Failing to learn and learning to fail (intelligently):

How great organizations put failure to work to improve and innovate

Mark D. Cannon
Vanderbilt University
mark.d.cannon@vanderbilt.edu

Amy C. Edmondson
Harvard University
aedmondson@hbs.edu

February 5, 2004

Introduction

The idea that people and the organizations in which they work should learn from failure has great popular support and even seems obvious. However, organizations that systematically and effectively learn from failure are very rare. This paper provides insight into what makes learning from failure so difficult to put into practice – that is, we address the question of why organizations *fail to learn* from failure. We identify pernicious barriers embedded in both technical and social systems that make collective learning processes unusual in organizations, and present recommendations for what managers can do to overcome these barriers.

We also address the question of how organizations can *learn to fail* intelligently – as a deliberate strategy to promote innovation and improvement. To do this, we clarify the key processes through which organizations can learn from failure – a set of activities that includes developing the ability to *generate* failures deliberately—through experimentation. Deliberate experimentation is thus seen as a means of accelerating an organization's learning. We illustrate the importance of understanding and practicing three specific learning processes, and describe ways of anticipating and removing barriers to engaging in them. These guidelines are offered as keys to unlocking an organization's ability to learn from failure.

Over the past decade or so, our research has revealed impediments to organizational learning from failure on multiple levels of analysis. Mark Cannon has investigated individuals' psychological responses to their own failures, demonstrating the aversive emotions people experience and how they inhibit learning. Amy Edmondson has identified group and organizational factors that limit collective learning from failure – in teams and organizations. We have worked together for a number of years to conceptualize and develop recommendations for how to enable organizational learning from failure, drawing from our own and others' research.

This article reports on the results of that work. First, we provide insights about what makes organizational learning from failure difficult – paying particular attention to what we see as a lack of understanding of the essential processes involved in learning from failure in a complex system such as a corporation, hospital, university, or government agency. Drawing from field research, we suggest that organizational learning from failure is feasible but involves skillful management of three distinct but interrelated processes: identifying failure, analyzing and discussing failure, and experimentation. Managed skillfully, these processes help managers take advantage of the lessons that failures offer, which are ignored or suppressed in most organizations.

Second, we argue that most managers underestimate the power of both *technical* and *social* barriers to organizational learning from failure. This article explains and illustrates these barriers and suggests ways of overcoming them. We use these ideas to develop a framework to illuminate the key drivers and activities through which successful organizational learning from failure can occur.

Organizational learning from failure

Large and well-publicized organizational failures such as the *Columbia* and *Challenger* Shuttle tragedies, the Colorado South Canyon Firefighter deaths, the drug error that killed a Boston Globe correspondent at Boston's Dana Farber Hospital, and the Parmelat and Enron accounting scandals argue for the necessity of learning from failure. Recognizing the need to understand and learn from consequential incidents such as these, executives or regulators often establish task forces or investigative bodies to uncover and communicate the causes and lessons of highly visible failures. Many of these efforts can be said to be "too little, too late" for the goal of organizational learning from failure. The multiple causes of large failures are usually deeply

embedded in the organizations where the failures occurred, have been ignored or taken for granted for years, and rarely are simple to correct.[1]

This article argues that an important part of most organizations' inability to learn from failure is due to a lack of attention to small, everyday organizational failures, especially as compared to the investigative commissions or formal "after action reviews" triggered by large catastrophic failures. Small failures are often the "early warning signs" which, if detected and addressed, may be the key to avoiding catastrophic failure in the future.[2]

Our research in organizational contexts ranging from the hospital operating room to the corporate board room suggests that an intelligent process of organizational learning from failure requires proactively identifying and learning from *small* failures, even those that appear at the time they occur or are discovered to be insignificant minor mistakes or problems, in addition to the more common attention to review of large failures. We find that when small failures are not identified widely, discussed and analyzed, it is very difficult for larger failures to be prevented. [3]

Organizational failure defined

Failure, in organizations and elsewhere, is deviation from expected and desired results. This includes both avoidable errors and unavoidable negative outcomes of experiments and risk taking.[4] We define failure broadly to include both large and small failures in domains ranging from the technical (a flaw in the design of a new machine) to the interpersonal (such as a failure to give feedback to an employee with a performance problem). Drawing from our own and others' research, we suggest that an organization's ability to learn from failure is best measured by how it deals with a range of large and small outcomes that deviate from expected results rather than focusing exclusively on how it handles major disasters.

