Fail Often, Fail Big, and Fail Fast?
Learning from Small Failures and R&D Performance
in the Pharmaceutical Industry

Rajat Khanna
Tulane University
7 McAlister Dr., New Orleans, United States
E-mail: rkhanna@tulane.edu
Phone: (919) 370-1325

Isin Guler
Sabanci University
Orhanli Tuzla
34956 Istanbul, Turkey
E-mail: isinguler@sabanciuniv.edu
Phone: +90-216-483-9670

Atul Nerkar
University of North Carolina at Chapel Hill
McColl Building, Chapel Hill, NC 27599, USA
E-mail: Atul_Nerkar@unc.edu
Phone: (919) 962-3272

Abstract

Do firms learn from their failed innovation attempts? Answering this question is important because failure is an integral part of exploratory learning. In this study, we explore whether and under what circumstances firms learn from their small failures in experimentation. Building on organizational learning literature, we examine the conditions under which prior failures influence firms' R&D output amount and quality. An empirical analysis of voluntary patent expirations (i.e., patents that firms give up by not paying renewal fees) in 97 pharmaceutical firms between 1980 and 2002 shows that the number, importance, and timing of small failures are associated with a decrease in R&D output (patent count) but an increase in the quality of the R&D output (forward citations to patents). Exploratory interviews suggest that the results are driven by a multi-level learning process from failures in pharmaceutical R&D. The findings contribute to the organizational learning literature by providing a nuanced view of learning from failures in experimentation.

Key words: learning from failure, small failures, experimentation, pharmaceutical industry, R&D performance, multilevel learning

INTRODUCTION

Failure is an integral part of the innovation process. Exploratory learning, a key building block of innovation, occurs through experimentation and search (March, 1991). The literature on organizational innovation emphasizes the importance of experimentation and the establishment of organizational structures and incentives that encourage it (e.g., Ahuja & Lampert, 2001; Cannon & Edmondson, 2005; Lee, Edmondson, Thomke, & Worline, 2004; Nohria & Gulati, 1996; Thomke & Kuemmerle, 2002). Inevitably, most experiments fail, and common wisdom suggests that such failures provide valuable feedback for future search efforts. However, learning from failure is far from automatic, given the psychological and organizational processes that attach negative meaning to failure (see Cannon and Edmondson, 2005 for a review.) This observation has led to a large body of prescriptive advice to embrace failure as a necessary part of the innovation process (e.g., Edmondson, 2011; McGrath, 2011). For instance, IDEO, an influential design firm known as one of the most innovative in the world, has a slogan that encourages experimentation and trial and error learning – *'fail often in order to succeed sooner'*. In spite of the central role of failure in experimentation, there is little empirical research examining whether and how firms learn from failure in innovation.

We aim to fill this gap by shedding light on learning from failure in the course of knowledge generation (March, 1991, cf. Sitkin, 1992). In particular, we ask whether and under what circumstances firms learn from their failed attempts at innovation. While prior literature provides many important insights about learning from failure (e.g., Baum & Dahlin, 2007; Haunschild & Sullivan, 2002; Hayward, 2002; Madsen & Desai, 2010; Miner, Kim, Holzinger, & Haunschild, 1999), most of these studies focus either on catastrophic failures such as plane or orbiter crashes (Haunschild & Sullivan, 2002, Madsen and Desai, 2010), or operational failures

such as acquisition integration problems or product defects (e.g., Hayward, 2002, Haunschild & Rhee, 2004, Henderson & Stern, 2004). Small and frequent failures arising in the natural course of experimentation are distinct from those studied in the prior literature, because failure in innovation is generally accepted as a likely, albeit unwelcome, outcome of the experimentation process. In contrast to operational failures, such as those in product safety, acquisitions, or product selection, where it is desirable to minimize the instances of failures, small failures in experimentation are often the only way to learn about causal relationships when a complete understanding of the underlying science is unavailable to boundedly rational decision makers (Sitkin, 1992; Fleming & Sorenson, 2004). Even though experimentation and failure are indispensable for innovation, learning from failed experiments in organizational contexts is far from straightforward (cf. Eggers, 2012a). Examining if and under what conditions firms can leverage prior failures to enhance innovation performance is therefore critical to our understanding of the innovation process (Sitkin, 1992).

In this study, we focus on a specific type of failure in the context of R&D efforts in the pharmaceutical industry. Specifically, we study how failed R&D projects in the pharmaceutical industry are associated with the subsequent R&D performance of firms. We observe firms' voluntary patent expirations before the legally allowed 20-year period, provided by the United States Patent and Trademark Office (USPTO), as an indicator of small failures in experimentation. Since the late 1980s, firms have been required to pay maintenance fees every four years to keep their patents active. Recent empirical evidence shows that firms let a substantial fraction of their patents expire before the regular expiration date by not paying maintenance fees (Serrano, 2010). Since pharmaceutical firms have patents ranging from several hundreds to several thousands, discontinuing some patents earlier than their legally permitted

expiration date may seem trivial. However, the effort and money invested in the discovery phase of each patent as well as the relative ease of renewing a patent suggest that firms make discontinuation decisions mindfully. Hence, the voluntary discontinuation of any patent is a self-admitted failure event. We ask whether and under what conditions these events create learning opportunities in R&D.

We examine how such small failures in experimentation influence an important component of innovation outcomes - R&D performance. Successful innovation entails high R&D performance, as well as competent commercialization (e.g., Fleming, 2002; Nerkar & Roberts, 2004; Schumpeter, 1934). In this study, we focus on the former step and ask how small failures influence subsequent R&D outcomes. In doing so, we distinguish between two important dimensions of R&D performance: R&D output amount (patent count) and R&D output quality (forward citations to patents). Organizational learning literature suggests a positive influence of learning on both dimensions, but most prior research on R&D either focuses on one dimension of R&D performance or combines the two into one aggregate measure. We submit that R&D output and quality are distinct outcomes, and test the influence of failures on each outcome separately in order to get a more complete picture of organizational learning following small failures.

The paper makes three important contributions. First, we contribute to the literature on learning from failure by examining small failures in exploration. Our study specifically focuses on the experimentation that underlies the concepts of search in innovation, and asks how failures in experimentation influence subsequent innovation outcomes. We extend and test ideas on learning from failures in exploration, building on prior work on intelligent failures (Sitkin, 1992) and learning from performance feedback (Greve, 2003, March & Simon, 1958). Second, we investigate conditions under which firms are more likely to learn from small failures in

experimentation. Prior literature suggests that experiences differ in terms of the learning they provide (Eggers, 2012b). In particular, we focus on the timing and importance of failures as sources of learning. We argue that early failures in R&D provide better learning opportunities for firms than do failures that come later in the R&D process. We also suggest that failures of projects that are relatively more important to the firm lead to higher learning.

This study also contributes to the literature on organizational innovation by examining failed innovation attempts as a determinant of a firm's subsequent R&D performance. While prior literature has studied many different factors that can influence R&D outcomes (e.g., Ahuja, 2000; Griliches, 1994; Henderson & Cockburn, 1994; Rothaermel & Thursby, 2007), few studies have investigated learning from prior failures as a contributing factor. We aim to fill this gap by disentangling the impact of failed experiments on both R&D output and quality.

We find that small failures in experimentation are important for the R&D performance of pharmaceutical firms but have opposite effects on R&D output and quality. Specifically, small failures in experimentation lead to a decrease in subsequent R&D output but an increase in the quality of R&D outcomes. We develop an understanding of this counterintuitive result through our field interviews. In particular, the distinction between idea generation and idea selection suggests a multi-level model of learning in pharmaceutical R&D.

ORGANIZATIONAL LEARNING FROM FAILURES

Early studies on organizational learning focused primarily on increasing efficiencies as a function of experience (e.g., Argote & Epple, 1990; Yelle, 1979). Recent work in this field has examined the effect of experience on such outcomes as service quality and survival rates of firms in service and banking industries (Baum & Dahlin, 2007; Baum & Ingram, 1998). Past experience also shapes the trajectory of a firm's innovations by increasing its absorptive capacity

(Cohen & Levinthal, 1990) and its competence (March, 1991; Sorensen & Stuart, 2000). Organizations primarily learn through a search process triggered by the feedback received from the environment (Greve, 2003; Levinthal & March, 1993; March, 1991).

Although a large body of literature has established that organizations learn from their experiences, relatively few studies distinguish whether the experience in question was a success or failure (cf. Sitkin, 1992). Failure offers firms many opportunities to learn, but learning from failure is far from guaranteed. Most organizations find it challenging to learn from failures (Cannon & Edmondson, 2005; Edmondson, 2002), due to the lack of sufficient information about the failure as well as a difficulty to agree on its causes (Eggers, 2012a). For instance, it is not uncommon that organizational members interpret causes of failure in a way that is most beneficial to themselves (Baumard & Starbuck, 2005).

If firms can at least partially overcome these challenges, failures can be an important source of learning. Failures may lead to process improvements, increase reliability, reduce rates of future failure, and decrease failure-related costs (Madsen & Desai, 2002; Haunschild & Sullivan, 2002; Baum & Dahlin, 2007; Kim & Miner, 2007). They enhance learning by challenging the understanding of the cause and effect relationships, helping firms replace existing routines and knowledge with more useful and accurate ones (e.g., Henderson & Stern, 2004; March, Sproull & Tamuz, 1991; Haunschild & Sullivan, 2002).

Moreover, failures may change the scope and direction of the organization's search activities. Success leads decision makers to remain on the same trajectory (Audia, Locke, & Smith, 2000), restricting the breadth of search to the neighborhood of existing knowledge (March, 1981). In contrast, failure to reach aspiration levels may trigger problemistic search, causing firms to look for solutions or alternatives that can address the problem of decreased

performance (Cyert & March, 1963; Greve, 2003). Failures may also provide firms with information to focus search in new directions (Wildavsky, 1988). In firms that hold a portfolio of innovations, trial-and-error learning often influences the composition of the portfolio of projects under consideration (e.g., Bower, 1970; Burgelman, 1983, 1991; Henderson & Stern, 2004).

In sum, as this brief review of the literature suggests, there is general agreement that failures are an important source of organizational learning, although few studies have specifically focused on learning from failures in the context of experimentation. In the next section, we describe our empirical context. We go on to develop and test our hypotheses about small failures in experimentation and R&D outcomes in subsequent sections.

