the elements $e1$, $e2$, and $e3$ are connected to other elements, firms can use the

information from failure to assess the potential of projects that have elements in common with the failure.

These arguments are supported by the empirical analysis that shows that as the relatedness of failures increases, the subsequent R&D performance also increases because firms are able to evaluate the potential of ongoing projects by analyzing the information from failures. On the other hand, high relatedness may not always result in information that can be analyzed and utilized by firms. One of the factors that can limit the learning from a failure is the complexity of the information generated by failures. Beyond a certain point, increase in relatedness of failures will lead to causal ambiguity, which can obstruct a firm's ability to interpret the information from failures effectively. In present settings, failures are ranked high on relatedness when their elements are present in most of the other ongoing projects. As the relatedness of failures increases above certain level, firms are not able to identify elements or their combinations responsible for failures on account of the increased complexity.

Figure 2b illustrates the difficulty a firm can face due to failures of higher relatedness. As shown in the figure, red nodes, which represent elements underlying failures, are present across almost all existing projects (indicated by the presence of linkages between red and blue nodes) in Rhône-Poulenc, therefore making it challenging for the firm to evaluate these projects. As most of the ongoing projects consist of elements that are part of failures, it is extremely difficult for a firm to discriminate between these projects and allocate its attention across projects accordingly. This is the reason we observe a curvilinear relationship between relatedness of failures and R&D performance. Initially, as the relatedness goes up, firms learn about potential elements and their combinations and can use this information to

improve existing projects, but too high relatedness makes the process of identifying unproductive elements and their combinations considerably challenging, hurting the R&D performance consequently.

Similarly, as argued before and evident in empirical analysis, *relatedness among failures* is positively associated with R&D performance till a certain point, after which it begins to affect the performance negatively. Firms can either experiment with elements in isolation or combine them to varying degrees. By combining elements in different ways, a firm learns about outcome of these combinations and can use this learning to evaluate ongoing projects. However, after a point, increase in relatedness among failures decreases the likelihood of determining potential elements and their combinations. Figure 3b provides an example of failures that are high on relatedness, i.e. elements are tightly connected to each other. The higher density makes the process of identifying elements and combinations that are liable for failures complex and challenging. This is why I find a curvilinear relationship between R&D performance and relatedness among failures.

In addition to examining nature of relationships among failures, the dissertation also explores search behavior associated with failures. Specifically, I argue that failures in knowledge elements that are familiar to firms will lead to less learning opportunities as compared to failures in elements new to the firm. Higher extent of failure in familiar elements compared to new elements signals a firm's inability to leverage previous knowledge about these elements. Experimentation with new elements, however, is likely to reveal new information that can contribute to a firm's inventive capabilities. Empirical analysis supports these arguments and shows that firms increase their R&D performance more following failures that are a result of experimentation with new elements rather than familiar elements.

Another aspect of search behavior that is explored in this dissertation is the type of knowledge that is used in failed projects. Firms can rely on the knowledge local to them, i.e. previously used or experiment with knowledge previously unknown to them. Knowledge elements discussed in previous section are akin to ingredients in a dish, whereas, type of knowledge discussed in this section is the dish's recipe. A firm can use similar elements but can rely on different pieces of knowledge and vice versa. Experimentation with local knowledge is likely to lead to successful outcomes because of firms' experience with it. Experimentation with distant knowledge, however, is often uncertain and likely to result in unsuccessful outcomes. Even though firms are better off relying on local knowledge because of higher rate of success, previous research has emphasized the need to explore distant knowledge for long-term viability (Gupta, Smith, & Shalley, 2006; Katila & Ahuja, 2002; March, 1991).

I argued earlier that failures associated with local knowledge are suggestive of a firm's inability to leverage prior experience. This can also be an indication of obsolescence of the firm's current knowledge. Failures in experiments that use distant knowledge are expected to occur and provide firms with valuable feedback necessary for novel inventions. Therefore, failure in experiments that use distant knowledge is likely to generate more learning opportunities for firms as compared to local knowledge. I find support for the above arguments in empirical analysis, and show that failure experience associated with distant knowledge tends to increase firms' subsequent R&D performance.

Results related to the moderating role of decomposability in relationship between relatedness of failures and R&D performance are partially supported. I find that positive impact of *relatedness of failures* on R&D performance is higher when decomposability is

moderate as a firm is able to make necessary changes within modules without causing unnecessary disturbance in other modules, and at the same time benefits of such changes are spread across the entire knowledge base because of the presence of cross-module linkages, something not possible when decomposability is low or high. However, negative impact of *relatedness of failures* beyond a certain point is also highest when decomposability is moderate, contrary in contrast to the hypothesized relationship. The opposite results discussed here could be because of the superstitious learning by firms (Argyris, 1986, Zollo, 2009). At extreme values of relatedness of failures, a firm is presented with information on a wide array of knowledge elements in its knowledge base. Such information, though seemingly beneficial, can lead a firm to interpret the relationships among elements incorrectly.

When a firm experiences a failure comprising knowledge elements (let us say x and y) that are widely used in the firm's other inventions as well, the process of identifying unfruitful combinations of elements x and y with other elements can be difficult because of the presence of large number of such combinations. A firm can incorrectly identify combinations responsible for failures, leading to superstitious learning. Changes introduced based on the superstitious learning will hurt the firm most when decomposability of its knowledge base is moderate because of the spread of such learning across the entire knowledge base; results are consistent with those found here. When decomposability is high, incorrect changes will be confined to the respective modules and will not hurt a firm as much. In case of low decomposability, a firm will find it difficult to introduce any changes (Simon, 1962, 1991), thereby limiting the negative impact of superstitious learning.

7. CONTRIBUTIONS AND LIMITATIONS

The dissertation makes a contribution to the present literature in several important ways. First, it contributes to the literature on learning from failures that, for the most part, has overlooked the small failures that occur during experimentation. Scholars have explored large or catastrophic failures and their implications for organizational learning (Haunschild & Sullivan, 2002; Madsen & Desai, 2010). Researchers have also examined operational failures, and though, in some instances these failures can be categorized as small, they often result from mistakes or are errors on a firm's part, and often reduce over time as the firm learns from them (Hayward, 2002; Haunschild & Rhee, 2004; Henderson & Stern, 2004). However, small failures studied in this dissertation differ fundamentally from catastrophic failures and operational failures studied earlier. Failures that have been studied in literature so far are not expected to occur and decrease over time as firms learn to avoid them. The key learning from such failures is to use the feedback from them to avoid similar failures in future as much as possible.

Small failures in experimentation, however, differ in terms of the way firms learn from them. Small failures that occur during experimentation are not desirable or welcome by a firm, but they certainly are expected. Even a small set of knowledge elements can result in extraordinarily large number of combinations, and the only way firms can identify true combinations is through experimentation. Most of these experiments fail, and firms are aware of it. Given such nature of the R&D process, the aim is not to minimize failed experiments, but use them as effectively as possible to determine the knowledge elements and their

combinations that can lead to successful inventions. By focusing on small failures in experimentation, present study not only adds another dimension to existing learning from failures literature, but also helps to increase our understanding of innovation process.

Second, the current study makes another important contribution to learning from failure literature by exploring the heterogeneity in failures with respect to their relational characteristics in a firm's knowledge base, and how this heterogeneity can have an effect on learning outcomes for the firm. I bring together research on relatedness, complexity, and learning from failures to answer the above question.

Most studies, if not all, do not discriminate among failures when it comes to learning from them (Hayward, 2002; Henderson & Stern, 2004), but not all failures are same and can differ on multiple dimensions. In this dissertation, I focus on relatedness of failures in a firm's knowledge base and its implications for learning. Small failures can differ in terms of relationship of elements underlying these failures with other elements in a firm's knowledge base. Previous research has examined implications of such relationships for knowledge flows and employee mobility (Sorenson, Rivkin, & Fleming, 2006; Ganco, 2013). Whether failures of varying relatedness result in heterogeneous learning outcomes is not explored in the literature. By examining the relatedness of small failures and its effect on learning, this study further increases our understanding of mechanisms involved in learning process.