Barriers to Learning From Failure

Learning from failure is a hallmark of innovative companies[5] but, as noted above, is more common in exhortation than in practice. Most organizations do a poor job of learning from failures, large and small. In our research, we found that even companies that had invested significant money and effort into becoming "learning organizations" (with the ability to learn from failure) struggled when it came to the day-to-day mindset and activities of learning from failure.[6]

Instead, fundamental attributes of organizations usually conspire to make a rational process of diagnosing and correcting factors that give rise to failures difficult to execute. A prominent tradition in managerial research examines the importance of considering both social and technical attributes of organizations as systems.[7] Recognizing that organizations are simultaneously social systems and technical systems, management researchers have long considered the need to examine how features of tasks and technologies as well as social, psychological and structural features shape the outcomes of organizational systems. We draw from this basic framework to organize the barriers to learning from failure embedded in both technical and social systems in organizations, before describing specific learning processes through which these fundamental barriers can be overcome through diligent action.

Barriers embedded in technical systems

As discoveries in the learning sciences and the history of science illustrate, limitations in human intuition and "sense-making" can lead people to draw false conclusions, inhibiting both individual and collective learning. Technical barriers are present when individuals lack the basic scientific "know how" to effectively engage in the analytical and scientific aspects of learning from failure, or when technologies are particularly difficult to understand and diagnose.

Task design can obscure or make failures transparent (e.g., excess *work in process* inventory slows the discovery of manufacturing process errors, as further discussed below). Similarly, cross-disciplinary work, if not managed or conducted by people with broad understanding of the issues and the potential interrelationships across fields can make it difficult for non experts within a field to understand the work well enough to identify things that have gone wrong.

Lastly, technical or skill based barriers include the inability to understand or engage in the following aspects of rigorous inquiry: the scientific method, problem diagnosis, experimental design, rigorous quantitative or qualitative analysis, statistical process controls, and statistical analysis.

Barriers embedded in social systems

Organizational structures and senior management behaviors can explicitly or implicitly discourage the behaviors involved in identifying and analyzing failure and in experimentation[8]. Organizations often punish or at least fail to reward these behaviors through raises, bonuses, promotions, career opportunities, privileges and other formal and informal sanctions. In addition, organizational structures, policies and procedures can encourage or inhibit these behaviors.

Encouragement is found in organizations with policies and procedures to encourage and fund experimentation. Other smart organizations provide structures or forums in which employees can analyze and discuss the results of these experiments. More typically, however, organizational incentives, policies, procedures, and structures do more to discourage these processes than to promote them. This contributes to an organizational culture in which even intelligent failure is discouraged. Employees are quick to figure out what behaviors are rewarded and which are not and act accordingly. When learning from failure means putting oneself at risk,

many employees decide not to bother. Similarly, the ability to facilitate an effective discussion in which parties with conflicting perspectives search to understand complex failures cannot be taken for granted in most organizations, yet it is critical to learning from failure.

The properties of social systems are derived from psychological and organizational factors. Social systems cannot be programmed to act in predicted, desired ways; they are influenced by complex interactions of human emotions and cognitions. Most social systems have properties that inhibit coping effectively, calmly, and rationally with failure.

Social systems are made up of individual, psychologically complex, human beings. And, most individuals experience strong negative feelings in response to their own failures, such that acknowledging and examining them can be a very bitter pill. We have an instinctive tendency to deny, distort, ignore, or disassociate ourselves from our own failures.[9] The fact that someone has achieved an executive or leadership position in an organization does not mean that they have learned to honestly and openly confront personal failures or limitations.[10]

Recent evidence of this comes from Finkelstein's extensive, six-year investigation of major failures that included almost 200 interviews and detailed analysis of major failures at over 50 companies.[11] After analyzing this breadth of data, he explained:

> Ironically enough, the higher people are in the management hierarchy, the more they tend to supplement their perfectionism with blanket excuses, with CEOs usually being the worst of all. For example, in one organization we studied, the CEO spent the entire forty-five-minute interview explaining all the reasons why others were to blame for the calamity that his company. Regulators, customers, the government, and even other executives within the firm--all were responsible. No mention was made, however, of personal culpability. (p. 181-2)

Thus, social systems present barriers to learning from failure that have intrapsychic, interpersonal, and instrumental elements.

Intrapsychically, honest acknowledgement of one's failures is not only unpleasant it also can strike a blow to one's self-esteem, self-image, and identity. The fundamentally human desire

to have positive self-regard can be hindered by an honest evaluation of one's failures. Likewise, an objective assessment of our failures and shortcomings can strike a blow to our positive self-image and highly valued aspects of our own identity.[12] Moreover, acknowledging one's failure and limitations may undermine the confidence and motivation that one needs for peak performance. Many psychologists have argued that unrealistically positive self-perceptions can actually facilitate motivation, determination, persistence, and energy level. Thus, ignoring our failures may not only protect us from the unpleasantness of the emotion, but may also help us avoid discouragement and maintain a high level of motivation.[13]

Being held in high regard by other people, especially those with whom one spends a lot of time, is a strong fundamental human desire. Open acknowledgement of one's mistakes and failures is thus interpersonally unappealing. Even though others may learn from and appreciate others' disclosures of failure does not guarantee that they will maintain a positive impression of the individual who made the mistake. People thus tend to avoid situations that might lead to public embarrassment or private derision.