CONTEXT: PATENT FAILURES IN THE PHARMACEUTICAL INDUSTRY

Patented R&D efforts of pharmaceutical firms provide the empirical context for our study of learning from small failures in experimentation. This setting provides an interesting and appealing context for several reasons. First, the highly research-intensive nature of the pharmaceutical industry (Henderson & Cockburn, 1994) makes it suitable for examining the R&D process and outcomes. More specifically, patents have a particular strategic importance in this industry (Grabowski & Vernon, 1992; Scott Morton, 2000). Pharmaceutical research is expensive and risky. It takes up to one billion dollars to bring a new drug into the market (Grabowski, 2002; Henderson & Cockburn, 1994; PhRMA, 2007), and the lengthy nature of the clinical trial process requires pharmaceutical firms to disclose sensitive proprietary knowledge, putting them at risk of imitation (De Carolis, 2003; Polidoro & Toh, 2011). Pharmaceutical firms rely on patents to protect the knowledge created within the firm (Gilbert & Shapiro, 1990; Grabowski, 2002; Klemperer, 1990). Luckily, patents provide a relatively effective method of guarding proprietary intellectual property for pharmaceutical firms (Levin, Klevorick, Nelson,

Winter, Glibert, & Griliches, 1987). Pharmaceutical firms patent every innovation possible (Cohen, Nelson, & Walsh, 2000; Levin et al., 1987; Paruchuri, Nerkar, & Hambrick, 2006) and start patenting early in the research process (Penner-Hahn & Shaver, 2005). A survey-based study found that pharmaceutical firms have the highest propensity to patent their innovations, with around 80% of their innovations protected by patents, as compared to an average of 35% across all industries studied in the survey (Arundel & Kabla, 1998). Often, pharmaceutical firms reward scientists on the number of patents produced (Stern, 2004). Previous research has established a positive correlation between the number of patents and profitability in science-based industries, such as pharmaceuticals (Cockburn & Griliches, 1988; Jaffe, 1986).

The second reason for choosing this context is that many patented inventions in the pharmaceutical industry end up as small failures in experimentation. Unlike other industries, such as electronics or computers, where patents are granted for products, patents in the pharmaceutical industry are granted early in the R&D process for research ideas that may or may not become products (Lehman, 2003). As in many innovative industries, the outcomes of R&D in the pharmaceutical industry are highly skewed (Scherer & Ross, 1990). Although patented ideas have crossed the first hurdle within the firm, they have a long way before they can turn into drugs and provide returns. In fact, most patents in the pharmaceutical industry do not lead to products. Firms typically apply for patents for lead compounds before or during the preclinical trials stage (Heled, 2012; Ward, 1992). Patented compounds then go through a long process of pre-clinical and human trials to identify suitable drugs (PhRMA, 2007). This process involves several stages. First, scientists try to understand the disease in question and its underlying causes through studies of changes in genes and how these changes can lead to the disease. Second, scientists look for a 'target', an altered gene or molecule, which can interact with a potential

drug. Third, scientists validate the identified target for its role in the disease and successful interaction with the drug molecule through extensive experiments. Once scientists have an understanding of the disease and the potential drug, they begin the process of finding the lead compounds that can be used to treat the disease. Fourth, after conducting preliminary safety tests and optimization studies, firms select a small number of compounds that are further tested in preclinical trials, a process that establishes the safety of drugs in animals before these drugs can be tested in humans. On average, of the 5,000-10,000 compounds tested in the fourth step, only around 250 are selected for preclinical testing (lead compounds). After extensive tests in preclinical studies, 1 to 5 lead compounds (drug candidates) are selected for further study in clinical trials. Firms file an investigational new drug (IND) application with the Food and Drug Administration (FDA) to begin clinical trials in humans for these drug candidates. About one out of every five drug candidates successfully clears all three phases of clinical trials and is commercialized in the market for the treatment of the disease (Grabowski, 2002). In sum, failures in pharmaceutical research mirror closely the idea of small failures in experimentation in this paper and are nicely captured in the patent data (Thomke & Kuemmerle, 2002; Thomke, 2003).

We examine the impact of small failures on two distinct dimensions of R&D performance - R&D output and R&D quality. R&D output has been used as a measure of innovation performance in many studies, as it represents a higher rate of innovation output (e.g., Ahuja, 2000; Gambardella, 1992; Cockburn & Henderson, 1998; Rothaermel & Thursby, 2007; Somaya, Williamson, & Zhang, 2007). R&D output "represents an externally validated measure of novelty (Griliches, 1990)" (Ahuja, 2000: 433) and has economic significance (Scherer & Ross, 1990). An increase in the R&D output not only suggests that the firm engages in more

experimentation, but also that the firm's engagement in experimentation is likely to lead to a larger diversity of solutions, which in turn increases the probability of finding a high-quality solution (Terwiesch & Xu, 2008; Terwiesch & Ulrich, 2009). The second measure of R&D performance is R&D output quality. Unlike output in manufacturing, where units without defect are identical, R&D output varies greatly in terms of quality (Cardinal, 2001). It is therefore important to understand whether and how prior failures influence the quality of R&D outcomes.

LEARNING FROM SMALL FAILURES IN EXPERIMENTATION

Small failures in experimentation provide opportunities for learning in multiple ways. First, frequency and scale of small failures provide firms with valuable feedback that helps shape the direction of the R&D portfolio. Firms that experience small failures on a frequent basis can form a more developed understanding of causal relationships without incurring large costs and can reallocate resources accordingly (Eggers, 2012a). Second, small failures encourage learning by initiating a search for the causes of such failures without threatening the existence of the firm or the decision makers, as larger, more visible failures may do. Since these failures do not endanger the survival of the firm, it is possible for the firm to both attend to such failures and engage in the search process. The experimental nature of these failures encourages decision makers within firms to analyze the outcomes more objectively (Baumard & Starbuck, 2005). Firms can learn from experimentation-driven failures internally before they make further investments in a particular product or innovation. 3M, one of the firms in our sample, does not favor following any master plan or engaging in complex strategic planning but rather promotes a culture in which employees are not afraid to "*try a lot of stuff and keep what works*" (Collins & Porras, 1994, p. 159).

Moreover, small failures in experimentation are associated with deliberate and mindful learning processes (Levinthal & Rerup, 2006, Sitkin, 1992). When firms have an understanding of the underlying scientific principles, they can come up with innovative ideas deductively (Fleming & Sorenson, 2004). But when the underlying science is poorly understood, and uncertainty with respect to outcomes is high, the firm needs to engage in experimentation in order to innovate (Terwiesch & Ulrich, 2009). As scientists can gain a better sense of the causal relationships by testing and ruling out hypotheses, new ideas generated through experimentation will on average be of higher quality.

In addition, experimentation stimulates learning by increasing the number and variety of solutions generated, and in turn, the quality of the final solution (Terwiesch & Xu, 2008; Thomke, 2003; Terwiesch & Ulrich, 2009). As firms actively engage in experimentation, they are likely to experience more failures, which in turn could lead to more innovative and successful outcomes (Thomke, 2003). Since such failures stimulate distant search and further experimentation, they are likely to result in a larger variability in the innovation output. Firms are then more likely to produce outliers in their innovation portfolios in terms of quality. Increasing variability enhances the likelihood of hitting 'home-runs' rather than producing consistently mid-range innovations in terms of impact. We therefore hypothesize that:

> *H1a: As a firm's small failure experience increases, its subsequent R&D output will increase.*

> *H1b: As a firm's small failure experience increases, its subsequent R&D output quality will increase.*

Failure of Important Projects

Small failures in experimentation may vary in terms of how important they were to the firm before the failure occurred. For instance, a failed project that has been endorsed by top management might be considered to be a failure of greater importance than one that was initiated by lower levels of management. In the present study, we define the importance of a small failure in terms of the project's expected performance before it eventually failed. In other words, decision makers in the firm had more expectations from a more important project before failure. As such, these failures are likely to be more visible within the firm. In the context of pharmaceutical R&D, a project's importance is highly related with its scientific value, or its impact on subsequent research. Note that the importance of a project is distinct from its size, which is often defined in terms of the investment at stake or consequences of failure. It is possible for a project that has required a lot of investment to generate little scientific or commercial value, just as it is possible for a small project in terms of outlays to generate high value. By definition, we focus only on small failures that are similar in the levels of investment and limited in consequences. Still, they may vary in terms of perceived potential and value.

We argue that important failures will elicit more learning for the following reasons. First, more important failures attract greater attention from decision makers within the firm. Given that managerial attention is scarce and selective, it is likely to be allocated to projects that were prominent before failure (Hoffman & Ocasio, 2001). Organizational decision makers may neglect to acknowledge failures in projects of lower importance, focusing on successful projects instead (Cannon & Edmondson, 2005; Madsen & Desai, 2010). Managers may fail to attend to the weak cues created by failures that deemed peripheral to the firm (Eggers, 2012a; Rerup, 2009). Failures of projects of little importance are not likely to challenge decision makers' core

beliefs (Baumard & Starbuck, 2005). In contrast, failures with a higher profile elicit surprise, are recognized more easily, and lead to changes in behavior more often, consequently affecting performance (Van de Ven, 1986).

In the context of small failures in experimentation, small failures of low importance to the firm may run the risk of going unnoticed, deliberately ignored, or being perceived at the aspiration level. In contrast, projects that were considered important to the firm before the failure trigger more extensive search for the causes and a more careful reallocation of R&D resources, which in turn is likely to improve the firm's subsequent R&D outcomes. We therefore expect firms that experience small failures of higher importance to enjoy a subsequent increase in their R&D performance.

> *Hypothesis 2a: As a firm experiences small failures of higher importance, its subsequent R&D output will increase.*

> *Hypothesis 2b: As a firm experiences small failures of higher importance, its subsequent R&D output quality will increase.*

Fail Early or Fail Late?

Prior research has argued that timing of experience is an important component of learning in the context of innovation (Eggers, 2012b). The effect of the timing of failures on R&D performance, however, is relatively understudied. As discussed in the previous section, failures provide firms with information on what might be wrong, and work as feedback that can improve performance going forward. The timing of the failure is important to the firms' learning outcomes, as it determines how quickly firms get feedback about a project. Experimental research shows that subjects who are given quick feedback will eliminate incorrect choices and

learn faster. Also, efficacy of feedback diminishes with the time elapsed in providing feedback (Skinner, 1954). Delayed feedback may be muddled with noise, making it hard for decision makers to assess relationships between actions and outcomes (Denrell, Fang, & Levinthal, 2004). Kettle and Haubl (2010) emphasize the importance of early feedback and its positive impact on performance. Sitkin (1992) suggests that small failures lead to most learning when they elicit quick feedback, so that the firm can learn, try new solutions, and generate new feedback.