Third, the dissertation makes an important contribution to the research on knowledge structures by examining their role in learning. Research on factors that can influence a firm's ability to incorporate learning from failures is rather limited. In most cases, it has been assumed that ability to assimilate the learning from failures is homogenous across firms. However, it may not be the case. Many factors can play a role in effective learning from

failures, and in this study I focus on one such factor, i.e. structure of a firm's knowledge base. To incorporate learning from failures, a firm is often required to make changes in its knowledge base. These changes may include bringing in new knowledge elements or changing combination of existing elements. It has been shown that there is heterogeneity across knowledge structures with respect to their capacity to change, and some structures are able to go through changes more easily than the other structures (Lyles & Schwenk, 1992; Yayavaram & Ahuja, 2008). Scholars have explored the role of decomposability in acquiring structural properties such as recombination, adaptation, and persistence (Baldwin & Clark, 2000; Schilling, 2000; Weick, 1976). Recent research has also begun to explore the effect of decomposability of a firm's knowledge base on the usefulness of subsequent inventions (Yayavaram & Ahuja, 2008). However, research has not yet explored whether the structure of knowledge base plays a role in learning from small failures in experimentation. Overall, current study is distinct from other studies on learning from failures, because it takes into account heterogeneity in learning ability across firms, something largely ignored in previous studies.

Fourth, this study contributes to the literature on organizational innovation by investigating failed innovation attempts and their influence on R&D performance. Research has explored various factors that affect R&D outcome (e.g., Henderson & Cockburn, 1994; Rothaermel & Thursby, 2007; Griliches, 1994; Ahuja, 2000), but failures in innovation process have not been studied from the perspective of learning. Though, there are studies that discuss the importance of experiment in innovation process, they are either largely qualitative in nature or do not discriminate between success and failure experience in experimentation (Sitkin, 1992). This study goes one step further and explores how failures in experimentation

can influence learning from them. By studying how different failures, depending on their position in a firm's knowledge base, can result in heterogeneous R&D performance, this study helps understand the process of innovation better.

Overall, research in strategic management has long focused on the heterogeneity in performance across firms (Helfat, 2000). One of the ways firms, especially in high technology industry, are able to gain competitive advantage, is through innovation (Schumpeter, 1934). Literature consists of studies that demonstrate the relationship between innovation and performance (e.g., Lawless & Anderson, 1996; Christensen, 1997). Similarly, studies have also provided important insights related to decision making and innovation within firms (Vandeven, 1986; Dougherty, 1992). Most of the existing research that seeks to explain the link between heterogeneity in firm performance and innovation focuses on the output of innovation, i.e. new products and new processes (Griliches, 1994; Hall, 1993; Hall & Mairesse, 1995; Hill & Rothaermel, 2003; Lee, Lee, & Pennings, 2001; Li & Atuahene-Gima, 2001; Mansfield, 1965; Nelson, 1962; Pavitt, 1993; Schmookler, 1966; Schumpeter, 1934, 1982; Vandeven, 1986; Vandeven & Polley, 1992). However, very little research has been done to explore factors such as problem solving and experimentation, which are inputs to the process of innovation. By studying factors responsible for innovation within firms, current study increases our understanding of heterogeneity in innovation capabilities across firms, and throws more light on the link between innovation and firm performance.

The current study is not without limitations. First, its findings are limited to the pharmaceutical industry. Will the results hold in other industries in which patents do not represent experimentation but rather products themselves? What happens when small failures are not perceived as a result of experimentation? In that case, are individuals reluctant in

acknowledging failures, and hence do not learn as much? Future research aimed at exploring different industries can improve our understanding of how the process of learning from failures may differ in different settings. Second, I capture only failures that are patented by a firm. Patents do not entirely represent a firm's R&D efforts, and it is also possible that some ideas fail even before reaching the patenting stage. However, given the high rate of patenting in the pharmaceutical industry (Arundel & Kabla, 1998), I expect these ideas to be relatively small in number, and therefore with little to no effect on overall findings in this dissertation. Finally, I use patent productivity and quality as measures of R&D performance. However, other measures such as commercial value of innovations are also important indicators of innovative performance. Future research that explores the value of learning from small failures beyond the patent portfolio, and whether such learning translates into more product launches or commercial success could enhance our understanding of the effect of failures at both scientific and commercial levels. Such studies could also provide a more complete picture of the subject matter.

FIGURES AND TABLES

TABLE 1

List of 76 Pharmaceutical Firms Used in This Dissertation

3M	Farmitalia Carlo Erba, S.P.A.	Roussel-Uclaf
Abbott Laboratories	Fujisawa Pharmaceutical Co., Ltd.	Sandoz Pharmaceuticals Corp.
Adir Et Compagnie	G. D. Searle & Co.	Sankyo Company Limited
Alcon Laboratories, Inc.	Genentech, Inc.	Sanofi, S.A.
Allergan, Inc.	Glaxo Group Limited	Schering Corp.
Alza Corporation	Hoechst-Roussel Pharmaceuticals, Inc.	Smith Kline & French Laboratories Ltd.
American Cyanamid Company	Hoffmann-La Roche Inc.	Smithkline Beckman Corporation
American Home Products Corporation	Imperial Chemical Industries Plc	Stauffer Chemical Company
Amgen Inc.	International Flavors + Fragrances Inc	Sterling Drug Inc.
Astrazeneca Ab	Janssen Pharmaceutica N.V.	Sumitomo Chemical Company, Limited
Basf Aktiengesellschaft	John Wyeth & Brother, Ltd.	Syntex (U.S.A.) Inc.
Bayer Aktiengesellschaft	Johnson and Johnson	Schering Corporation
Bausch and Lomb	Kao Kabushiki Kaisha (Kao Corporation)	Smithkline Beecham Corporation
Beecham Group P.L.C.	L'Oreal	Takeda Chemical Industries Ltd.
Boehringer Ingelheim, Gmbh	Massachusetts Institute Of Technology	The Procter & Gamble Company
Boehringer Mannheim G.M.B.H.	Merrell Dow Pharmaceuticals	The Regents of the University Of California
Burroughs Wellcome Co.	Merck & Co., Inc.	The United States of America As Representatives
Bristol-Myers Squibb Company	Novo Nordisk A/S	United States of America, Department of
Chesebrough-Pond's USA Co., Division of	Novartis Ag	United States of America, Health & Human
Chiron Corporation	Ortho Pharmaceutical Corporation	University of Texas
Ciba-Geigy Corporation	Otsuka Pharmaceutical Co., Ltd.	Upjohn Company
Colgate-Palmolive Company	Procter + Gamble Company	Warner-Lambert Company
Dow Chemical Company	Pfizer Inc.	Wisconsin Alumni Research Foundation
E. I. Du Pont De Nemours And Company	Richardson-Vicks Inc.	Wyeth
Eisai Co., Ltd.	Rohm And Haas Company	Zeneca Limited
Eli Lilly And Company		

TABLE 2
Descriptive Statistics and Correlation Matrix

Variable		<i>n</i>	Mean	s. d.	Min	Max	1	2	3	4	5	6	7	8	9	10	11	12
1.	R&D productivity	1731	185	189	1	1473	1.00											
2.	R&D quality	1731	139	160	0	1929	0.74	1.00										
3.	Relatedness of failures	1200	2.8	3	0	35	0.32	0.31	1.00									
4.	Relatedness among failures	1200	1.1	1.7	0	20	0.32	0.41	-0.03	1.00								
5.	Familiarity of knowledge elements	1200	4.5	6.5	0	79	0.34	0.23	-0.05	0.20	1.00							
6.	Familiarity of knowledge	1200	0.32	0	1	442	0.52	0.38	0.18	0.20	0.57	1.00						
7.	Decomposability	1731	0.5	0.16	0	0.87	-0.50	-0.28	-0.19	-0.14	-0.44	-0.38	1.00					
8.	Number of failures	1200	121	177	1	1481	0.37	0.15	0.18	0.07	0.21	0.26	-0.55	1.00				
9.	Number of alliances	1400	2.8	6.5	0	52	0.28	0.11	0.03	0.04	0.07	0.15	-0.33	0.42	1.00			
10.	Number of scientists	1731	58	61	1	407	0.82	0.54	0.24	0.23	0.19	0.37	-0.46	0.50	0.44	1.00		
11.	Technological diversity	1731	0.03	0.07	0	1	-0.27	-0.20	0.05	-0.11	-0.14	-0.13	0.10	-0.08	-0.05	-0.27	1.00	
12.	Number of countries	1731	127	107	1	625	0.65	0.68	0.17	0.33	0.38	0.50	-0.53	0.39	0.31	0.83	-0.33	1.00