Instrumentally, organizational rewards may depend upon appearing not to make failures. Individual employees may incur tangible costs if their actions create unfavorable impressions on people who influence decisions regarding promotions, raises, or desirable project assignments.

These potent barriers are all but hard wired into social systems. They greatly reduce the ability of most organizations to learn from failure. Thus, not only do few organizations systematically capture failure's lessons, most managers lack a clear understanding of what a proactive process of learning from failure looks like.

Without a clear model of what it takes to learn from failure, organizations are at a disadvantage. In full acknowledgement of the magnitude of the challenge, we suggest that breaking the process down into more tangible component activities greatly enhances the

likelihood of gleaning failures' lessons. In the next section, we identify and explain three distinct processes through which effective organizations can proactively learn from failure.

The Processes of Organizational Learning from Failure

Learning from failure is a process. Considering the component activities through which this process occurs is an initial step in making it happen. We offer three basic organizational activities through which collective learning from failure occurs: (1) identifying failure, (2) analyzing failure, and (3) deliberate experimentation. They are presented in order of increasing challenge, both organizationally and in terms of the specific skills required. In this section, we describe these activities in some detail and provide illustrations of how some organizations have successfully overcome the pervasive technical and social barriers associated with enacting them.

Identifying failure

Proactive and timely identification of failures is an obvious and essential first step in the process of learning from them. One of the revolutions in manufacturing, the drive to reduce inventory to the lowest possible levels, was stimulated as much by the desire to make problems and errors quickly visible as by the desire to avoid other, inventory-associated costs. Surfacing errors before they are compounded, incorporated into larger systems, or made irrevocable is an essential step in achieving high quality.[14]

Indeed, one of the tragedies in organizational learning is that catastrophic failures are often preceded by smaller failures that were not identified as failures worth examination and learning. In fact, these small failures are often the key "early warning signs" that could provide a wake up call needed to avert disaster down the road. Social system barriers are often the key driver of this kind of problem. Rather than acknowledge and address a small failure, individuals

have a tendency to deny the failure, distort the reality of the failure, or cover it up, and groups and organizations have the tendency to suppress awareness of failures.

For example, Finkelstein[15] mentions Jill Barad at Mattel as an illustration of failing to acknowledge and learn from mistakes in a timely manner. In Mattel's ill-fated acquisition of the Learning Company, Barad first overlooked the problems that the organization was having prior to the acquisition. An opportunity to acknowledge the failure came when the third quarter 1999 earnings turned out to be a loss of $105 million, rather a profit of $50 million as she expected. However, rather than address the failure, she remained optimistic and predicted significant profits for the next quarter; instead, there was a loss of $184 million. Once again, rather than acknowledge the failure and learn from it, she repeated the same mistake for the next two quarters as well, thus making the same mistake for a total of four quarters.

By contrast, the CEO of a mechanical contractor recognized the value of exposing failure and publicizing in order to help employees learn from each other and not repeat the same mistake. The CEO:

> pulled a $450 "Mistake" out of the company's dumpster, mounted it on a plaque, and named it the "no-nuts award"--for the missing parts. A presentation ceremony followed at the company barbecue. "You can bet no one makes *that* mistake anymore," the CEO says. "The winner, who was initially embarrassed, now takes pride in the fact that his mistake has saved this company a lot of money."[16]

Overcoming the psychological barriers to identifying failure requires courage to face the unpleasant truth. The key organizational barrier to identifying failure has mostly to do with overcoming the inaccessibility of data that would be necessary to identify failures. To overcome this barrier, organizational leaders must take the initiative to develop systems and procedures that make available the data necessary to identify and learn from failure.

As an example of this in the medical arena, Dr. Kim Adcock of Kaiser Permanente proactively collected and organized data to identify failure of physicians who read mamograms.

Due to the inherent difficulties in reading mammograms accurately, the medical profession has come to expect a 10-15% error rate even among the expert readers. Consequently, discovering that a reader has missed one or even several tumors does not necessarily say anything about that reader's diagnostic ability and my not provide much incentive for learning from failure. By contrast, when Dr. Adcock became radiology chief at Kaiser Permanente Colorado, he utilized the longitudinal data available in the HMO's records to proactively identify failure and produce detailed, systematic feedback including bar charts and graphs for each individual x-ray reader.[17] For the first time, each reader could learn whether he or she was falling near or outside of the acceptable range of errors. Dr. Adcock also provided readers with the opportunity to return to the misread x-rays so they could investigate why they missed a particular tumor and learn not make the same mistake again.