In the context of innovation, the timing of a failure refers to the point in the R&D process at which a firm faces and acknowledges the failure. For instance, in the pharmaceutical industry, a compound may fail early in development, even before preclinical trials, or many years later, during late-stage clinical trials. Whether a firm will face failure relatively early or late in the development of an innovation depends on both the external environment and internal practices. If internal practices within the firm are oriented towards identifying failures, then the likelihood of spotting failures early and learning from them will be higher than when there is no specific attention given to the process. Sometimes firms are so constrained by the signals from the external environment that they have no choice but to wait before labeling a technology as a success or failure. Interim feedback helps firms identify potential problems and motivates them to engage in search to find solutions (Jordan & Audia, 2012). This idea also resonates with anecdotal evidence on the approach of highly innovative firms. For instance, Rich DeVaul, head of the Rapid Evaluation Team at Google noted: "Why put off failing until tomorrow or next week if you can fail now (Gertner, 2014)?"

Based on the discussion above, we expect the timing of feedback to influence R&D performance. Early feedback in the R&D process allows firms to manage available resources and limit allocation of resources to unproductive arenas. Also, early failures allow firms to

experiment in more ways as opposed to failures that come late in the R&D process. In contrast, when feedback comes late in the R&D process, it may be hard to pinpoint the exact decisions or actions that led to the failure, therefore confounding learning. Moreover, later failures may lead to escalation of commitment and cause a firm to continue related investments (Staw, 1976). R&D in high-tech industries is path dependent, and it is difficult for firms to change direction after significant progress has been made. A firm that receives early feedback on a technology may find it easier to reconfigure R&D investments and implement the learning from failures more effectively.

Hypothesis 3a: As a firm experiences small failures earlier in the R&D process, its subsequent R&D output will increase.

Hypothesis 3b: As a firm experiences small failures earlier in the R&D process, its subsequent R&D output quality will increase.

METHODS

We test our hypotheses with data on patent expirations in the pharmaceutical industry. We operationalize a small failure as a firm's decision to discontinue a patent, leading to its expiration before the end of its legal life of 20 years. In the late 1980s, USPTO made it necessary for firms to pay maintenance fees every 4, 8 and 12 years to keep their patents active. Any firm or entity failing to pay the fees at the scheduled time has its patent expired. The most likely reason for early discontinuation of a patent is the lack of relative value as perceived by the firm (Serrano, 2010).

After the introduction of the maintenance fees law in 1980, pharmaceutical firms have discontinued a large number of patents, with expired patents reaching up to 50% of total patents held (Serrano, 2010). Figure 1 shows the premature expiration of patents between 1985 and 2002

for 97 pharmaceutical firms analyzed in the current study. A quick look at the figure suggests that immediately following the introduction of patent law changes in 1980, firms discontinued only a small proportion of existing patents; but this share has increased dramatically in recent years, ranging between 40% and 60%. The small number of patents discontinued in the beginning is also suggestive of the critical role of firms' decision making. In the absence of relevant information on their patents during initial years, firms did not discontinue as many patents. With time, firms have started to use patent expirations to manage their patent portfolio.

Insert Figure 1 about here

Serrano (2010)'s finding that patents that are potentially less valuable are more likely to be discontinued suggests that a firm's decision to discontinue its patents prematurely is not random but rather deliberate in nature. As we discussed in the previous sections, patents are critical to the success of R&D in the pharmaceutical industry (Cockburn & Griliches, 1988; Jaffe, 1986). Interestingly, patent maintenance fees are negligible compared to the costs of acquiring a patent. Total fees to maintain a patent that has already been granted are less than $15,000 in total. Given the high potential value of a patent and the low maintenance fee, it is reasonable to think that firms would discontinue a patent only if they have good reason to believe that the patent has very limited value. In the time between the patent grant and the renewal date, the firm receives information about the future value of the patent. This information could be external, e.g., about other technological advancements, or internal, e.g., about the perceived viability of the project. Discontinuation of a patent suggests that the firm has received negative feedback. The high rate of discontinuation is consistent with the fact that most patented compounds fail before introduction to the market (Grabowski, 2002). Early patenting activity

represents the large number of "bets" taken by the pharmaceutical firms, while early expirations represent culling out unpromising opportunities from the research portfolio as a result of incoming information. With a reasonable understanding that the patent discontinuations represent small failure events, the uncertainty in estimating the value of a given patent in advance, and availability of specific checkpoints to evaluate the fates of patents make this setting appealing for a study of learning from small failures in experimentation.

Data and Sample

We obtained the data on patent expirations from USPTO. To be consistent with previous research, we considered 3-digit USPTO classes 514 and 424 to identify patents in pharmaceutical industry (Anand, Oriani, & Vassolo, 2010; Guler & Nerkar, 2012). These patents belonged to more than 200 firms. Since we examine the effect of patent discontinuations on R&D outcomes, we only kept firms in our sample that were active in patenting. Therefore, we removed firms that did not patent for more than 20 years between 1980 and 2002, leading to the final sample consisting of 97 pharmaceutical firms. Table 1 lists 30 major pharmaceutical firms out of the 97 firms[1] in our sample. The USPTO provides information on all expired patents and the stage at which patents were discontinued, i.e. first stage (after 4 years of application date), second stage (after 8 years of application date), or third stage (after 12 years of application date). As we measure learning at the firm level, we aggregated the number of discontinued patents at the firm-year level, leading to a final panel that contains 2015 firm-year observations[2]. After the introduction of the maintenance fee regime in December 1980 by the USPTO, the first set of

---

[1] Complete list of firms is available from authors upon request.
[2] Pharmaceutical industry experienced a big wave of merger and acquisitions (M&As) during the period of our study. In such cases, we kept the firms as separate entities before the M&A occurred and combined them in a single entity afterwards.

patents to expire did so in 1985 (after the first deadline of 4 years), so our dataset captures every discontinued patent following the regime change. Of 156,267 patents granted to the 97 firms in our sample, 56,630 patents (36.24%) were discontinued as of 2002. We tracked all variables between 1980 and 2002 in order to capture all patents that were at risk of discontinuation in 1985 and tracked all expired patents between 1985 and 2002.

Insert Table 1 about here

Dependent Variables

The dependent variable for this analysis is the R&D performance of the firm, measured as both firms' R&D output and the quality of their R&D output. We measured R&D output as the number of successful patent applications by a firm in a given year. Patent output has been used as a measure of R&D performance in many studies (Ahuja, 2000; Gambardella, 1992; Cockburn & Henderson, 1998; Nicholls-Nixon & Woo, 2003; Penner-Hahn & Shaver, 2005; Rothaermel & Thursby, 2007; Somaya, Williamson, & Zhang, 2007). Even though patents provide an imperfect measure of a firm's innovation output, "patents are tangible manifestations of a firm's ideas, techniques, and products, and are therefore an important indicator of innovation (DeCarolis & Deeds, 1999)" (Somaya, Williamson, & Zhang, 2007: 922). There are two main issues with using patents as a proxy for innovation (Griliches, 1990). First, not all innovations are patented. This is less of a problem in the pharmaceutical industry, where patents provide an effective way of protecting intellectual property (Levin et al., 1987) and are strategically essential to firms (Grabowski & Vernon, 1992; Scott Morton, 2000; Henderson & Cockburn, 1996). Second, not all patented inventions become innovations. Even so, patents have been shown to have a strong correlation with other firm-level innovation outputs, such as the number

of new product introductions (Comanor & Scherer, 1969; Basberg, 1982) and sales from new products (Comanor & Scherer, 1969) as well as innovative activity (Acs & Audretsch, 1989).

It is possible that patents that contribute to R&D performance in period t can become failures in periods after t, i.e. patents that firms produce in a given period can be discontinued prematurely at a future time. Although specifications of our model take into account this endogeneity, and the model provides efficient estimates, we used only the number of patents that did not get prematurely discontinued in calculating our dependent variable[3].

We measure the quality of R&D output for each firm as the total number of citations to all successful patents, i.e. patents that did not get discontinued prematurely. Numerous studies provide evidence of correlations between the importance of patents and citations to patents and have established the use of citations as a legitimate proxy for the quality of innovative or inventive performance (Jaffe, Trajtenberg, & Henderson, 1993; Pavitt, 1988; Trajtenberg, 1990). While prior literature favors a citation-weighted measure of R&D output, we examine R&D output and quality separately in order to gain a better understanding of how learning from failure influences R&D performance. Both measures are log-transformed due to skewness.

Independent Variables

Quantity, Timing, and Relative Importance of Small Failures: The quantity of a firm's small failures is measured as the number of patents that were discontinued due to non-payment of maintenance fees by the firm each year between 1985 and 2002. To test our arguments related to the timing of failures, we calculated the number of discontinued patents at each stage of patent discontinuation (at 4, 8, and 12 years). On average, we expect a patent discontinued at 4 years to

---

[3] The conclusions are robust to using the raw patent count as a measure of R&D output.

have elicited feedback earlier than a patent discontinued at 12 years[4]. In order to measure the

importance of a discontinued patent, we calculated the citations to the patent up until the year of

expiration. Number of forward citations is a commonly used measure of a patent's value and its

impact on future inventions (Jaffe, Trajtenberg, & Henderson, 1993; Pavitt, 1988; Trajtenberg,

1990), and captures the expectations of decision makers about the potential of a given patent

before expiration.

Control Variables

We used several control variables that could have confounding effects on R&D

performance. First, we controlled for the size of firms' R&D units, as large R&D units are likely

to have higher output. We calculated the size of the R&D unit by counting the number of

scientists with patent applications in each firm for every year (McFadyen & Cannella, 2004). As

firms may also source innovation externally (Ahuja, 2000; Ahuja & Katila, 2001; Sampson,

2007), we controlled for the count of the alliances for each firm in our sample for each year.

Geographic diversity may also affect innovation in multinational firms (MNCs) (Kobrin, 1991;

Lahiri, 2010). We therefore included a count of the number of countries that were represented in

the patent applications of a firm in a given year.