TABLE 3
Fixed Effects Negative Binomial Estimates for R&D performance (DV = R&D Productivity $_{it+1}$)

Variable	1	2	3	4	Variables Roles
R&D productivity $_{it}$	0.0005 *** (0.000)	0.0006 *** (0.000)	0.0004 *** (0.000)	0.0004 *** (0.000)	Lagged DV
Relatedness of failures $_{it}$		0.10 *** (0.01)			H1
Relatedness of failures squared $_{it}$		-0.005 *** (0.001)			H1
Relatedness among failures $_{it}$			0.11 *** (0.01)		H2
Relatedness among failures squared $_{it}$			-0.007 *** (0.001)		H2
Familiarity of knowledge elements $_{it}$				-0.15 *** (0.000)	H3
Familiarity of knowledge $_{it}$					H4
Relatedness of failures $_{it} \times$ Low decomposability $_{it}$					H5a
Relatedness of failures squared $_{it} \times$ Low decomposability $_{it}$					H5a
Relatedness of failures $_{it} \times$ Moderate decomposability $_{it}$					H5b
Relatedness of failures squared $_{it} \times$ Moderate decomposability $_{it}$					H5b
Relatedness among failures $_{it} \times$ Low decomposability $_{it}$					H5a
Relatedness among failures squared $_{it} \times$ Low decomposability $_{it}$					H5a
Relatedness among failures $_{it} \times$ Moderate decomposability $_{it}$					H5b
Relatedness among failures squared $_{it} \times$ Moderate decomposability $_{it}$					H5b
Low decomposability $_{it}$	-0.46 * (0.18)	-0.71 ** (0.23)	-0.51 ** (0.17)	-0.08 (0.06)	Control
Moderate decomposability $_{it}$	0.30 * (0.13)	0.28 * (0.13)	0.28 * (0.12)	0.09 (0.05)	Control
Number of failures $_{it}$	-0.0004 *** (0.000)	-0.0005 *** (0.000)	-0.0003 *** (0.000)	-0.0005 *** (0.000)	Control

Number of alliances $_{it}$	-0.007 *** (0.002)	-0.01 *** (0.002)	-0.006 *** (0.002)	-0.008 *** (0.002)	Control
Number of scientists $_{it}$	0.004 *** (0.000)	0.006 *** (0.000)	0.004 *** (0.000)	0.004 *** (0.000)	Control
Technological diversity $_{it}$	-24.96 *** (1.69)	-24.18 *** (1.71)	-24.75 *** (1.67)	-27.8 *** (1.9)	Control
Number of countries $_{it}$	-0.001 (0.007)	-0.003 (0.007)	-0.002 (0.007)	-0.002 (0.007)	Control
Constant	2.10 *** (0.12)	1.97 *** (0.13)	2.04 *** (0.12)	2.31 *** (0.09)	
$Wald \chi^2$	1270	1337	1518	1278	
χ_p^2	0.0000	0.0000	0.0000	0.0000	
# Observations	1130	1130	1130	1130	
# Groups	75	75	75	75	

Variable	5	6	7	8	Variables Roles
R&D productivity $_{it}$	0.0006 *** (0.000)	0.0005 *** (0.000)	0.0005 *** (0.000)	0.0004 *** (0.000)	Lagged DV
Relatedness of failures $_{it}$		0.06 *** (0.01)		0.06 * (0.03)	H1
Relatedness of failures squared $_{it}$		-0.003 *** (0.000)		-0.003 ** (0.001)	H1
Relatedness among failures $_{it}$			0.14 *** (0.000)	0.13 ** (0.05)	H2
Relatedness among failures squared $_{it}$			-0.01 *** (0.002)	-0.008 *** (0.001)	H2
Familiarity of knowledge elements $_{it}$				-0.12 *** (0.024)	H3
Familiarity of knowledge $_{it}$	-0.0014 *** (0.000)			0.0004 ** (0.0002)	H4
Relatedness of failures $_{it} \times$ Low decomposability $_{it}$		0.29 *** (0.06)		0.27 *** (0.05)	H5a
Relatedness of failures squared $_{it} \times$ Low decomposability $_{it}$		-0.024 *** (0.007)		-0.021 ** (0.006)	H5a
Relatedness of failures $_{it} \times$ Moderate decomposability $_{it}$		0.39 ** (0.16)		0.33 * (0.18)	H5b
Relatedness of failures squared $_{it} \times$ Moderate decomposability $_{it}$		-0.038 *** (0.009)		-0.023 ** (0.01)	H5b
Relatedness among failures $_{it} \times$ Low decomposability $_{it}$			0.13 *** (0.04)	0.09 * (0.042)	H5a
Relatedness among failures squared $_{it} \times$ Low decomposability $_{it}$			-0.007 **	-0.006 *	H5a

Low decomposability $_{it}$			(0.003)		(0.003)	
Relatedness among failures $_{it} \times$			0.28 ***		0.22 *	H5b
Moderate decomposability $_{it}$			(.06)		(.011)	
Relatedness among failures squared $_{it} \times$			-0.018 **		-0.011 *	H5b
Moderate decomposability $_{it}$			(0.09)		(0.09)	
Low decomposability $_{it}$	0.02	-0.34	-0.34		-0.31	Control
	(0.06)	(0.24)	(0.25)		(0.39)	
Moderate decomposability $_{it}$	0.14 **	0.12	0.14		0.10	Control
	(0.05)	(0.09)	(0.11)		(0.10)	
Number of failures $_{it}$	-0.0004 ***	-0.0004 ***	-0.0004 ***		-0.0004 ***	Control
	(0.000)	(0.000)	(0.000)		(0.000)	
Number of alliances $_{it}$	-0.007 ***	-0.009 ***	-0.009 ***		-0.01 ***	Control
	(0.002)	(0.002)	(0.002)		(0.002)	
Number of scientists $_{it}$	0.004 ***	0.006 ***	0.005 ***		0.006 ***	Control
	(0.000)	(0.000)	(0.000)		(0.000)	
Technological diversity $_{it}$	-24.9 ***	-25.18 ***	-25.15 ***		-21.68 ***	Control
	(1.67)	(1.73)	(1.70)		(1.24)	
Number of countries $_{it}$	-0.002	-0.002	-0.002		-0.004	Control
	(0.008)	(0.007)	(0.007)		(0.007)	
Constant	2.28 ***	2.22 ***	2.04 ***		2.89 ***	
	(0.09)	(0.13)	(0.12)		(0.13)	
Wald χ^2	1358	1477	1652		2600	
χ_p^2	0.0000	0.0000	0.0000		0.0000	
# Observations	1130	1130	1130		1130	
# Groups	75	75	75		75	

Note: Standard errors in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001

TABLE 4

Fixed Effects Negative Binomial Estimates for R&D Performance (DV = R&D Quality $it + 1$)

Variable	1	2	3	4	Variables Roles
R&D productivity it	0.002 *** (0.000)	0.002 *** (0.000)	0.002 *** (0.000)	0.002 *** (0.000)	Lagged DV
Relatedness of failures it		0.072 *** (0.02)			H1
Relatedness of failures squared it		-0.004 * (0.002)			H1
Relatedness among failures it			0.11 *** (0.03)		H2
Relatedness among failures squared it			-0.01 ** (0.003)		H2
Familiarity of knowledge elements it				-0.04 * (0.03)	H3
Familiarity of knowledge it					H4
Relatedness of failures $it \times$ Low decomposability it					H5a
Relatedness of failures squared $it \times$ Low decomposability it					H5a
Relatedness of failures $it \times$ Moderate decomposability it					H5b
Relatedness of failures squared $it \times$ Moderate decomposability it					H5b
Relatedness among failures $it \times$ Low decomposability it					H5a
Relatedness among failures squared $it \times$ Low decomposability it					H5a
Relatedness among failures $it \times$ Moderate decomposability it					H5b
Relatedness among failures squared $it \times$ Moderate decomposability it					H5b
Low decomposability it	-1.35 *** (0.28)	-1.52 *** (0.28)	-1.48 *** (0.28)	-1.25 *** (0.26)	Control
Moderate decomposability it	0.57 *** (0.13)	0.53 *** (0.13)	0.45 *** (0.09)	0.36 *** (0.11)	Control
Number of failures it	0.001 *** (0.000)	0.001 *** (0.000)	0.001 *** (0.000)	0.001 *** (0.000)	Control