On a larger scale, Electricitie De France, which operates 57 nuclear power plants, provides an example of identifying and learning from potential mistakes.[18] The organization tracks each plant for anything even slightly out of the out of the ordinary and has a policy of quickly investigating and publicly reporting any anomalies throughout the entire system so that the whole system can learn.

Feedback seeking is also an effective way of identifying many types of failures. Feedback from customers, employees and other sources can expose failures, including communication breakdowns as well as failure to meet goals or satisfy customer requirements. Proactively seeking feedback from customers outside the company may be necessary in order for manufacturers and service providers to identify and address failures in a timely manner.

For example, only 5 to 10 percent of dissatisfied customers choose to complain following service failure; instead, most simply switch providers.[19] This is one of the reasons service companies fail to learn from failures and therefore lose customers. Service management

researchers Tax and Brown[20] cite General Electric and United Parcel Service as two organizations that proactively seek data that will help them identify failures. G.E. places an 800 number directly on each of its products and encourages customers to inform the company of any problems. GE has an Answer Center that is open twenty-four hours a day, 365 days a year, receiving approximately 3 million customer calls a year. UPS provides an example of how to seek feedback from within the company. The company has built in a half hour per week to the schedule of each of its drivers for receiving their feedback and answering questions. These simple techniques exemplify methods of identifying failure in a timely way so that the organizations can learn, respond quickly, and retain customers.

At the same time, this is not easy. Employees, consciously and not, often fail to exploit and may even actively avoid opportunities to expose and learn about their failures. Effective identification of failure entails exposing failures as early as possible, to allow learning in an efficient and cost effective way. This often requires a proactive effort on the part of managers to surface available data on failures and use it in a way that promotes learning.

A recent tragic example of the consequences of delayed and minimized identification of failure is found in the *Columbia* disaster. As discussed in the Accident Investigation Board's report, NASA managers spent 17 days downplaying the possibility that foam strikes on the left side of the shuttle represented a serious problem – a true failure – and so did not view the events as a trigger for conducting detailed analyses of the situation.[21] Instead the strikes were deemed ordinary events, within the boundaries of past experience, an interpretation that would later seem absurd given the large magnitude of the debris. The shared belief that there was little they could have done contributed to a lack of proactivity in analysis and exploration of possible remedies. Sadly, post-event analyses have suggested the possibility that fruitful actions could

have been taken had the failure been identified and explored early in this window of opportunity.

Because psychological and organizational factors conspire to reduce failure identification, a fundamental reorientation in which individuals and groups are motivated to engage in the emotionally challenging task of seeking out failures is needed. Obviously, organizations that have a habit of "shooting the messenger" that identifies and reveals a failure will discourage this process.

Creating an environment in which people have an incentive, or at least do not have a disincentive, to identify and reveal failures is the job of leadership.[22] For example, the Chief Operating Officer at Children's Hospital in Minneapolis, Julie Morath, developed a "blameless reporting" system to encourage employees not only to reveal medical errors right away, but also to share additional information that could be used in analyzing causes of the error.[23] Similarly, The US Air Force specifically motivates speaking up early by penalizing pilots for not reporting errors within 24 hours. Errors reported immediately are not penalized; those not reported but discovered later anyway are treated severely. Finally, in contrast to the potent psychological and organizational barriers that discourage identifying failure, the technical barriers to identifying failure are fairly minor.

Analyzing failure

It hardly needs to be said that organizations cannot learn from failures if people do not discuss and analyze them. Yet, this remains an important insight. The potential learning cannot be realized unless thoughtful analysis and discussion of failure occurs. For example, for the X-ray readers mentioned above, it was not enough just to know that one is making more than the acceptable number or errors. Readers need to analyze the x-rays that were mistakenly read to look for patterns in their reading errors or spot particular weaknesses that could be corrected.

On a larger scale, the US Army is known for conducting After Action Reviews that enable participants to analyze, discuss and learn from both the successes and failures of a variety of military initiatives. Similarly, hospitals use "Morbidity and Mortality" (M&M) conferences (in which physicians convene to discuss significant mistakes or unexpected patient deaths) as a forum for identifying, discussing and learning from failures. Many of these vehicles for analysis only address substantial failures, however, rather than identifying and learning from smaller ones.

An example of effective analysis of failure is found in the meticulous and painstaking analysis that goes into understanding the crash of an airliner. Hundreds of hours may go into gathering and analyzing data to sort out exactly what happened and what can be learned. Compare this kind of analysis to what takes place in most organizations after a failure.

As noted above, social systems tend to discourage this kind of analysis. First, individuals experience negative emotions when examining their own failures and this can chip away at self-confidence and self-esteem. Most people prefer to put past mistakes behind them rather than to revisit and unpack them for greater understanding. Second, conducting an analysis of a failure requires a spirit of inquiry and openness, patience, and a tolerance for ambiguity. However, most people admire and are rewarded for decisiveness, efficiency and action rather than for deep reflection and painstaking analysis.