In this paper we are specifically interested in small failures in the form of patent

expirations in the pharmaceutical industry. Some patent expirations, however, may in fact

represent larger failures, such as a major failure in the firm's research agenda, or termination of a

project due to litigation. We included two controls to ensure that our sample indeed captures

learning from small failures. First, we controlled for the technological focus of the firm's

---

[4] While this assumption need not always be true (firms might receive early feedback and continue to maintain some patents), this noise is likely to make our results more conservative. Please see the Discussion section for a more detailed consideration of this possibility.

discontinued patents at the firm-year level. It is possible that a large number of failures

concentrated in a few technological classes in fact represent one big failure in the firm's research

portfolio, as opposed to many small failures. The technological focus variable was calculated

with the Herfindahl index for the technological classes of each firm's prematurely discontinued

patents in any given year. The value of this index ranges from zero to one. A higher value

indicates the presence of a large number of patents in a small number of technological classes,

and a value close to zero indicates that the discontinued patents were from different

technological classes. Second, we included the number of litigations faced by each firm in our

sample to account for the possibility of larger failures in the sample. We constructed this variable

using the number of lawsuits filed against each firm at both the state and federal levels in a given

year, as reported in Lexis-Nexis Legal Research database.

We calculated the moving average of each control variable for the past three years to

account for long lasting effects and smooth out sharp changes in these variables. In addition to

the time-varying variables, time-invariant variables specific to the firm are taken into account by

the Arellano-Bond method (described below) via first differencing. Figure 2 provides a brief

description of all variables and includes a sample data point for Abbott in 1998.

Insert Figure 2 about here

Empirical Model

The following models test the relationships between small failures and R&D performance

in a panel dataset for 97 firms for a period of 22 years:

$$P_{it} = \beta_1 P_{i,t-1} + \beta_2 F_{i,t-1} + \beta_3 C_{it} + u_{it} \tag{1}$$

$$Q_{it} = \beta_1 Q_{i,t-1} + \beta_2 F_{i,t-1} + \beta_3 C_{it} + v_{it} \tag{2}$$

In equation (1) P_{it} is R&D output, measured as the number of patents filed by firm i in year t, $P_{i,t-1}$ is the lagged value of R&D output for firm i, and $F_{i,t-1}$ is the vector of expired patent characteristics, including number, importance and timing, for firm i in period t - 1. C_{it} is the matrix of control variables and includes size of the R&D team, number of alliances, technological focus, number of litigations, and geographical diversity. Equation (2) has the same right-hand side variables; but the dependent variable (Q_{it}) is R&D output quality, measured as citations to the patents of firm i in period t, and $Q_{i,t-1}$ is firm i's R&D output quality in period t - 1.

At present, several empirical issues with equation (1) and (2) arise in the estimation. First, there is a possibility that small failures are endogenous to R&D performance, i.e., firms choose to give up patents because they believe that they are going to produce more patents in the next period. As such, causality can run in both directions, leading to correlation between the error term (u_{it} and v_{it} in equation (1) and (2) respectively) and independent variables ($F_{i,t-1}$).

Second, time-invariant fixed effects such as geography and firm culture may be correlated with the independent variables. Under such circumstances, unobserved fixed effects are combined with the error term, leading to a correlation between independent variables and the error term, a primary reason for biased coefficients. If Z_i's are the firm-specific, time-invariant fixed effects, and e_{it} is the observation-specific error term, the error term in equation (1) is:

$$u_{it} = Z_i + e_{it} \tag{3}$$

Third, both Equations (1) and (2) have lagged dependent variables to account for the dynamics in the process of patenting. Adding the lagged dependent variable on the right-hand side of the equation captures the effect of what firms patented in period t - 1 on what firms patent in period t instead of incorrectly attributing it to the other explanatory variables. However, this

can significantly inflate the coefficient on the lagged dependent variable and deflate coefficients

on other explanatory variables (Kelly, 2002).

To deal with these issues, we used the General Method of Moments (GMM) based

estimation introduced by Arellano and Bond (1991). The Arellano-Bond method is standard

when it comes to estimating dynamic panel models where time-variant fixed effects can

substantially influence the coefficients on the independent variables of interest. The Arellano-

Bond method, first proposed by Holtzeakin, Newey, & Rosen (1988), uses appropriate lags of

both dependent and independent variables as instruments for first-differenced dependent and

explanatory variables respectively. By using previous lags as instruments, the Arellano-Bond

model provides efficient estimates of parameters. Since both dependent and independent

variables are first differenced, time-invariant fixed effects are subtracted out in the model.

Previous research has used the Arellano-Bond model to resolve similar issues (David,

Yoshikawa, Chari, & Rasheed, 2006; Knott, Posen, & Wu, 2009; Milanov & Shepherd, 2013;

Uotila, Maula, Keil, & Zahra, 2009). After first differencing, equation (1) looks like this:

$$\Delta P_{it} = \beta_1 \Delta P_{i,t-1} + \beta_2 \Delta F_{i,t-1} + \beta_3 \Delta C_{it} + \Delta e_{it} \tag{4}$$

Since Z_i is time-invariant, it is subtracted out after differencing. Since the farthest lag that

appears for output is $P_{i,t-2}$ in equation (4), lags of four or higher can be used as instruments in

the model[5]. Analogously, lags of four or higher of endogenous variables can be used as

instruments for $\Delta F_{i,t-1}$. To check the validity of instruments used in the model, we performed

Sargan's test of overidentifying restrictions. We were not able to reject the null hypothesis that

overidentifying restrictions are valid (p-value of ~ 0.8), suggesting that instruments are

---

[5] We do not use the first lag as an instrument in our models because we found autocorrelation of first order. If, for
example, we have the term $P_{i,t-1}$ in the model, because of first order autocorrelation, we do not use $P_{i,t-2}$ as an
instrument but instead use $P_{i,t-3}$ and further lags as instruments.

uncorrelated with the residuals; hence estimates from the model are not biased (Davidson &

McKinnon, 1993). In addition, after comparing Wald chi-square values, which are quite high for

all our models, with critical values provided in Stock and Yogo (2005), we find no evidence of

weak instruments in the present study. The Arellano-Bond estimator is also designed for large N

and small T, which makes the use of this model for the present study more relevant. The same

approach is used to estimate Equation (2).

RESULTS

Table 2 presents the descriptive statistics and partial correlation matrix for our variables.

Since both measures of R&D performance used in this study are count variables and are skewed,

we used the logs[6] of R&D output and R&D output quality in order to correct for the non-

normality of the distribution (Chaganti & Damanpour, 1991; Ruef & Patterson, 2009). As

evident from Table 2, correlations between some of the independent variables are quite high. In

order to ensure that multicollinearity is not an issue in our model, we estimated the variance

inflation factors (VIFs). VIFs for all variables are less than 10, with an average value of 3.35 for

models testing Hypothesis 1 and 2.22 for models testing Hypotheses 2 and 3. Also, as the total

expired patents equals the sum of expired patents after 4, 8, and 12 years, we excluded total

expired patents from the model when testing for the separate effects of expired patents after 4, 8,

and 12 years.

---

[6] Usually it is appropriate to use Poisson or negative binomial models when the dependent variable is a non-negative integer. However, as explained, the most appropriate model for our data was the Arellano-Bond model, and this model does not allow a negative binomial or Poisson specification. We therefore use a log-transformed dependent variable in our models. Sensitivity analyses with count models are reported below.

Insert Table 2 about here

Tables 3 and 4 present the results of the Arellano-Bond models that use the logarithm of R&D output and R&D output quality as dependent variables, respectively. Our first set of analyses uses R&D output as a measure of R&D performance. Model 1 in Table 3 is the baseline model with only controls. Models 2 and 3 include the total number of expired patents and importance of failures, respectively. The total number of expired patents in Model 3 has a significant and negative effect on output. This is contrary to Hypothesis 1a, which predicts a positive relationship between number of failures and R&D output. Models 4, 5, and 6 contain number of expired patents after four, eight, and twelve years respectively. We included all independent variables except the total number of expired patents (to avoid multicollinearity) in Model 7 (full model.) The importance of failures has a negative and significant coefficient in Model 7, contrary to Hypothesis 2a, which predicts a positive relationship between the importance of failures and R&D output. The results do not support Hypothesis 3a that early small failures will lead to a higher R&D output, either. They suggest no effect of expired patents on output after four and twelve years, and a negative effect of expired patents on output after eight years. Control variables for R&D unit size and geographical diversity have expected signs and are significant at $p < 0.001$. Contrary to the expectation, the number of alliances has a negative and significant effect on output. With respect to the output measure of R&D performance, we found results contrary to our hypotheses. We discuss the implications of these findings in the subsequent section.

Insert Table 3 about here

Table 4 provides results with the logarithm of the quality of innovations (measured as the number of citations to the firm's patents) as the dependent variable. As before, Model 1 contains only control variables. Models 2 and 3 include the total number of failures (expired patents) and importance of failures (citations to expired patents), respectively. Model 3 provides support for Hypothesis 1b that there is a positive relationship between a firm's failures and its R&D output quality, as the coefficient on the total number of expired patents is positive and significant at $p < 0.001$. Models 4, 5, and 6 contain the number of expired patents after four, eight, and twelve years, respectively. We include all independent variables except the total number of expired patents (to avoid multicollinearity) in Model 7 (full model.) Hypothesis 2b, which predicts a positive relationship between the importance of failures and R&D output quality, is supported, as the coefficient on the importance of failures in Model 7 is positive and significant at $p < 0.001$. Hypothesis 3b, which predicts that early failures would increase R&D output quality more than later failures, is also supported, as the coefficient on expired patents after four years is positive and significant, whereas the coefficient of expired patents after eight years is not significant, and that of expired patents after 12 years is negative and significant. R&D unit size and number of alliances have a positive effect on innovation quality. Geographic diversity has no effect. In other words, we found support for all of our hypotheses with respect to the quality measure of R&D performance but not with respect to the output measure.

Insert Table 4 about here

Based on calculations from Models 3 and 7 (in Tables 3 and 4), pharmaceutical firms produced 0.2% fewer patents, while the quality of R&D output increased by 0.3% on average for each failure. Along the same lines, firms reduced the number of patent applications by 0.03% but increased patent quality by 0.04% following a unit increase in the importance of a failure. Firms in our sample filed for approximately two fewer patents and received six more citations per year for every 10 patents expired. Likewise, for an increase of 100 citations for expired patents, firms produced almost two fewer patents, and citations to their patents increased by eight per year.