Number of alliances $_{it}$	-0.013 *** (0.003)	-0.013 *** (0.003)	-0.012 *** (0.003)	-0.008 ** (0.003)	Control
Number of scientists $_{it}$	0.002 ** (0.000)	0.002 ** (0.000)	0.002 ** (0.000)	-0.002 *** (0.000)	Control
Technological diversity $_{it}$	-13.5 *** (1.56)	-11.95 *** (1.56)	-12.55 *** (1.55)	-8.1 *** (1.5)	Control
Number of countries $_{it}$	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.02 (0.02)	Control
Constant	0.44 ** (0.16)	0.13 (0.18)	0.26 (0.17)	0.56 *** (0.17)	
$Wald \chi^2$	654	687	692	684	
χ_p^2	0.0000	0.0000	0.0000	0.0000	
# Observations	1130	1130	1130	1130	
# Groups	75	75	75	75	

Variable	5	6	7	8	Variables Roles
R&D productivity $_{it}$	0.002 *** (0.000)	0.002 *** (0.000)	0.002 *** (0.000)	0.002 *** (0.000)	Lagged DV
Relatedness of failures $_{it}$		0.11 *** (0.03)		0.08 ** (0.03)	H1
Relatedness of failures squared $_{it}$		-0.005 ** (0.002)		-0.004 * (0.003)	H1
Relatedness among failures $_{it}$			0.13 *** (0.04)	0.13 ** (0.07)	H2
Relatedness among failures squared $_{it}$			-0.01 ** (0.004)	-0.01 ** (0.004)	H2
Familiarity of knowledge elements $_{it}$				-0.02 * (0.01)	H3
Familiarity of knowledge $_{it}$	-0.0017 *** (0.000)			0.002 *** (0.000)	H4
Relatedness of failures $_{it} \times$ Low decomposability $_{it}$		0.16 + (0.10)		0.09 + (0.06)	H5a
Relatedness of failures squared $_{it} \times$ Low decomposability $_{it}$		-0.03 * (0.01)		-0.02 * (0.01)	H5a
Relatedness of failures $_{it} \times$ Moderate decomposability $_{it}$		0.29 * (0.14)		0.18 * (0.13)	H5b
Relatedness of failures squared $_{it} \times$ Moderate decomposability $_{it}$		-0.1 ** (0.04)		-0.05 * (0.04)	H5b
Relatedness among failures $_{it} \times$ Low decomposability $_{it}$			0.17 ** (0.05)	0.14 ** (0.05)	H5a
Relatedness among failures squared $_{it} \times$			-0.01 * (0.05)	-0.01 * (0.05)	H5a

Low decomposability $_{it}$			(0.004)		(0.005)	
Relatedness among failures $_{it} \times$			0.32 **		0.26 *	H5b
Moderate decomposability $_{it}$			(0.14)		(0.14)	
Relatedness among failures squared $_{it} \times$			-0.015 *		-0.011 *	H5b
Moderate decomposability $_{it}$			(0.008)		(0.007)	
Low decomposability $_{it}$	-1.22 *** (0.34)	-1.45 *** (0.31)	-1.05 *** (0.32)		-1.57 *** (0.50)	Control
Moderate decomposability $_{it}$	0.49 *** (0.16)	0.62 ** (0.25)	0.58 ** (0.25)		0.43 ** (0.24)	Control
Number of failures $_{it}$	0.001 *** (0.000)	0.001 *** (0.000)	0.001 *** (0.000)		0.001 *** (0.000)	Control
Number of alliances $_{it}$	-0.008 * (0.003)	-0.018 *** (0.004)	-0.017 *** (0.004)		-0.012 *** (0.004)	Control
Number of scientists $_{it}$	-0.002 *** (0.000)	0.003 ** (0.000)	0.003 ** (0.000)		0.002 ** (0.000)	Control
Technological diversity $_{it}$	-10.03 *** (1.4)	-11.72 *** (1.56)	-12.10 *** (1.58)		-10.25 *** (1.56)	Control
Number of countries $_{it}$	-0.03 (0.02)	-0.01 (0.01)	-0.01 (0.01)		-0.01 (0.01)	Control
Constant	1.13 *** (0.10)	0.16 (0.19)	0.27 (0.17)		0.15 (0.26)	
$Wald \chi^2$	716	696	703		845	
χ_p^2	0.0000	0.0000	0.0000		0.0000	
# Observations	1130	1130	1130		1130	
# Groups	75	75	75		75	

Note: Standard errors in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001

FIGURE 1
Number of Expired Patents for Firms Included in This Study between 1985 and 2002





FIGURE 2

Network of Subclasses with Red Nodes Representing Elements Underlying Failed Patents and All other Elements Denoted by Blue Nodes

a. Fisons Pharmaceuticals (1995)



b. Rhône-Poulenc (1995)





FIGURE 3

Network of Subclasses with Nodes Representing Elements Underlying Failed Patents

a. Bristol-Myers Squibb (1991)



b. Bristol-Myers Squibb (1995)





FIGURE 4

Knowledge Base of Firms with Nodes of Same Colors Belonging to the Same Module

a. Rhône-Poulenc (1991)



b. Fisons Pharmaceuticals (1991)