Third, psychologists have spent decades documenting heuristics and psychological biases and errors that reduce the accuracy of human perception, sense making, estimation, and attribution.[24] These can also hinder the human ability to analyze failure effectively. People tend to be more comfortable attending to evidence that enables them to believe what they want to believe, denying responsibility for failures, attributing the problem to others or the system, and like to move on to something more pleasant. Rigorous analysis of failure requires that people at

least temporarily put aside these tendencies to explore unpleasant truths and take personal responsibility.

To illustrate, we have observed a number of failed consulting relationships in our field research in which the consultants simply blamed the failure to the client, concluding that the client was not really committed to change, or that the client was defensive or difficult.[25] By contrast, a few, highly learning oriented consultants were able to engage in discussion and analysis that involved raising questions about how they themselves contributed to the problem. In these analytic sessions, the consultants raised questions such as "Are there things I said or did that contributed to the defensiveness of the client?" or "Was my presentation of ideas and arguments clear and persuasive?" Did my analysis fall short in some way that led the client to have legitimate doubts?" Raising these questions requires both profound personal curiosity to learn what the answers might be and increases the chances that the consultants will learn something useful from the failed relationship. However, blaming the client is much more comfortable and efficient.

Recent research in the hospital setting shows that health care organizations typically fail to analyze or make changes even when people are well aware of failures.[26] Whether medical errors or simply problems in the work process, few hospitals organizations dig deeply enough to understand and capture the potential learning from failures. Processes, resources, and incentives to bring multiple perspectives and multiple minds together to carefully analyze what went wrong and how to prevent the occurrence of similar failures in the future are lacking in most organizations.

Thus, formal processes or forums for discussing, analyzing and applying the lessons of failure elsewhere in the organization are needed to ensure that effective analysis and learning from failure occurs. Such groups are most effective when people have technical skills and

expertise for analysis and diverse views and perspectives allowing them to brainstorm and explore different interpretations of a failure's causes and consequences. Because this usually involves the potential for conflict and for differences to get out of hand and become personal, people skilled in interpersonal or group process, or the use of expert facilitators, can help to keep the process productive.

Next, skills for managing a group process of analyzing a failure with a spirit of inquiry and with sufficient understanding of the scientific method is an essential input to learning from failure as an organization.[27] Without rigorous analysis and deep probing individuals tend to prematurely leaping to unfounded conclusions and misunderstand complicated problems. Some understanding of system dynamics, the ability to see patterns, statistical process controls, and group dynamics can be very helpful here.

To illustrate how this works, Children's Hospital's Morath implemented supports for effective analysis of failures, both large and small.[28] First, she strengthened her technical knowledge of how to probe more deeply into the causes of failure in hospitals by attending the Executive Sessions on Medical Errors and Patient Safety and Harvard University and through a variety of other educational opportunities and experiences. She learned that rather than being the fault of a single individual, medical errors tend to be embedded in complex interdependent systems and have multiple roots. In addition to strengthening her technical knowledge, she overcame organizational barriers by making structural changes within the organization to create a context in which failure could be identified, analyzed and learned from. To create a forum for learning from failure at Children's Hospital, she developed a Patient Safety Steering Committee (PSSC). Not only was the PSSC proactive in seeking to identify failures, it ensured that all failures were subject to analysis so that learning could take place. For example, the PSSC determined that "Focused Event Studies" would be conducted not only after serious medical

accidents, but even after much smaller scale errors or "near misses." These "Focused Event Studies" were forums designed explicitly for the purpose of learning from the mistakes by probing deeply into their causes.

In addition, cross-functional teams, known as "safety action teams" spontaneously formed in certain clinical areas to better understand how failures occurred, to proactively improve medical safety. One clinical group developed something they called a "Good Catch Log" to record information that might be useful in better understanding and reducing medical errors. Other teams in the hospital quickly followed their example, finding the idea compelling and practical.

To supplement or help build an organization's ability to analyze its own failures, outside sources of technical assistance in analyzing failure can be engaged. For example, Frederick Reichheld, a Bain and Company partner who has studied failures in the form of customer and employee defections[29] demonstrates a deep, probing analysis of failure. He illustrated, noting that the fact that most of the customers who defected from a particular bank gave "interest rates" as the reason for switching banks would seem to suggest that their original bank's interest rates were not competitive. Additional investigation theb demonstrated that there were no significant differences in interest rates across the banks. Careful probing through interviews indicated that many customers defected because they were irritated by the fact that they had been aggressively solicited for a bank-provided credit card and then subsequently turned them down for the card. A superficial analysis of customer defection would have led to the conclusion that the bank's interest rates were not competitive. A deeper analysis led to an alternative conclusion, that the bank's marketing department needed to do a better job of screening in advance the customers to whom it promoted bank-provided credit cards.