Toward a Multi-Level Model of Learning from Failure

In order to better understand why an increase in small failures lead to an increase in the quality of R&D (citations) but a decline in R&D output (number of patents), we followed up with some interviews in the pharmaceutical industry. We conducted informal interviews with patent attorneys in pharmaceutical firms as well as in-depth interviews with four scientists who have worked with and patented in the organizations in our sample. The interviews were unstructured but all had the following open-ended questions: (1) What drives firm patenting behavior? (2) In your opinion, why do firms give up patents? (3) Are individual scientists aware that firms have not renewed their patents, and how do they react to such information? We refrained from asking leading questions that would either support or negate our hypotheses. Below are some conclusions from our interviews.

First, the decisions to file and renew patents are made at the firm level, through the intellectual property (IP) office. With respect to patent applications, all interviewees confirmed what prior research had documented. IP is critical for competitive advantage, and firms develop specific IP strategies to protect the drugs that they hope to launch in the market. Scientists are asked to file disclosures with the IP office as soon as they have reached a milestone in their

work. This could be a new molecule, a new mechanism, or a new application. The IP office makes decisions as to whether the firm will file for a patent or not. The decision to file for a patent in large, established firms is never made by the individual scientist.

Similarly, the decision to renew (or not) a patent is also made by the firm's IP office. Typically such decisions are based on a multitude of factors that include but are not limited to the performance of the R&D program / laboratory from which the patent originates (this could include failure in field trials, additional patents that negate earlier work, patents from competitors that negate the work done by the firm) and balancing of the overall IP portfolio in line with its corporate strategy. It is the patent attorneys that compile and distill the knowledge from failed patents and redirect the firm's patent portfolio. One interviewee suggested that one of the many responsibilities of a patent counsel is to look for commonalities between failed patents and existing patents to see if the existing portfolio can be improved. Interviewees concurred that individual scientists generally do not receive feedback about each discontinued patent, but they do receive information about future research and new technologies from the IP office.

How do non-renewed patents influence the scientists' work? The interviewees said that the non-renewal of a single patent was not a major decision either for the firm or for the scientist. One of the scientists we interviewed had a patent that was not renewed but he did not seem too concerned about it. That said, the same scientist mentioned that if many patents from the same inventor were not renewed, this would serve as feedback to the inventor. Also, most scientists are not likely to change their short term behavior based on the non-renewal of their patents, i.e. they would continue to turn in the patent disclosures to the IP office even though the firm could choose to reduce filing patents from them. All scientists explained that the IP office worked with respective program / laboratory managers to collect and analyze information, and decide whether

particular programs were effective, i.e. create knowledge that could to lead to IP that was useful. There was less of an emphasis on telling individual scientists to reduce or increase patent disclosures in particular areas.

The interviews suggest a multi-level perspective of organizational learning in the pharmaceutical industry. The R&D process can be conceived as a two-step process: Patentable ideas are generated in the first step, and these ideas are filtered (at multiple gates) at the second step – akin to an innovation tournament (Terwiesch & Ulrich, 2009). While individual scientists are the ones who work on research programs and produce patentable work, it is at the firm level (through the IP office) that these ideas are filtered based on value and place in the corporate R&D portfolio. The accumulated feedback is used to improve the firm's filters in selecting projects with higher expected returns. Adaptation of the firm's selection filters occurs relatively quickly, since IP officers constantly analyze past failures, have a clear view of the firm's overall R&D strategy, and can execute non-renewals relatively easily. In contrast, this knowledge may not immediately influence the generation of patents for several reasons: (1) individual scientists may not necessarily possess the knowledge about the overall health and direction of the research portfolio, and where their work stands relative to this portfolio. They may receive feedback from failed patents as a signal to change research direction only after several such failures have accumulated. (2) Even when scientists receive such feedback, the path-dependent nature of the search process (e.g., March & Simon, 1958; Nelson & Winter, 1982; Levinthal & March, 1993) is likely to lead them to continue working on what they did before. As put by one of our interviewees, "a general manager wants his scientists to start working on new technologies more often than scientists are willing to.' In other words, starting a brand new research agenda takes time. (3) Scientists' personal motives and incentives may cause them to receive the feedback in

an unfavorable way. Prior work suggests that scientists value independence and intrinsic rewards, and their output is highly correlated with the perceived importance and value of their work (e.g., Sauerman & Cohen, 2010). A concern with job security may cause threat-rigidity responses (Staw, Sandelands, & Dutton, 1981) and reduce output (Sauerman & Cohen, 2010). Feedback about non-renewals may therefore cause scientists to struggle with output.

In sum, the multi-level process of learning in the R&D process operates as follows. While the selection mechanism operates at the firm level and adapts to negative feedback relatively promptly through non-renewals, the idea generation mechanism operates at the individual scientist level and is slower to adapt to negative feedback. As a result, while IP officers are able to influence the overall quality of the R&D portfolio through non-renewals, they have little influence on the generation of new ideas. The lag in transferring knowledge to scientists, combined with the path-dependent nature of search and scientists' motives, cause a dip in the R&D output as individual scientists slowly adapt to incoming feedback.

An examination of the lag structure for our two dependent variables is consistent with this model. In supplementary analyses of different lag structures, we found that the decline in R&D output was more immediate, but the impact on R&D output quality was more gradual. This suggests that selection processes adapt to feedback from failures more promptly, causing a decline in the number of patents, followed by an increase in R&D output quality. Idea generation follows with a lag, and thus fails to compensate for the decline in R&D output in the short run.

Alternative Explanations and Robustness

We conducted several tests in order to rule out alternative explanations. A first potential explanation is that our results are due to the risk preferences of the organizational decision makers and not due to learning. Theory suggests two possible ways in which firms' risk

preferences may influence their responses to failure. On the one hand, performance feedback theory suggests that firms are likely to become more risk seeking following a deviation from their aspiration level (Cyert & March, 1963; Greve, 2003). On the other hand, firms may move away from riskier projects by increasing their risk threshold due to the 'hot-stove effect' (Denrell & March, 2001). In order to see if firms changed their criteria for project choice in response to failure, we examined whether and how small failures at time $t - 1$ influenced the total number of subclasses and number of novel subclasses in which the firms patented at time t. The total numbers of subclasses and novel subclasses represent the breadth of the firm's search, and have been used as measures of technological uncertainty involved in the search process as well as the level of technological exploration (Fleming, 2001; Rosenkopf & Nerkar, 2001). The results show no significant relationship between the number of total subclasses or novel subclasses and number of failures, providing preliminary evidence that firms do not become significantly more or less risk-seeking in their project choice in response to failure.

A second alternative explanation is that, even though we observe premature patent expirations as small failures, they may in fact be components of a larger, systemic failure, such as a failure of a research agenda or a drug that is already in the market. As explained above, we control for both the number of litigations involving the firm and the technological focus of the failed patents to rule out this possibility. According to Tables 3 and 4, the former variable was not significant in either model, while the latter was negatively significant in both the R&D output and quality models, suggesting that firms might struggle in R&D activities when failures are more concentrated in certain research areas. Still, these controls did not influence our core results that past failures decreased R&D output and increased R&D output quality.

Next, we address the alternative explanation that not all patent expirations are failures, and not all continuations are successes. Firms may renew patents for competitive or legal reasons, such as fending off a threat of litigation or for cross-licensing. We find it unlikely that such alternative explanations will systematically bias our results for several reasons. First, competitive and legal uses of patents in the pharmaceutical industry are limited compared to other industries such as electronics or computers (Hall, Helmers, von Graevenitz, & Rosazza-Bondibene, 2013; Lehman, 2003). Second, our conversations with patent attorneys in pharmaceutical firms supported the notion of failed patents as a vehicle for learning in a way consistent with our results. Third, we did not find any support for such uses in our empirical design, since number of litigations was not significant or weakly significant in our models. Finally, the noise in our measure of failure is likely to reduce the strength of our results, making our estimates conservative.

Another alternative explanation is that firms may learn from the experiences of all other firms, not just their own. This would suggest a different mechanism of vicarious learning rather than the one we put forth. In unreported models, we controlled for the number, importance and timing of failures of all firms in the industry other than the focal firm. These population-level controls revealed some interesting patterns. For instance, the total number of failures at the population level is associated with a drop in the firm's R&D output, but does not significantly affect the firm's R&D output quality, possibly suggesting that firms may find it harder to learn from other firms' failures. At the same time, these controls did not significantly change our results, ruling out population-level learning as an alternative explanation.

We also tested whether learning from failures in our context is cumulative in nature. Since previous values of the variables are used as instruments to estimate current values of both

the dependent and the independent variables in the Arellano-Bond model, using cumulative

measures can potentially lead to overestimation of the contribution of the previous periods.

Instead we used the cumulative measure of our independent variables with an annual forgetting

factor of 0.8 in the fixed effects negative binomial model (Darr, Argote, & Epple, 1995). The

results are consistent with those of the previous models.

As we only observe discontinuations at the legally determined four, eight and twelve

years, we may be missing the time at which discontinuation decisions are actually made. For

instance, it is possible that a molecule associated with a patent fails at most four years earlier

than we observe it in our data. In order to take into account this possibility, we assumed that the

discontinuation decisions were made at an average of two years before we observed them, and

used independent variables from $t-2$ to test our hypotheses. This test led to fewer observations

but the results were robust.

We also tested for the presence of curvilinearity in our models. Although we find no

support for a curvilinear relationship between the number and importance of failures and R&D

output, we do find the presence of a curvilinear relationship between number and importance of

failures and R&D output quality. However, the inflection point of the curve is estimated to be at

5,263 patent expirations, well beyond the maximum of 1,076 patent expirations observed in our

sample. We also repeated our models with negative binomial regression models with firm fixed

effects (Hausman, Hall, & Griliches, 1984; Henderson & Cockburn, 1994), and found that the

results are similar to those from previous specifications.

In short, the empirical patterns that we observed in the data, the supplementary evidence

ruling out alternative mechanisms, and the qualitative evidence combined, suggest that our

results reflect learning processes, and that learning from failure in pharmaceutical firms is a

multi-level process. Project selection mechanisms at the firm level adapt to feedback from failed patents faster than generation mechanisms at the individual scientist level, causing us to observe an increase in R&D output quality and a decline in R&D output.