FIGURE 5

Sample Data Points for Pfizer in Year 2000 and Description of Variables

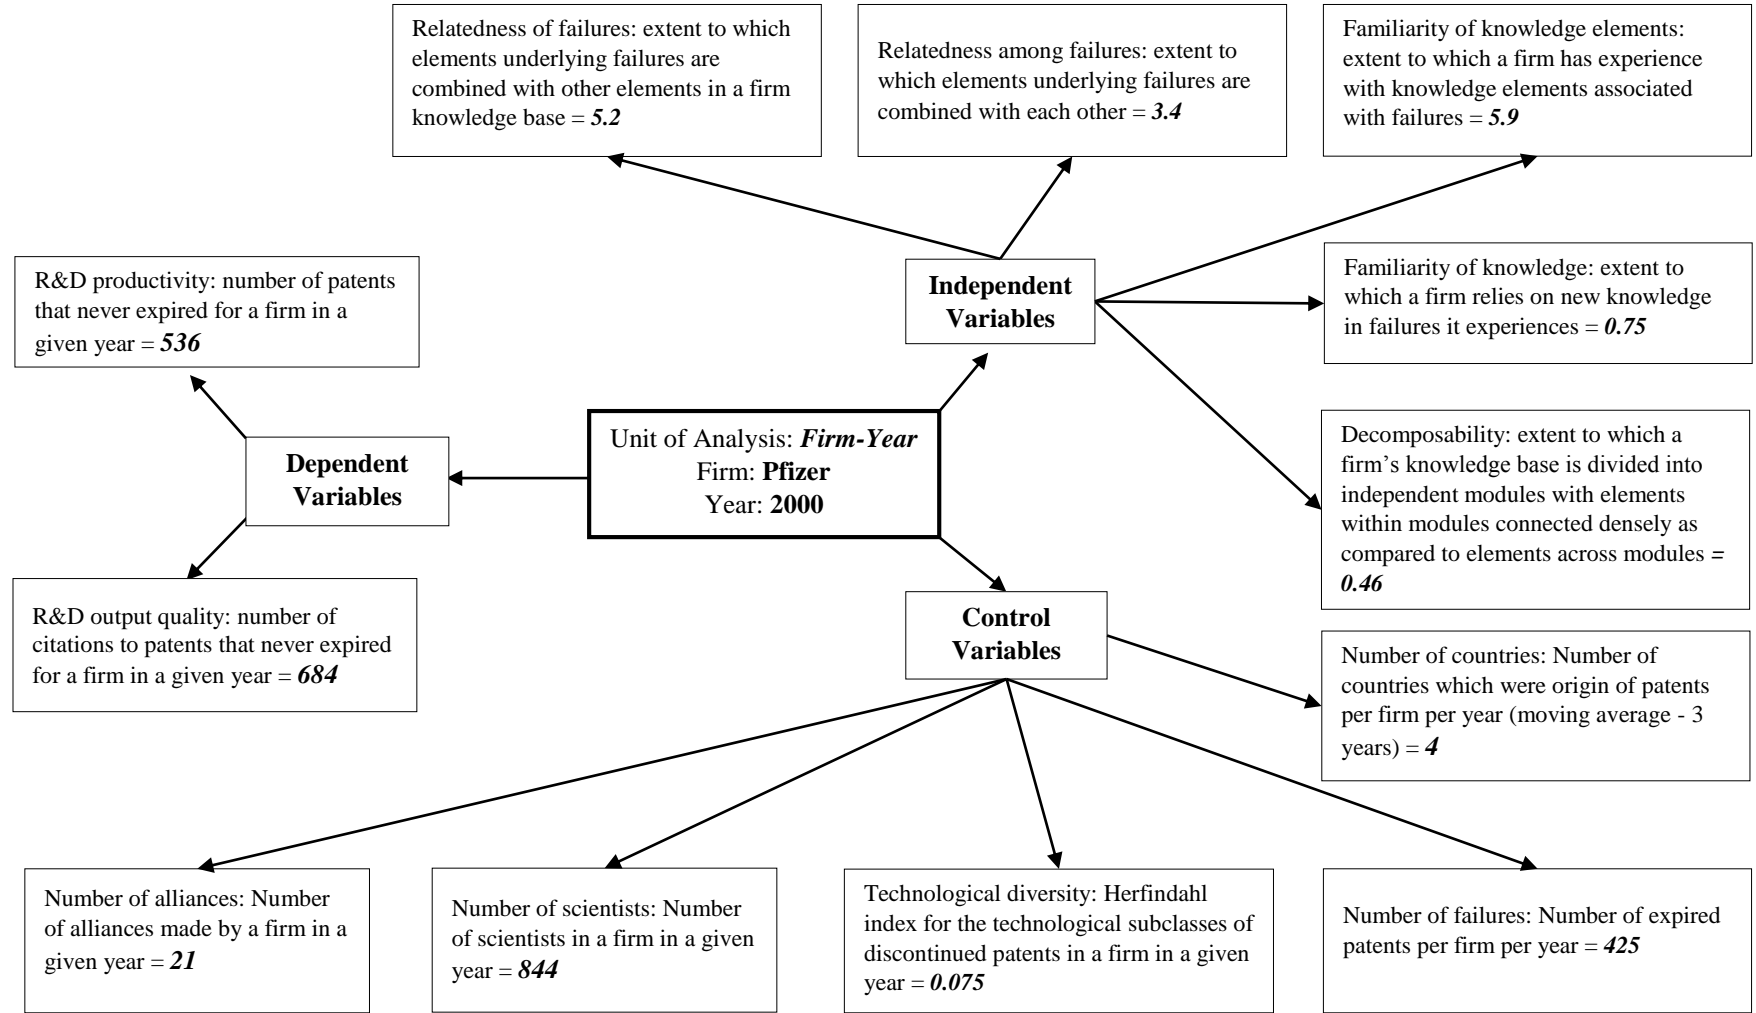
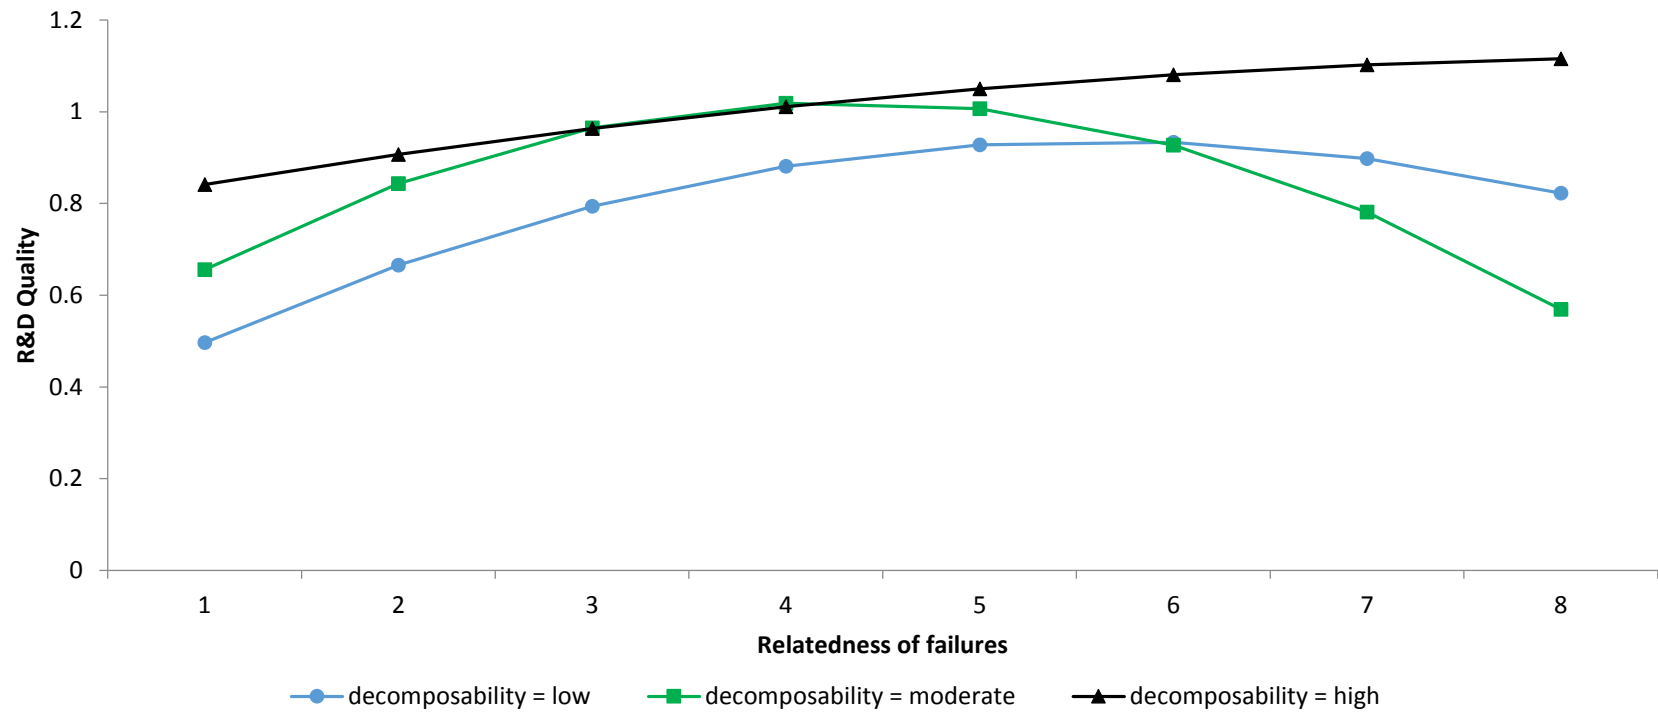


FIGURE 6

Nature of Relationship between Relatedness of Failures and R&D Outcome Quality at Different Levels of Decomposability



REFERENCES

- Adler PS, Clark KB. 1991. Behind the learning-curve - a sketch of the learning-process. *Management Science* **37**(3): 267-281.
- Ahuja G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly* **45**(3): 425-455.
- Ahuja G, Katila R. 2001. Technological acquisitions and the innovation performance of acquiring firms: A longitudinal study. *Strategic Management Journal* **22**(3): 197-220.
- Ahuja G, Lampert CM. 2001. Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal* **22**(6-7): 521-543.
- Allen TJ. 1966. Studies of the problem-solving process in engineering design. *IEEE Transactions on Engineering Management* EM **13**(2): 72-83.
- Altshuller GS. 1999. The innovation algorithm: TRIZ, systematic innovation and technical creativity. Technical Innovation Center, Inc.
- Anand J, Oriani R, Vassolo RS. 2010. Alliance Activity as a Dynamic Capability in the Face of a Discontinuous Technological Change. *Organization Science* **21**(6).
- Anderson PA. 1983. Decision-making by objection and the cuban missile crisis. *Administrative Science Quarterly* **28**(2): 201-222.
- Anderson P, Tushman ML. 1990. Technological discontinuities and dominant designs - a cyclical model of technological-change. *Administrative Science Quarterly* **35**(4): 604-633.
- Argote L. 1996. Organizational learning curves: Persistence, transfer and turnover. *International Journal of Technology Management* **11**(7-8): 759-769.
- Argote L, Epple D. 1990. Learning-curves in manufacturing. *Science* **247**(4945): 920-924.
- Argote L, Beckman SL, Epple D. 1990. The persistence and transfer of learning in industrial settings. *Management Science* **36**(2): 140-154.
- Argyris C. 1990. *Overcoming organizational defenses: Facilitating organizational learning*. Boston: Allyn and Bacon.
- Argyris C, Schon D. 1978. *Organizational learning: A theory of action approach*. Reading, MA: Addison Wesley.