Another example of the importance of analysis relates to employee turnover.[30] A company's managers became concerned when they observed high turnover among sales people and conducted an investigation. Many of the employees gave "working too many hours" as the reason for their defection. Initially, it looked like the turnover may not have been such a bad thing—after all who needs employees who are not committed to working hard? However, after additional data gathering management learned that many of the employees who quit were among their most successful salespeople. Further data collection demonstrated that the successful salespeople who quit had found jobs that required, on average, 20 percent fewer hours. Once again, gaining a true understanding of the situation required deeper probing and analysis.

In addition to systematic analysis, discussing failures is important for several reasons. First, discussion provides an opportunity for others who may not have been directly involved in the failure to learn from it. Second, others may bring new perspectives and insights that deepen the analysis and help to counteract self-serving biases that may color the perceptions of those most directly involved in the failure. After experiencing failure, people typically attribute too much blame to other people and forces beyond our control. If this tendency goes unchecked, it reduces an organization's ability to dig out the key learning that could come from the experience.

Involving individuals who were not directly tied to a failure in the analysis can improve the quality of the discussion. Such individuals are less emotionally tied to a particular sense-making regarding the event, often have a different perspective to bring to the analysis, and can ask probing questions.

Lastly, the learning value that might result from analyzing and discussing simple mistakes is often overlooked. In fact, many scientific discoveries have resulted from those who were attentive to simple mistakes in the lab. For example, Peter Drucker notes[31] that researchers in one of the early German polymer labs occasionally made the mistake of leaving a Bunsen burner

lit over the weekend. Upon discovering this mistake on Monday mornings, the chemists simply discarded the overcooked results and went on with their day. Ten years later, a chemist in a polymer lab at DuPont made the same mistake. However, rather than simply discarding the mistake, the Dupont chemist gave some analysis to the result, discovered that the fibers had congealed, and this was discovery provided the first step toward the invention of nylon. With similar attention to the minor failure in the German lab, the Germans might have had a decade head start in nylon and could have dominated the market.

Experimentation

The third process for learning from failure is perhaps both the least obvious and the most provocative. A handful of exceptional organizations not only seek to identify and analyze failures, they seek to generate them – for the express purpose of learning and innovating. This means they devote some portion of their collective energy to deliberate experimentation – trying new things out to find out what works and what doesn't. Through deliberate experimentation, organizations can generate novel solutions to problems and new ideas for products, services and innovations. In this way, they put new idea to the test – in a controlled context.

Experiments are understood to have uncertain outcomes and to be designed for learning. Despite the increased rate of failure that accompanies deliberate experimentation, organizations that experiment effectively actually are likely to be more innovative, productive, and successful than those that do not take such risks.[32] Similarly, other research has shown the research and development teams that experimented frequently performed better than other teams.[33]

Social systems can make deliberate experimentation difficult, because most organizations reward success, not failure. Knowingly setting out to generate some failures alongside some successes, although reasonable, is difficult when failures are stigmatized. Conducting

experiments also requires acknowledging that the status quo is imperfect, requiring the possibility of change. A psychological bias known as the confirmation trap, meaning that people tend to seek to confirm their views rather than to learn why they might be wrong, makes this particularly difficult.[34] Deliberate experimentation requires that people not just assume their views are correct, but actually put their ideas to the test and design (even very informal, small) experiments in which their views can be disconfirmed.

A good example of the ability to overcome these psychological barriers is provided by the influential, award-winning design firm, IDEO. They communicate this perspective with slogans such as: "Fail often in order to succeed sooner." And "Enlightened trial-and-error succeeds over the planning of the lone genius."[35] These sayings are accompanied by frequent small experiments, and much good humor about the associated failures.[36]

Similarly, PSS/World Medical encourages experimentation in a variety of ways and sometimes even goes so far as to encourage employees to experiment with career moves. PSS/World Medical has a "soft landing" policy means that if an employee tries out a new position, but does not succeed after a good faith effort, the employee can have his or her former job back. This "soft landing" policy is an implicit recognition that experiments have uncertain outcomes and that people will be more willing to experiment if the organization protects their interests.[37]

Technical skills are critical in implementing a deliberate experimentation process. First, since analyzing failures is part of this process, key individuals need skills for analyzing the results of experiments. Additionally, rigor is needed to design experiments that will effectively confirm or disconfirm hypotheses to generate useful learning. Under some conditions, this can be extremely challenging. For example, customer satisfaction at a large resort is affected by many interdependent aspects of the customer's experience. As the resort experiments with different

possible innovations to enhance customer satisfaction, how do they determine what the impact of these different innovations was?