DISCUSSION AND CONCLUSION

Previous research on learning from past experience, including failures, has argued that experience in general leads to higher productivity and lower unit cost in manufacturing and service industries (Argote, 1996; Argote & Epple 1990). However, results obtained in the current study tell a different story. Specifically, we find that as the number of failures increases, R&D output, measured as the number of patents filed by the firm, decreases, whereas the quality of the patents, measured as citation to those patents, increases. Similarly, the importance and timing of small failures cause opposing changes in R&D output and R&D output quality. This difference in the effect of number, timing and importance of failures on R&D output and quality points to the potentially different drivers of these two dimensions of R&D performance.

We offer a potential reconciliation of our results based on our interviews with industry participants. As opposed to many studies that depict organizations as monolithic entities, our interviews suggest a multi-level model of learning, in line with process models of resource allocation (Bower, 1970; Burgelman, 1983, 1991; Noda & Bower, 1996; cf. Gaba & Joseph, 2013 for a consideration of divergent aspiration levels at multiple levels.) To reiterate, we suggest that patent generation occurs at the individual scientist level, whereas patent selection for the overall portfolio takes place at the level of the IP office. While the selection mechanisms adjust more easily to feedback from failed projects, the idea generation process adjusts more slowly, due to the lag in relaying feedback to individual scientists, the path dependency of search processes, and the motives of individual scientists. As a result, fast adaption of selection leads to

higher quality patents, whereas slow adaptation of generation leads to an immediate drop in R&D output.

We contribute to the organizational learning literature in the following ways. First, our study confirms the critical role of failures for organizational performance and adds to prior studies by examining small failures in experimentation – an understudied but important type of failure (e.g., Sitkin, 1992). We show in this paper how such failures can provide firms with critical feedback in exploration. Figure 1 suggests that learning from failures did not reduce the number of subsequent failures in our study, as opposed to the patterns in other types of failures covered in prior research (e.g., Haunschild & Rhee, 2004; Haunschild & Sullivan, 2002; Madsen & Desai, 2010). This may suggest that firms do not use small failures to increase the reliability of their processes but to expand their search and identify new directions. This interesting observation underscores the unique nature of small failures in experimentation and the importance of studying them separately from other failures.

In addition, our finding that small failures lead to a decrease in R&D output but an increase in quality is a novel contribution to the literature on innovation. While prior studies on R&D outcomes have examined different factors influencing R&D outcomes, the role of prior failures has so far been neglected. Our findings on the timing and importance of small failures also provide novel insights on how firms learn from failures in the context of innovation and add to the understanding of the innovation process.

We believe there are certain boundary conditions for these findings. For instance, firms in the pharmaceutical industry are faced with extreme uncertainty and long development times. The outcomes of the R&D process are highly skewed, with few projects generating most of the returns. The abundance of generated ideas and the need to filter them necessitate the separation

of idea generation and selection processes. Our results are likely to hold in other innovative industries with similar characteristics, such as venture capital, corporate venturing, or creative endeavors such as film production. Future studies can examine to what extent these results are generalizable to other R&D intensive industries. Do firms in different industries perceive patents differently in terms of their value? Do such firms vary in how they learn from such failures? These are some of the questions to be explored to broaden the understanding of learning from small failures in the context of innovation.

Our findings also have implications for practicing R&D managers. While the study by no means recommends failures, it encourages experimentation and conscious attention to evaluation of patents as early as possible. Higher rate of experimentation will increase the number of failures but also the likelihood of finding the right bet. Similarly, our findings suggest that firms should try to remove subpar projects from their portfolio as early as possible, and not be afraid to eliminate important projects from the portfolios.

Despite its contributions, the study has several limitations. First, we could not observe the underlying reasons for patent failures or discriminate between these underlying reasons as determinants of learning. We leave it to future studies to disentangle the relationships between different reasons for patent expirations and the subsequent learning outcomes. At the same time, the process by which these failures lead to improvements in R&D performance also goes unobserved. While we believe that the improvements are both in ongoing projects as well as in the reallocation of R&D resources, we did not distinguish between these two mechanisms and their relative importance in this study. Future studies can greatly enhance our understanding of underlying mechanisms governing learning from failures. Last, future studies could provide

more robust analysis of the multi-level model of learning at the idea generation and selection steps.

It is worth noting that patent expirations represent a distinct type of failure in the pharmaceutical R&D process. By studying all patents, as opposed to those that failed in clinical trials, we are able to get a more complete picture of small failures in experimentation. It is also possible that some ideas fail even before reaching the patenting stage. Given the high rate of patenting in the pharmaceutical industry, these early ideas probably show little promise to begin with. Future studies may find it fruitful to compare these different kinds of failures in terms of learning outcomes.

REFERENCES

Acs, Z., & Audretsch, D. B. 1989. Patents as a measure of innovative activity. *Kyklos*, 42(2): 171-180.

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

Anand, J., Oriani, R., & Vassolo, R. S. 2010. Alliance activity as a dynamic capability in the face of a discontinuous technological change. *Organization Science,* 21(6): 1213-1232.

Arellano, M. & Bond, S. 1991. Some tests of specification for panel data - monte-carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2): 277-297.

Argote, L., & Epple, D. 1990. Learning-curves in manufacturing. *Science*, 247(4945): 920-924.

Argote, L. 1996. Organizational learning curves: Persistence, transfer and turnover. *International Journal of Technology Management*, 11(7-8): 759-769.

Arundel, A., & Kabla, I. 1998. What percentage of innovations are patented? Empirical estimates for European firms. *Research policy*, 27(2): 127-141.

Audia, P. G., Locke, E. A., & Smith, K. G. 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.

Basberg, B. L. 1982. Technological change in the Norwegian whaling industry: A case-study in the use of patent-statistics as a technology indicator. *Research Policy,* 11(3): 163-171.

Baum, J. A. C., & Ingram, P. 1998. Survival-enhancing learning in the Manhattan hotel industry, 1898-1980. *Management Science*, 44(7): 996-1016.

Baum, J. A. C., & Dahlin, K. B. 2007. Aspiration performance and railroads ' patterns of learning from train wrecks and crashes. *Organization Science*, 18(3): 368-385.

Baumard, P., & Starbuck, W. H. 2005. Learning from failures: Why it may not happen. *Long Range Planning*, 38(3): 281-298.

Beukel, K., Valentin, F., & Jensen, R. L. 2011. How Patent Types effect Patent Value. *In 6th Annual Conference of the EPIP Association*.

Bower, J. 1970. *Managing the Resource Allocation Process,* Irwin, Homewood, IL.

Burgelman, R. A. 1983. A process model of internal corporate venturing in the diversified major firm. *Administrative Science Quarterly*, 223-244.

Burgelman, R. A. 1991. Intraorganizational ecology of strategy making and organizational adaptation: Theory and field research. *Organization science*, 2(3): 239-262.

Cannon, M. D., & Edmondson, A. C. 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.

Cardinal, L. B. 2001. Technological innovation in the pharmaceutical industry: The use of organizational control in managing research and development. *Organization Science,* 12(1): 19-36.

Chaganti, R., & Damanpour, F. 1991. Institutional ownership, capital structure, and firm performance. *Strategic Management Journal*, 12(7): 479-491.

Christie, A. F., Dent, C., McIntyre, P., Wilson, L., & Studdert, D. M. 2013. Patents associated with high-cost drugs in Australia. *PloS one*, 8(4): e60812.

Cockburn, I., & Henderson, R. M. 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. *The Journal of Industrial Economics*, 46(2): 157-182.

Cockburn, I., & Griliches, Z. 1988. Industry effects and appropriability measures in the stock markets valuation of r-and-d and patents. *American Economic Review*, 78(2): 419-423.

Cohen W. M., & Levinthal D. A. 1990. Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly* 35(1): 128-152.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. 2000. Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not). *National Bureau of Economic Research*.

Collins, J. C., & Porras, J.1994. *Built to Last: successful habits of visionary companies*. New York: HarperBusiness.

Comanor, W. S., & Scherer, F. M. 1969. Patent statistics as a measure of technical change. *The Journal of Political Economy*, 392-398.

Cyert, R. M., & March, J. G. 1963. *A behavioral theory of the firm,* Englewood Cliffs, NJ: Prentice Hall.

Darr, E. D., Argote, L., & Epple, D. 1995. The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management science*, 41(11): 1750-1762.

David, P., Yoshikawa, T., Chari, M. D. R., & Rasheed, A. A. 2006. Strategic investments in Japanese corporations: Do foreign portfolio owners foster underinvestment or appropriate investment? *Strategic Management Journal*, 27(6): 591-600.

Davidson, R., & Mackinnon, J. G. 1993. *Estimation and Inference in Econometrics (2nd edn).* New York: Oxford University Press.

DeCarolis, D. M. 2003. Competencies and imitability in the pharmaceutical industry: An analysis of their relationship with firm performance. *Journal of Management,* 29(1): 27-50.

DeCarolis, D. M., & Deeds, D.L., 1999. The impact of stocks and flows of organizational knowledge on firm performance: an empirical investigation of the biotechnology industry. *Strategic Management Journal,* 20: 953 – 968.

Denrell, J., & March, J. G. 2001. Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5): 523-538.

Denrell, J., Fang, C., & Levinthal D. A. 2004. From T-Mazes to labyrinths: Learning from model-based feedback. *Management Science,* 50(10): 1366-1378.

Edmondson, A. C. 2002. The local and variegated nature of learning in organizations: A group-level perspective. *Organization Science*, 13(2): 128-146.

Edmondson, A. C. 2011. Strategies of learning from failure. *Harvard Business Review,* 89(4): 48-55.

Eggers, J. P. 2012a. Falling Flat: Failed Technologies and Investment under Uncertainty. *Administrative Science Quarterly,* 57(1): 47-80.

Eggers, J. P. 2012b. All experience is not created equal: learning, adapting, and focusing in product portfolio management. *Strategic Management Journal,* 33(3): 315-335.

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management science*, 47(1): 117-132.

Fleming, L. 2002. Finding the organizational sources of technological breakthroughs: the story of Hewlett-Packard's thermal ink-jet. *Industrial and Corporate Change*, 11(5): 1059-1084.

Fleming, L., & Sorenson, O. 2004. Science as a map in technological search. *Strategic Management Journal*, 25(8-9): 909-928.