- Arundel A, Kabla I. 1998. What percentage of innovations are patented? empirical estimates for European firms. *Research Policy* **27**(2): 127-141.
- Audia PG, Locke EA, Smith KG. 2000. The paradox of success: An archival and a laboratory study of strategic persistence following radical environmental change. *Academy of Management Journal* **43**(5): 837-853.
- Baldwin CY, Clark KB. 2000. *Design rules: The power of modularity*: Vol. 1. MIT Press. Cambridge, MA
- Barnett WP, McKendrick DG. 2004. Why are some organizations more competitive than others? Evidence from a changing global market. *Administrative Science Quarterly* **49**(4): 535-571.
- Barron J. 1988 *Thinking and Deciding*. Cambridge University Press, New York.
- Basalla G. (Ed.) 1988. *The evolution of technology*. Cambridge University Press.
- Baum JAC, Dahlin KB. 2007. Aspiration performance and railroads ' patterns of learning from train wrecks and crashes. *Organization Science* **18**(3): 368-385.
- Baum JAC, Ingram P. 1998. Survival-enhancing learning in the Manhattan hotel industry, 1898-1980. *Management Science* **44**(7): 996-1016.
- Baumard P, Starbuck WH. 2005. Learning from failures: Why it may not happen? *Long Range Planning* **38**(3): 281-298.
- Belderbos R, Carree M, Lokshin B. 2004. Cooperative R&D and firm performance. *Research Policy* **33**(10): 1477-1492.
- Breschi S, Lissoni F, Malerba F. 2003. Knowledge-relatedness in firm technological diversification. *Research Policy* **32**(1): 69-87.
- Brouwer E, Kleinknecht A. 1999. Innovative output, and a firm's propensity to patent. An exploration of CIS micro data. *Research Policy* **28**(6): 615-624.
- Bruderer E, Singh JV. 1996. Organizational evolution, learning, and selection: A genetic-algorithm-based model. *Academy of Management Journal* **39**(5): 1322-1349.
- Cameron KS. 1984. Organizational adaptation and higher-education. *Journal of Higher Education* **55**(2): 122-144.
- Campbell DT. 1991. Methods for the experimenting society. *American Journal of Evaluation* **12**(3): 223-260.
- Cannon MD, Edmondson AC. 2005. Failing to learn and learning to fail (Intelligently): How great organizations put failure to work to innovate and improve. *Long Range Planning* **38**(3): 299-319.

- Capon N, Farley JU, Lehmann DR, Hulbert JM. 1992. Profiles of product innovators among large US manufacturers. *Management Science* **38**(2): 157-169.
- Carter EE. 1971. The behavioral theory of the firm and top-level corporate decisions. *Administrative Science Quarterly*. 413-429.
- Carter A. 1995. *Measuring the performance of the knowledgebased economy*, Working Paper. Brandeis University.
- Christensen CM. 1997. *The innovator's dilemma: when new technologies cause great firms to fail*. Harvard Business School Press: Boston, MA.
- Chuang YT, Baum JAC. 2003. It's all in the name: Failure-induced learning by multiunit chains. *Administrative Science Quarterly* **48**(1): 33-59.
- Clark KB, Fujimoto T. 1991. *Product development performance*. Harvard Business School Press, Boston, Massachusetts.
- Clark KB, Wheelwright SC. 1992. Organizing and leading heavyweight development teams. *California Management Review* **34**(3): 9-28.
- Cockburn IM, Henderson RM, Stern S. 2000. Untangling the origins of competitive advantage. *Strategic Management Journal* **21**(10-11): 1123-1145.
- Cooper AC, Schendel D. 1976. Strategic responses to technological threats. *Business Horizons* **19**(1): 61-69.
- Cyert RM, March JG. 1963. *A behavioral theory of the firm*. Englewood Cliffs, NJ.
- Denrell J. 2003. Vicarious learning, undersampling of failure, and the myths of management. *Organization Science* **14**(3): 227-243.
- Dess GG, Beard DW. 1984. Dimensions of organizational task environments. *Administrative Science Quarterly* **29**(1): 52-73.
- Dewar RD, Dutton JE. 1986. The adoption of radical and incremental innovations - an empirical-analysis. *Management Science* **32**(11): 1422-1433.
- Dosi G. 1988. Sources, procedures, and microeconomic effects of innovation. *Journal of Economic Literature* **26**(3): 1120-1171.
- Dougherty D. 1992. Interpretive barriers to successful innovation in large firms. *Organization Science* **3**(2): 179-202.
- Drews J. 2000. Drug discovery: A historical perspective. *Science* **287**(5460): 1960-1964.
- Duncker K. 1945. *On problem solving*. Psychological Monographs. 58(5, Whole No. 270).

- Eisenhardt KM, Tabrizi BN. 1995. Accelerating adaptive processes - product innovation in the global computer industry. *Administrative Science Quarterly* **40**(1): 84-110.
- Ethiraj SK, Levinthal D. 2004. Modularity and innovation in complex systems. *Management Science* **50**(2): 159-173.
- Ettlie JE, Bridges WP, Okeefe RD. 1984. Organization strategy and structural differences for radical versus incremental innovation. *Management Science* **30**(6): 682-695.
- Fleming L. 2001. Recombinant uncertainty in technological search. *Management Science* **47**(1): 117-132.
- Fleming L. 2012. *Breakthroughs and the "long tail" of innovation*. MIT Sloan Management Review. v10.
- Fleming L, Sorenson O. 2001. Technology as a complex adaptive system: evidence from patent data. *Research Policy* **30**(7): 1019-1039.
- Fleming L, Sorenson O. 2004. Science as a map in technological search. *Strategic Management Journal* **25**(8-9): 909-928.
- Foye WO. 2008. *Foye's principles of medicinal chemistry*. T. L. Lemke, & D. A. Williams (Eds.). Lippincott Williams & Wilkins.
- Frenken K. 2000. A complexity approach to innovation networks. The case of the aircraft industry (1909-1997). *Research Policy* **29**(2): 257-272.
- Ganco M. 2013. Cutting the Gordian knot: The effect of knowledge complexity on employee mobility and entrepreneurship. *Strategic Management Journal* **34**(6): 666-686.
- Gavetti G, Levinthal D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly* **45**(1).
- Gilfillan SC. 1935. *Inventing the ship*. Follett.
- Goleman D. 1985. *Vital Lies, Simple Truths: The Psychology of Self-Deception*. New York: Simon and Schuster.
- Greve HR. 1998. Managerial cognition and the mimetic adoption of market positions: What you see is what you do. *Strategic Management Journal* **19**(10): 967-988.
- Greve HR. 2003. A behavioral theory of R&D expenditures and innovations: Evidence from shipbuilding. *Academy of Management Journal* **46**(6): 685-702.
- Griliches Z. 1994. Productivity, R&D, and the data constraint. *The American Economic Review*: 1-23.
- Gulati R, Sytch M, Tatarynowicz A. 2012. The rise and fall of small worlds: Exploring the dynamics of social structure. *Organization Science* **23**(2): 449-471.