Designing experiments in complex, interdependent systems is challenging even for research experts. In addition to knowledge of experimental design and analysis, people need resources to run experiments in different parts of the organization and to capture the learning.

The 3M corporate has been unusually successful in providing incentives and policies that encourage deliberate experimentation. The company has earned a reputation for successful product innovation by encouraging deliberate experimentation and by cultivating a culture that is tolerant and even rewarding of failures; failures are seen as a necessary step in a larger process of developing successful, innovative products. Apocryphal stories such as Arthur Fry and the failed super-adhesive that spawned the Post-it industry are spread far and wide, both within and outside of the company. Setting goals, such as that of having 25 percent of a division's revenues come from products introduced within the last five years means that divisions must be actively experimenting to develop new products.

Bank of America provides an interesting example of how an organization realized to become more innovative as a service firm would require some deliberate, intelligent experimentation – in this case, visible to the customer. In 2000, Bank of America decided to become an industry leader in innovation and established a program to develop a process and culture of innovation in two-dozen real-life "laboratories."[38] Laboratories, in this case, referred to fully operating banking branches in which new product and service concepts, such as virtual tellers, were being tested by employees.

Senior executives addressed organizational barriers by funding and developing an Innovation & Development team to manage this process. A successful program entailed hiring

individuals with the technical research skills to address a number of complicated questions such as:

> how to gauge success of a concept, how to prioritize which concepts would be tested, how to run several experiments as once, and how to avoid the novelty factor itself form altering the experimental outcome[39]

Successful experiments-determined on the basis of consumer satisfaction or revenue growth-were then recommended for a national rollout.

Senior management strongly supported innovation at these branches, and further recognized that trying out innovative ideas necessarily produced failures along the way (in fact, a projected failure rate of 30% was targeted as one that would indicate sufficient attempts at truly novel ideas). However, employee rewards were primarily based on indices measuring routine performance (such as opening new customer accounts). Employees' personal compensation often suffered when they spent time experimenting with new ideas, or when their experiments failed. As a result, employees were reluctant to try out radical experiments until management made changes to assure that reward systems were aligned with the organization's espoused value of innovation.

Similarly, experimental research in social psychology has demonstrated that espoused goals of increasing innovation through experimentation are not as effective when rewards penalize failures, as when reward and values are aligned with the goal of promoting experimentation.[40] As both field and laboratory examples show, although experimentation is an essential activity underlying innovation, it is both technically and socially challenging to make it happen intelligently. One of the advantages of most forms of experimentation is that failures can take place off line – in dry runs, simulations, and other kinds of practice situations in which the failures are not costly.[41] However, even in these situations, interpersonal fears can lead to reluctance to take risks, limiting the effectiveness of the experiments.[42] Moreover, some

experiments must take place *on line*, or in real settings in which customers interact directly with the failures.

The next section brings the above observations together into a framework for enabling organizational learning from failure. Our basic premise is that, although the barriers to a systematic process of learning from failure in organizations are deep rooted and numerous, by breaking this process down into component activities, with associated recommendations and practices, organizations can slowly but surely improve their track record of learning from their own failures. Further, we suggest that by focusing first on small failures rather than catastrophic ones, organizations and their managers can minimize the inherently threatening nature of failure to gain experience and momentum in this learning process.

Moving Forward: Putting Failure to Work to Innovate and Improve

As noted above, most organizations are not well prepared to learn from small failures. Only after a real crisis, often after ignoring "early warning signs," is active learning from failure common, at which point the learning is too late to prevent the serious harm. The first part of this article offered reasons why a more systematic process of organizational learning from failure is uncommon; we then described the specific activities through which this process can take place, to break down the encompassing and abstract notion of organizational learning from failure into specific, actionable activities. Awareness of these specific activities, however, is insufficient for ensuring that they occur in many organizations.

This section builds on the previous arguments to explore what organizations can do to overcome the barriers – both technical and social – to successfully enact specific activities of learning from failure. Table 1 summarizes this advice, which relates the two types of barriers to the three critical activities for learning from failure to suggest six actionable recommendations.

Insert Table 1 about here

The bottom row of Table 1 presents recommendations for building psychological and organizational capabilities for identifying failure, analyzing failure, and experimentation; the top row is directed at overcoming technical barriers to these activities and emphasizes training, education, and the judicious use of technical expertise.

For overcoming technical barriers, we first recommend helping employees to see that identifying failure requires a proactive and skillful search, that human intuition is often insufficient to extract the key learning from failure, and that intelligent experimental design is a critical tool for innovation and learning. With this basic awareness, employees are better able to recognize when they either need more specialized training themselves or need to engage the assistance of someone else who has such training.