Gambardella, A. 1992. Competitive advantages from in-house scientific research: the US pharmaceutical industry in the 1980s. *Research Policy*, 21(5): 391-407.

Gertner, J. 2014. The truth about google x: an exclusive look behind the secretive lab's closed doors. *Fastcompany,* http://www.fastcompany.com/3028156/united-states-of-innovation/the-google-x-factor. Accessed on April 30, 2014.

Gilbert, R., & Shapiro, C. 1990. Optimal patent length and breadth. *The Rand Journal of Economics,* 106-112.

Grabowski, H. G. 2002. Patents, Innovation, and Access to New Pharmaceuticals. *Journal of International Economic Law,* 5(4): 849-860.

Grabowski, H. G., & Vernon, J. M. 1992. Brand loyalty, entry, and price competition in pharmaceuticals after the 1984 Drug Act. *Journal of Law and Economics*, 35: 331.

Greve, H. R. 2003. *Organizational learning from performance feedback: A behavioral perspective on innovation and change*. Cambridge University Press.

Griliches, Z. 1990. Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28: 1661-707.

Griliches, Z. 1994. Productivity, R&D, and the data constraint. *The American Economic Review*, 1-23.

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.

Hall, B., Helmers, C., von Graevenitz, G., & Rosazza-Bondibene, C. 2013. Technology entry in the presence of patent thickets. *NBER*.

Haunschild, P. R., & Ni Sullivan, B. 2002. Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Administrative Science Quarterly*, 47(4): 609-643.

Haunschild, P. R., & Rhee, M. 2004. The role of volition in organizational learning: The case of automotive product recalls. *Management Science*, 50(11): 1545-1560.

Hausman, J. A., Hall, B. H., & Griliches, Z. 1984. Econometric models for count data with an application to the patents-R&D relationship. *NBER*.

Hayward, M. L. A. 2002. When do firms learn from their acquisition experience? Evidence from 1990-1995. *Strategic Management Journal*, 23(1): 21-39.

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.

Henderson, A. D., & 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., & Cockburn, I. 1994. Measuring competence - exploring firm effects in pharmaceutical research. *Strategic Management Journal*, 15: 63-84.

Henderson, R., & Cockburn, I. 1996. Scale, scope, and spillovers: the determinants of research productivity in drug discovery. *The Rand journal of economics*, 32-59.

Hoffman, A. J., & Ocasio, W. 2001. Not all events are attended equally: Toward a middle-range theory of industry attention to external events. *Organization Science,* 12(4): 414-434.

Holtzeakin, D., & Newey, W., & Rosen, H. S. 1988. Estimating vector autoregressions with panel data. *Econometrica*, 56(6): 1371-1395.

Jaffe, A.B. 1986. Technological opportunity and spillovers of r&d: Evidence from firms' patents, profits and market value. *American Economic Review,* 76:984-1001.

Jaffe, A. B., Trajtenberg, M., & Henderson, R. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly journal of Economics,* 108(3): 577-598.

Jordan, A. H., & Audia, P. G. 2012. Self-enhancement and learning from performance feedback. *Academy of Management Review*, 37(2): 211-231.

Kelly, N. J. 2002. *The nature and degree of bias in lagged dependent variable models*. Annual Meeting of the Southern Political Science Association, GA.

Kettle, K. L., & Haeubl, G. 2010. Motivation by anticipation: Expecting rapid feedback enhances performance. *Psychological Science*, 21(4): 545-547.

Kim, J. Y., & Miner, A. S. 2007. Vicarious learning from the failures and near-failures of others: Evidence from the US commercial banking industry. *Academy of Management Journal,* 50(3): 687-714.

Klemperer, P. 1990. How broad should the scope of patent protection be?. *The Rand Journal of Economics,* 113-130.

Knott, A. M., Posen, H. E., & Wu, B. 2009. Spillover Asymmetry and Why It Matters. *Management Science*, 55(3): 373-388.

Kobrin, S. J. 1991. An empirical-analysis of the determinants of global integration. *Strategic Management Journal*, 12: 17-31.

Lahiri, N. 2010. Geographic distribution of R&D activity: how does it affect innovation quality? *Academy of Management Journal*, 53(5): 1194-1209.

Lee, F., Edmondson, A. C., Thomke S., & Worline M. 2004. The mixed effects of inconsistency on experimentation in organizations. *Organization Science* 15(3): 310-326.

Lehman, B. 2003. *The pharmaceutical industry and the patent system*. International Intellectual Property Institute.

Levin, R. C., Klevorick, A. K., Nelson, R. R., Winter, S. G., Gilbert, R., & Griliches, Z. 1987. *Appropriating the returns from industrial research and development*. Brookings papers on economic activity, 783-831.

Levinthal, D. A., & March, J. G. 1993. The myopia of learning. *Strategic Management Journal*, 14: 95-112.

Levinthal, D., & Rerup, C. 2006. Crossing an apparent chasm: Bridging mindful and less-mindful perspectives on organizational learning. *Organization Science* 17(4): 502-513.

Madsen, P. M., & 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.

March, J. G. 1981. Footnotes to organizational-change. *Administrative Science Quarterly*, 26(4).

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2(1): 71-87.

March, J. G., & Simon H. A. 1958. *Organizations*. New York: Wiley.

March, J. G., Sproull, L. S., & Tamuz, M. 1991. Learning from samples of one or fewer. *Organization science*, 2(1): 1-13.

McFadyen, M. A., & Cannella, A. A. 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, R. G. 2011. Failing by design. *Harvard Business Review*, 89(4): 77-83.

Milanov, H., & Shepherd, D. A. 2013. The importance of the first relationship: The ongoing influence of initial network on future status. *Strategic Management Journal*. doi: 10.1002/smj.2109.

Miner, A. S., Kim, J. Y., Holzinger, I. W., & Haunschild, P. 1999. Fruits of failure: Organizational failure and population-level learning. *Advances in Strategic Management, Vol 16 - 1999: Population-Level Learning and Industry Change*, 16: 187-220.

Nelson, R. R., & Winter, S. G. 1982. *An evolutionary theory of economic change*. Cambridge: Belknap.

Nerkar, A., & Roberts, P. W. 2004. Technological and product-market experience and the success of new product introductions in the pharmaceutical industry. *Strategic Management Journal*, 25(8-9): 779-799.

Nicholls-Nixon, C. L., & Woo, C. Y. 2003. Technology sourcing and output of established firms in a regime of encompassing technological change. *Strategic Management Journal*, 24(7): 651-666.

Noda, T., & Bower, J. L. 1996. Strategy making as iterated processes of resource allocation. *Strategic Management Journal*, 17(S1): 159-192.

Nohria, N., & Gulati, R. 1996. Is slack good or bad for innovation? *Academy of Management Journal,* 39(5): 1245-1265.

Paruchuri, S., Nerkar, A., & Hambrick, D. C. 2006. Acquisition integration and productivity losses in the technical core: Disruption of inventors in acquired companies. *Organization Science*, 17(5): 545-562.

Pavitt, K. 1988. *Uses and abuses of patent statistics*. In: van Raan, A.F.J. (Ed.), Handbook of Quantitative Studies of Science and Technology, Amsterdam: Elsevier.

Penner-Hahn, J., & Shaver, J. M. 2005. Does international research and development increase patent output? An analysis of Japanese pharmaceutical firms. *Strategic Management Journal*, 26(2): 121-140.

PhRMA. 2007. *Drug Discovery and Development: Understanding the R&D Process*. Pharmaceutical Research and Manufacturers of America (PhRMA): Washington, DC.

Polidoro, F., & Toh, P. K. 2011. Letting rivals come close or warding them off? The effects of substitution threat on imitation deterrence. *Academy of Management Journal,* 54(2): 369-392.

Rerup, C. 2009. Attentional Triangulation: Learning from Unexpected Rare Crises. *Organization Science,* 20(5): 876-893.

Rosenkopf, L., & Nerkar, A. 2001. Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal,* 22(4): 287-306.

Rothaermel, F. T., & Thursby, M. 2007. The nanotech versus the biotech revolution: Sources of productivity in incumbent firm research. *Research Policy*, 36(6): 832-849.

Ruef, M., & Patterson, K. 2009. Credit and Classification: The Impact of Industry Boundaries in Nineteenth-century America. *Administrative Science Quarterly*, 54(3): 486-520.

Sampson, R. C. 2007. R&D alliances and firm performance: the impact of technological diversity and alliance organization on innovation. *Academy of Management Journal.* 50(2): 364-386.

Sauermann, H., & Cohen, W. M. 2010. What makes them tick? Employee motives and firm innovation. *Management Science*, 56(12): 2134-2153.

Scherer, F. M., & Ross, D. 1990. *Industrial market structure and economic performance.* University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship.

Serrano, C. J. 2010. The dynamics of the transfer and renewal of patents. *The Rand Journal of Economics*, 41(4): 686-708.

Schumpeter, J. A. 1934. *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*. Vol. 55. Transaction Publishers.

Scott Morton, F. M. 2000. Barriers to entry, brand advertising, and generic entry in the US pharmaceutical industry. *International Journal of Industrial Organization*, 18(7): 1085-1104.

Sitkin, S. B. 1992. Learning through failure - the strategy of small losses. *Research in Organizational Behavior*, 14: 231-266.

Skinner, B. 1954. *The science of learning and the art of teaching*. Cambridge, Mass, USA.

Somaya, D., Williamson, I. O., & Zhang, X. 2007. Combining patent law expertise with R&D for patenting performance. *Organization Science,* 18(6): 922-937.

Sorensen J. B., & Stuart T. E. 2000. Aging, obsolescence, and organizational innovation. *Administrative Science Quarterly,* 45(1): 81-112.

Staw, B. M. 1976. Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, 16: 27-44.

Staw, B. M., Sandelands, L. E., & Dutton, J. E. 1981. Threat rigidity effects in organizational behavior: A multilevel analysis. *Administrative science quarterly*, 501-524.

Stern, S. 2004. Do scientists pay to be scientists? *Management Science*, 50(6): 835-853.

Stock, J. H., & Yogo, M. 2005. Testing for weak instruments in linear IV regression. D. W. K. Andrews, J. H. Stock, eds. *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*. Cambridge, UK: Cambridge University Press.

Terwiesch, C., & Xu, Y. 2008. Innovation contests, open innovation, and multiagent problem solving. *Management science,* 54(9): 1529-1543.