- Guler I, Nerkar A. 2012. The impact of global and local cohesion on innovation in the pharmaceutical industry. *Strategic Management Journal* **33**(5): 535-549.
- Gupta AK, Smith KG, Shalley CE. 2006. The interplay between exploration and exploitation. *Academy of Management Journal* **49**(4): 693-706.
- Hall B. 1993. Industrial research during the 1980s: Did the rate of return fall? *Brookings Paper on Economic Activity, Microeconomics* (2): 289-330.
- Hall BH, Mairesse J. 1995. Exploring the relationship between research-and-development and productivity in French manufacturing firms. *Journal of Econometrics* **65**(1): 263-293.
- Hargadon A, Sutton RI. 1997. Technology brokering and innovation in a product development firm. *Administrative Science Quarterly* **42**(4): 716-749.
- Haunschild PR, Ni Sullivan B. 2002. Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Administrative Science Quarterly* **47**(4): 609-643.
- Haunschild PR, Rhee M. 2004. The role of volition in organizational learning: The case of automotive product recalls. *Management Science* **50**(11): 1545-1560.
- Hayward MLA. 2002. When do firms learn from their acquisition experience? Evidence from 1990-1995. *Strategic Management Journal* **23**(1): 21-39.
- Hedberg BL, Nystrom PC, Starbuck WH. 1976. Camping on seesaws: Prescriptions for a self-designing organization. *Administrative Science Quarterly* 41-65.
- Heled Y. 2012. Why primary patents covering biologics should be unforceable against generic applicants under the Biologics Price Competition and Innovation act. *Annals of Health Law*, **21**: 211.
- Helfat CE. 1994. Evolutionary trajectories in petroleum firm research-and-development. *Management Science* **40**(12): 1720-1747.
- Helfat CE. 2000. Guest Editors' introduction to the Special Issue: Evolution of firm capabilities. *Strategic Management Journal*, Summer Special Issue **21**(10-11): 955-959.
- Henderson AD, Stern I. 2004. Selection-based learning: The coevolution of internal and external selection in high-velocity environments. *Administrative Science Quarterly* **49**(1): 39-75.
- Henderson R. 1995. Of life-cycles real and imaginary - the unexpectedly long old-age of optical lithography. *Research Policy* **24**(4): 631-643.
- Henderson R, Cockburn I. 1994. Measuring competence - exploring firm effects in pharmaceutical research. *Strategic Management Journal* **15**: 63-84.

- Henderson R, Clark KB. 1990. Architectural innovation - the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly* **35**(1): 9-30.
- Hill CWL, Rothaermel FT. 2003. The performance of incumbent firms in the face of radical technological innovation. *Academy of Management Review* **28**(2): 257-274.
- Huber GP. 1991. Organizational learning: the contributing processes and the literatures. *Organization Science* **2**(1): 88-115.
- Iansiti M. 2000. How the incumbent can win: Managing technological transitions in the semiconductor industry. *Management Science* **46**(2): 169-185.
- Jaffe AB. 1986. Technological opportunity and spillovers of r&d: Evidence from firms' patents, profits and market value. *American Economic Review*, **76**:984-1001.
- Jaffe AB. 1989. Characterizing the technological position of firms, with application to quantifying technological opportunity and research spillovers. *Research Policy* **18**(2): 87-97.
- Jaffe AB, Trajtenberg M, Henderson R. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics* **108**(3): 577-598.
- Jaikumar R, Bohn RE. 1992. A dynamic approach to operations management - an alternative to static optimization. *International Journal of Production Economics* **27**(3): 265-282.
- Jehn KA, Northcraft GB, Neale MA. 1999. Why differences make a difference: A field study of diversity, conflict, and performance in workgroups. *Administrative Science Quarterly* **44**(4): 741-763.
- Katila R, Ahuja G. 2002. Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal* **45**(6): 1183-1194.
- Kauffman S. 1993. *The origins of order: Self-organization and selection in evolution*. Oxford University Press.
- Khanna R, Guler I, Nerkar A. (forthcoming 2015). Fail often, fail big, and fail fast? learning from small failures and R&D performance in the pharmaceutical industry. *Academy of Management Journal*.
- Kim DJ, Kogut B. 1996. Technological platforms and diversification. *Organization Science* **7**(3): 283-301.
- Kobrin SJ. 1991. An empirical-analysis of the determinants of global integration. *Strategic Management Journal* **12**: 17-31.

- Kotabe M. 1990. The relationship between offshore sourcing and innovativeness of US multinational firms - an empirical-investigation. *Journal of International Business Studies* **21**(4): 623-638.
- Lane PJ, Lubatkin M. 1998. Relative absorptive capacity and interorganizational learning. *Strategic Management Journal* **19**(5): 461-477.
- Lapre MA, Mukherjee AS, Van Wassenhove LN. 2000. Behind the learning curve: Linking learning activities to waste reduction. *Management Science* **46**(5): 597-611.
- Lapre MA, Tsikriktsis N. 2006. Organizational learning curves for customer dissatisfaction: Heterogeneity across airlines. *Management Science* **52**(3): 352-366.
- Lawless MW, Anderson P. 1996. Generational technological change: effects of innovation and local rivalry on performance. *Academy of Management Journal* **39**(5): 1185-1217.
- Lee C, Lee K, Pennings JM. 2001. Internal capabilities, external networks, and performance: A study on technology-based ventures. *Strategic Management Journal* **22**(6-7): 615-640.
- Lee F, Edmondson AC, Thomke S, Worline M. 2004. The mixed effects of inconsistency on experimentation in organizations. *Organization Science* **15**(3): 310-326.
- Leonardbarton D. 1992. Core capabilities and core rigidities - a paradox in managing new product development. *Strategic Management Journal* **13**: 111-125.
- Levinthal DA. 1997. Adaptation on rugged landscapes. *Management Science* **43**(7): 934-950.
- Levinthal D, March JG. 1981. A model of adaptive organizational search. *Journal of Economic Behavior & Organization* **2**(4): 307-333.
- Levinthal DA, March JG. 1993. The myopia of learning. *Strategic Management Journal* **14**: 95-112.
- Levitan B, Lobo J, Kauffman SA, Schuler, R. 1999. Optimal organizational size in a stochastic environment with externalities. Santa Fe Institute.
- Levitt B, March JG. 1988. Organizational learning. *Annual Review of Sociology* **14**: 319-340.
- Li HY, Atuahene-Gima K. 2001. Product innovation strategy and the performance of new technology ventures in China. *Academy of Management Journal* **44**(6): 1123-1134.
- Lieberman MB. 1987. The learning-curve, diffusion, and competitive strategy. *Strategic Management Journal* **8**(5): 441-452.
- Lobo J, Kauffman SA, Schuler R. 1999. Optimal organizational size in a stochastic environment with externalities. *Santa Fe Institute*.

- Louis MR, Sutton RI. 1991. Switching cognitive gears - from habits of mind to active thinking. *Human Relations* **44**(1): 55-76.
- Lyles MA, Schwenk CR. 1992. Top management, strategy and organizational knowledge structures. *Journal of Management Studies* **29**(2): 155-174.
- Madsen PM, Desai V. 2010. Failing to learn? the effects of failure and success on organizational learning in the global orbital launch vehicle industry. *Academy of Management Journal* **53**(3): 451-476.
- Makri M, Hitt MA, Lane PJ. 2010. Complementary technologies, knowledge relatedness, and invention outcomes in high technology mergers and acquisitions. *Strategic Management Journal* **31**(6): 602-628.
- Mansfield E. 1965. Rates of Return from Industrial R&D. *American Economic Review* **55**(2): 310-322.
- March JG, Simon HA. 1958. *Organizations*. New York.
- March JG. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2**(1): 71-87.
- Marples DL. 1961. The decisions of engineering design. *IRE Transactions on Engineering Management* (2): 55-71.
- McFadyen MA, Cannella AA. 2004. Social capital and knowledge creation: Diminishing returns of the number and strength of exchange relationships. *Academy of Management Journal* **47**(5): 735-746.
- McGrath RG. 2011. Failing By design. *Harvard Business Review* **89**(4): 76-83.
- Mead C, Conway L. 1980. *Introduction to VLSI systems*. Addison-Wesley, Reading, MA.
- Miner AS, Bassof P, Moorman C. 2001. Organizational improvisation and learning: A field study. *Administrative Science Quarterly* **46**(2): 304-337.
- Mokyr J. 2002. *The gifts of Athena: Historical origins of the knowledge economy*. Princeton University Press.
- Mowery DC, Oxley JE, Silverman BS. 1998. Technological overlap and interfirm cooperation: implications for the resource-based view of the firm. *Research Policy* **27**(5): 507-523.
- Nelson RR. 1959. The simple economics of basic scientific-research. *Journal of Political Economy* **67**(3): 297-306.
- Nelson RR. 1962. *The rate and direction of inventive activity*. Princeton, NJ: Princeton University Press