Terwiesch, C., & Ulrich, C. 2009. *Innovation Tournaments: Creating and Selecting Exceptional Opportunities.* Boston, MA: Harvard Business School Press.

Thomke, S. 2003. *Experimentation matters: unlocking the potential of new technologies for innovation*. Harvard Business Press.

Thomke, S., & 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. *The Rand Journal of Economics,* 172-187.

Uotila, J., Maula, M., Keil, T., & Zahra, S. A. 2009. Exploration, exploitation, and financial performance: analysis of S&P 500 corporations. *Strategic Management Journal*, 30(2): 221-231.

Van de Ven, A. H. 1986. Central problems in the management of innovation. *Management science,* 32(5): 590-607.

Ward, M. R. 1992. Drug approval overregulation. *Regulation*, 15: 47.

Wildavsky, A. 1988*. Searching for safety*. New Brunswick, NJ: Transaction Books.

Yelle, L. E. 1979. The learning curve: Historical review and comprehensive survey. *Decision Science,* 10: 302-228.

FIGURES AND TABLES

FIGURE 1

Number of Expired Patents for Firms Included in this Study between 1985 and 2002

FIGURE 2
Sample Data Point for Abbott in Year 1998 and Description of Variables

Failures after 4 years: Number of patents which expired after 4 years per firm per year = **14**

Failures (total): Total number of patents expired per firm per year = **40**

Failures after 8 years: Number of patents which expired after 8 years per firm per year = **18**

R&D output (log): log of number of patents that did not prematurely expire for a firm in a given year = **4.52**

Failures after 12 years: Number of patents which expired after 12 years per firm per year = **8**

Independent Variables

Dependent Variables

Unit of Analysis: *Firm-Year*
Firm: **Abbott**
Year: **1998**

Importance of failures: Number of citations to the expired patents per firm per year = **541**

R&D output quality (log): log of number of citations to patents that never expired for a firm in a given year = **6.27**

Control Variables

Number of countries: Number of countries which were origin of patents per firm per year (moving average - 3 years) = **4**

Number of alliances: Number of alliances made by a firm in a given year (moving average - 3 years) = **16**

R&D unit size: Number of scientists in a firm in a given year (moving average - 3 years) = **678**

Technological focus: Herfindahl index for the technological classes of discontinued patents in a firm in a given year (moving average - 3 years) = **0.09**

Number of litigations: Number of lawsuits at both the state and federal level against a firm in a given year (moving average - 3 years) = **10**

TABLE 1

List of 30 Major Pharmaceutical Firms from the Sample of 97 Firms

Abbott	Bristol-Myers Squibb	Monsanto
Ajinomoto	Chiron	Novartis
Allergan	Eli Lilly	Novo Nordisk
American Cyanamid	Fujisawa Pharmaceutical	Pfizer
Amgen	Genentech	Pharmacia Ab
Astra	Genzyme	Roche
Aventis	Glaxo Wellcome	Schering
Baxter	Janssen Pharmaceutical	Shionogi
Bayer	Mallinckrodt Medical	Smithkline Beecham
Boehringer Ingelheim	Merck	Zeneca

TABLE 2

Descriptive Statistics and Partial Correlations

	Variable	Mean	s. d.	Min	Max	1	2	3	4	5	6	7	8	9	10	11	12
1.	R&D output (log)	3.31	1.55	0.00	7.09	1											
2.	R&D output quality (log)	4.36	2.01	0.00	9.46	-0.00	1										
3.	Failures (total)	45.67	77.05	1.00	1076	0.08	0.01	1									
4.	Failures after 4 years	16.37	25.16	1.00	180	-0.11	-0.02	0.87	1								
5.	Failures after 8 years	22.17	37.16	1.00	544	-0.11	-0.05	0.77	0.36	1							
6.	Failures after 12 years	17.65	29.37	1.00	368	0.12	0.06	0.57	-0.01	0.68	1						
7.	Importance of failures	303	687	0.00	8428	-0.16	0.43	-0.10	-0.07	0.32	0.64	1					
8.	Number of alliances	2.82	4.80	0.00	40	-0.07	0.10	0.05	0.07	0.01	-0.18	-0.00	1				
9.	R&D unit size	248	389	1.67	2835	0.05	-0.11	0.44	0.46	0.19	-0.05	0.33	0.19	1			
10.	Number of countries	3.56	2.67	1.00	13.67	0.41	0.29	0.26	0.22	-0.06	0.02	0.08	0.30	0.07	1		
11.	Technological focus	0.24	0.22	0.03	1.00	-0.13	-0.09	-0.05	-0.05	-0.05	-0.16	-0.12	-0.01	-0.09	-0.04	1	
12.	Number of litigations	4.41	7.15	0	52	0.03	0.19	0.04	-0.04	0.03	0.20	0.15	0.20	0.01	0.17	0.37	1

TABLE 3
Arellano-Bond Model Estimates for R&D Output [a, b]

Variable	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6		Model 7	
R&D output$_{t-1}$ (log)	0.48 ***	(0.06)	0.54 ***	(0.06)	0.78***	(0.09)	0.44***	(0.05)	0.68***	(0.08)	0.80***	(0.10)	0.72***	(0.10)
Failures (total)			-0.005 ***	(0.00)	-0.002*	(0.000)								
Importance of failures					-0.003***	(0.000)	-0.003***	(0.000)	-0.002**	(0.00)	-0.002**	(0.000)	-0.003***	(0.000)
Failures after 4 years							-0.003	(0.02)					-0.004	(0.003)
Failures after 8 years									-0.007***	(0.00)			-0.01***	(0.00)
Failures after 12 years											0.01	(0.004)	0.01	(0.01)
Number of alliances	-0.10 ***	(0.01)	-0.06***	(0.01)	-0.04***	(0.01)	-0.04***	(0.01)	-0.06***	(0.01)	-0.06***	(0.01)	-0.05***	(0.01)
R&D unit size	0.002***	(0.00)	0.001***	(0.00)	0.001***	(0.00)	0.001***	(0.00)	0.001***	(0.00)	0.001***	(0.00)	0.001***	(0.00)
Number of countries	0.22***	(0.03)	0.23***	(0.03)	0.24***	(0.03)	0.28***	(0.03)	0.25***	(0.03)	0.27***	(0.04)	0.30***	(0.04)
Technological focus	0.24$^+$	(0.12)	-0.04	(0.16)	-0.02	(0.18)	0.06	(0.15)	-0.12	(0.20)	-0.15	(0.39)	-2.10**	(0.65)
Number of litigations	-0.01	(0.00)	-0.00	(0.00)	-0.00	(0.00)	-0.00	(0.00)	0.00	(0.00)	-0.00	(0.01)	-0.00	(0.01)
Constant	-0.32 *	(0.15)	-0.26	(0.17)	-0.44$^+$	(0.23)	-0.40**	(0.16)	-0.18	(0.22)	-1.05***	(0.30)	-1.44***	(0.33)
Number of Instruments	239		269		126		300		143		103		115	
Number of Observations	1003		859		849		868		753		498		447	
Number of Groups	90		83		81		81		81		74		71	
Wald χ^2	1177		1138		973		1210		1018		835		876	
χ^2_p	0.0000		0.0000		0.0000		0.0000		0.0000		0.0000		0.0000	

[a] Numbers of observations are different across different models because of the difference in the number of observations for expired patents after 4, 8, and 12 years. For example, patents that expired after 4 years first appeared in year 1985 whereas patents that expired after 8 years first made it in our dataset in year 1989, leading to lower number of observations. Following the same logic, number of observations in models with the variable 'failures after 12 years' is even fewer.

[b] Standard errors in parentheses.

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 4
Arellano-Bond Model Estimates for R&D Output Quality [a, b]

Variable	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6		Model 7	
R&D output quality$_{t-1}$ (log)	0.60***	(0.03)	0.44***	(0.04)	0.36***	(0.04)	0.37***	(0.04)	0.59***	(0.04)	0.25 ***	(0.05)	0.20***	(0.05)
Failures (total)			0.007***	(0.00)	0.003**	(0.00)								
Importance of failures					0.005***	(0.00)	0.006***	(0.00)	0.002**	(0.000)	0.004***	(0.00)	0.004***	(0.00)
Failures after 4 years							0.01**	(0.00)					0.004^{+}	(0.002)
Failures after 8 years									0.005*	(0.002)			-0.002	(0.002)
Failures after 12 years											-0.007**	(0.00)	-0.01**	(0.002)
Number of alliances	0.05***	(0.01)	0.05***	(0.01)	0.04**	(0.01)	0.04**	(0.01)	0.04***	(0.01)	0.01	(0.01)	0.02^{+}	(0.01)
R&D unit size	0.00	(0.00)	0.00	(0.00)	0.001***	(0.00)	0.001***	(0.00)	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)
Number of countries	0.01	(0.03)	-0.00	(0.03)	-0.03	(0.03)	-0.04	(0.03)	0.00	(0.02)	-0.03	(0.03)	-0.04	(0.03)
Technological focus	-0.52*	(0.20)	-0.52*	(0.24)	-0.59**	(0.22)	-0.61**	(0.21)	-0.16	(0.17)	-0.42	(0.30)	-0.18*	(0.49)
Number of litigations	0.01^{+}	(0.00)	0.01^{+}	(0.00)	0.01^{+}	(0.00)	0.01^{+}	(0.00)	0.01	(0.01)	0.00	(0.00)	0.01^{+}	(0.00)
Constant	1.76***	(0.19)	2.29***	(0.20)	2.69***	(0.20)	2.59***	(0.21)	2.11***	(0.25)	4.28***	(0.31)	4.58***	(0.32)
Number of Instruments	172		188		204		204		103		211		270	
Number of Observations	936		813		805		822		778		500		451	
Number of Groups	86		82		80		81		82		73		71	
Wald χ^2	780		853		943		936		915		178		180	
χ_p^2	0.0000		0.0000		0.0000		0.0000		0.0000		0.0000		0.0000	

[a] Numbers of observations are different across different models because of the difference in the number of observations for expired patents after 4, 8, and 12 years. For example patents that expired after 4 years first appeared in year 1985 whereas patents that expired after 8 years first made it in our dataset in year 1989, leading to lower number of observations. Following the same logic, number of observations in models with the variable 'failures after 12 years' is even fewer.

[b] Standard errors in parentheses.

$+ p < 0.1$, $* p < 0.05$, $** p < 0.01$, $*** p < 0.00$