- Nelson RR, Winter SG. 1982. *An evolutionary theory of economic change*. Cambridge: Belknap.
- Nesta L, Saviotti PP. 2005. Coherence of the knowledge base and the firm's innovative performance: Evidence from the US pharmaceutical industry. *Journal of Industrial Economics* **53**(1): 123-142.
- Newman ME. 2004. Detecting community structure in networks. *The European Physical Journal B-Condensed Matter and Complex Systems* **38**(2): 321-330.
- Newman ME, Girvan M. 2004. Finding and evaluating community structure in networks. *Physical Review E* **69**(2): 026113.
- Nohria N, Gulati R. 1996. Is slack good or bad for innovation? *Academy of Management Journal* **39**(5): 1245-1264.
- Nystrom PC, Starbuck WH. 1984. To avoid organizational crises, unlearn. *Organizational Dynamics* **12**(4): 53-65.
- Pakes A. 1986. Patents as options - some estimates of the value of holding European patent stocks. *Econometrica* **54**(4): 755-784.
- Pavitt K. 1988. Uses and abuses of patent statistics.
- Pavitt K. 1993. *What do firms learn from basic research?* Technology and the Wealth of Nations, London: Pinter
- Perman S. 2006. Some brilliant mistakes. Business Week
<http://www.businessweek.com/stories/2006-09-13/some-brilliant-mistakes>
 Accessed on July 15, 2014
- PhRMA. 2007. *Drug discovery and development: Understanding the R&D Process*. Pharmaceutical Research and Manufacturers of America (PhRMA): Washington, DC.
- Pinfield LT. 1986. A field-evaluation of perspectives on organizational decision-making. *Administrative Science Quarterly* **31**(3): 365-388.
- Ramanujam R, Goodman PS. 2003. Latent errors and adverse organizational consequences: a conceptualization. *Journal of Organizational Behavior* **24**: 815-836.
- Reason JT. 1997. *Managing the risks of organizational accidents* (Vol. 6). Aldershot: Ashgate.
- Reed R, DeFillippi RJ. 1990. Causal ambiguity, barriers to imitation, and sustainable competitive advantage. *Academy of Management Review* **15**(1): 88-102.
- Rivkin JW. 2000. Imitation of complex strategies. *Management Science* **46**(6): 824-844.

- Rivkin JW, Siggelkow N. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Science* **49**(3): 290-311.
- Rosenkopf L, Nerkar A. 2001. Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal* **22**(4).
- Rothaermel FT, Thursby M. 2007. The nanotech versus the biotech revolution: Sources of productivity in incumbent firm research. *Research Policy* **36**(6): 832-849.
- Sahal D. 1981. *Patterns of technological innovation* (pp. 201-202). Reading, MA: Addison-Wesley.
- Sahal D. 1985. Technological guideposts and innovation avenues. *Research Policy* **14**: 61-82
- Sanchez R, Mahoney JT. 1996. Modularity, flexibility, and knowledge management in product and organization design. *Strategic Management Journal* **17**: 63-76.
- Sarasvathy SD. 2003. Entrepreneurship as a science of the artificial. *Journal of Economic Psychology* **24**(2): 203-220.
- Sarasvathy SD, Simon HA 2000. Effectuation, near-decomposability, and the creation and growth of entrepreneurial firms. *In First Annual Research Policy Technology Entrepreneurship Conference*.
- Schilling MA. 2000. Toward a general modular systems theory and its application to interfirm product modularity. *Academy of Management Review* **25**(2): 312-334.
- Schumpeter J. 1934. *The theory of economic development*. Harvard University Press: Cambridge, MA.
- Schumpeter J. 1939. *Business cycles* (Vol. 1, pp. 161-74). New York: McGraw-Hill.
- Schmookler J. 1966. *Invention and economic growth*. Cambridge: Harvard University Press.
- Schumpeter J. 1982. *The Theory of Economic Development: an Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Transaction Publishers: Piscataway, NJ.
- Serrano CJ. 2010. The dynamics of the transfer and renewal of patents. *Rand Journal of Economics* **41**(4): 686-708.
- Simon HA. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* **106**(6): 467-482.
- Simon HA. 1969. *The sciences of the artificial* (Vol. 136). MIT Press.
- Simon HA. 1991. *The architecture of complexity* (pp. 457-476). Springer US.
- Simon HA. 1996. *The sciences of the artificial* 3rd Edition, MIT Press, Cambridge, MA.

- Sitkin SB. 1992. Learning through failure - the strategy of small losses. *Research in Organizational Behavior* **14**: 231-266.
- Smith RP, Eppinger SD. 1997. Identifying controlling features of engineering design iteration. *Management Science* **43**(3): 276-293.
- Sorenson O. 2003. Social networks and industrial geography. *Journal of Evolutionary Economics* **13**(5): 513-527.
- Sorenson O, Rivkin JW, Fleming L. 2006. Complexity, networks and knowledge flow. *Research Policy* **35**(7): 994-1017.
- Stuart TE, Podolny JM. 1996. Local search and the evolution of technological capabilities. *Strategic Management Journal* **17**: 21-38.
- Teece DJ, Rumelt R, Dosi G, Winter S. 1994. Understanding corporate coherence - theory and evidence. *Journal of Economic Behavior & Organization* **23**(1): 1-30.
- Thomke SH. 1997. The role of flexibility in the development of new products: An empirical study. *Research Policy* **26**(1): 105-119.
- Thomke SH. 1998. Managing experimentation in the design of new products. *Management Science* **44**(6): 743-762.
- Thomke SH. 2001. Enlightened experimentation - The new imperative for innovation. *Harvard Business Review* **79**(2): 66-75.
- Thomke SH. 2003. *Experimentation matters: unlocking the potential of new technologies for innovation*. Harvard Business Press.
- Thomke SH, von Hippel E, Franke R. 1998. Modes of experimentation: an innovation process - and competitive - variable. *Research Policy* **27**(3): 315-332.
- Thomke SH, Kuemmerle W. 2002. Asset accumulation, interdependence and technological change: Evidence from pharmaceutical drug discovery. *Strategic Management Journal* **23**(7): 619-635.
- Trajtenberg M. 1990. A penny for your quotes - patent citations and the value of innovations. *Rand Journal of Economics* **21**(1): 172-187.
- Tucker AL, Edmondson AC. 2003. Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change. *California Management Review* **45**(2): 55-72.
- Tushman ML, Anderson P. 1986. Technological discontinuities and organizational environments. *Administrative Science Quarterly* **31**(3): 439-465.

- Tyre MJ, vonHippel E. 1997. The situated nature of adaptive learning in organizations. *Organization Science* **8**(1): 71-83.
- Ulrich KT, Eppinger SD. 1999. *Product design and development*, 2nd ed. McGraw-Hill, New York.
- Usher A. 1954. *A history of mechanical invention*. Dover, Cambridge, MA.
- Vandeven AH. 1986. Central problems in the management of innovation. *Management Science* **32**(5): 590-608.
- Vandeven AH, Polley D. 1992. Learning while innovating. *Organization Science* **3**(1): 92-116.
- Vega, M. 2006. Does technological diversification promote innovation? An empirical analysis for European firms. *Research Policy* **35**(2): 230-246.
- Vincenti W. 1990. *What engineers know, and how they know it?* Johns Hopkins University Press, Baltimore, MA.
- Vonhippel E, Tyre MJ. 1995. How learning by doing is done - problem identification in novel process equipment. *Research Policy* **24**(1): 1-12.
- Ward, M. R. (1992). Drug approval overregulation. *Regulation*, 15: 47.
- Weick KE. 1976. Educational organizations as loosely coupled systems. *Administrative Science Quarterly* **21**(1): 1-19.
- Weiner B. 1971. *Perceiving the causes of success and failure*. General Learning Press.
- Weiner B. 1985. An attributional theory of achievement-motivation and emotion. *Psychological Review* **92**(4): 548-573.
- Weitzman M. 1996. Hybridizing growth theory. In: *Proceedings of the American Economics Association* 207-212.
- Wildavsky AB. 1988. *Searching for safety* (Vol. 10). Transaction Publishers.
- Yayavaram S, Ahuja G. 2008. Decomposability in knowledge structures and its impact on the usefulness of inventions and knowledge-base malleability. *Administrative Science Quarterly* **53**(2): 333-362.
- Yayavaram S, Chen WR. 2013. Changes in firm knowledge couplings and firm innovation performance: The moderating role of technological complexity. *Strategic Management Journal* **36**(3): 377-396.
- Zhou YM. 2013. Designing for complexity: Using divisions and hierarchy to manage complex tasks. *Organization Science* **24**(2): 339-355.