Recommendation 1: Overcoming technical barriers to identifying failure

Organizations are complex systems, which can make small (and sometimes large) failures difficult to detect. Failures, as noted above, are deviations from the expected and desired; if a system has many complex parts and interactions, such deviations can be ambiguous. The *Columbia* shuttle's initial failure exemplifies this phenomenon; it was not clear to those involved until much later that the foam strike should indeed be identified as a failure. Further, people tend to have an erroneous level of confidence that their initial interpretations that nothing is really wrong are correct, also as seen in the *Columbia* example.

Thus, enhancing the ability of individuals to identify (especially small) failures often requires training. For example, training in statistical process controls is useful for identifying failure on an assembly line. Without this training, people are at a disadvantage in discovering whether variation indicates that something is really wrong or whether it is just natural noise in a

process that is actually in control Or, employees in organizations that are complicated and interdependent will benefit from training in systems thinking and scientific analysis. This enhances their ability to identify failure and pinpoint its source – and especially to realize the critical role of small failures in creating large consequences in complex systems.

Fortunately, overcoming technical barriers does not require that all employees have the required technical skill themselves. The judicious use of a few well-placed technical experts and systems thinkers may be enough to trigger more reliable identification of failure. For example, at Children's Hospital in Minneapolis, safety experts were brought in to help the hospital identify latent failures. Another example occurred following the Columbia launch when a group of technical experts in simulation technology from Boeing who were contracted to support the Shuttle programs at NASA tried to apply their expertise to figure out how large a problem the foam strike represented. Unfortunately, organizational factors limited the quality of the dialogue and made these efforts less effective than they might have been. We deal with this type of barrier in recommendations 4-6, below.

Recommendation 2: Overcoming technical barriers to analyzing failure

Most people have a propensity not to recognize that they lack complete information, or that their analysis is not rigorous, and thus leap quickly to questionable conclusions, while remaining confident that they are correct. Extracting failure's lessons is not always straightforward. Figuring out which aspects of a complex situation were contributing factors to something that did not go as expected is can be a complex undertaking. For this reason, we recommend training (for some not necessarily all relevant individuals) in skills and techniques for systematic analysis of complex organizational data. At Children's Hospital, the patient safety effort included considerable effort to ensure appropriate technical skills were in place to gain the most and most appropriate lessons from each mishap – whether large or small. Perhaps one of

the most difficult aspects of analyzing failure pertains to interpersonal dynamics, the focus of recommendation 5.

Recommendation 3: Overcoming technical barriers to effective experimentation

To produce valuable learning, experiments must be designed effectively. Yet, even PhD laboratory researchers with years of experience can struggle to get an experimental design just right. And, most organizational settings have limited ability to isolate variables and reduce noise, makes designing experiments for organizational learning challenging. Yet, at its most basic, designing experiments for learning requires careful thought to what kinds of data will be collected and how results of the experiment will be assessed. For example, the Bank of America examined financial and customer satisfaction metrics of its experimental bank branches. The key is to consider possible outcomes in advance and know how they might be interpreted.

It is not necessary to make all employees experts in experimental methodology; it is more important to know when help is needed from (outside or company) experts with sophisticated skills. Bank of America handled this smartly by developing the Innovation and Development team and staffing it with experts who understood the vulnerabilities associated with conducting research experiments in a real-world setting and how to work around them.

Although technical barriers and thus the above recommendations are important for facilitating organizational learning, barriers due to social systems are more subtle, pervasive, and difficult to address. Even without explicit incentives against failure, many organizations have norms and practices that are unfriendly to experimentation as well as to identifying, analyzing, and learning from failure. The next three recommendations tackle these issues directly.

Recommendation 4: Overcoming social system barriers to identifying failure

To promote timely identification of failure, organizations must avoid "shooting the messenger" and instead put constructive incentives for speaking up in place. That is, people

must feel able to speak up about the failures, both clear and ambiguous, of which they are aware. To do this, leaders need to cultivate an atmosphere of psychological safety to mitigate risks to self-esteem and others impression. Developing psychological safety begins with the leader modeling the desired behaviors.[43] Modeling means that leaders visibly engage in the behaviors that they wish to encourage subordinates and peers to enact.

Leader modeling serves two significant purposes. First, to communicate expected and appropriate behavior, it is important for leaders to "walk the talk." Second, leader modeling can help subordinates learn how to enact these processes. Because these behaviors may be unfamiliar in many organizations, having a model to observe can be very helpful in facilitating subordinate learning. Leaders can model effectively by: generating new ideas, disclosing and analyzing failure, inviting constructive criticism and alternative explanations, and capturing and utilizing learning.

Finally, psychological safety cannot be implemented by top down command; it is created instead, work group by work group, through attitudes and activities of local managers, supervisors and peers; the development of managerial coaching skills is one way to help build this type of learning environment.[44]

Recommendation 5: Overcoming social system barriers to analyzing failure

Developing an environment in which people feel safe enough to identify and speak up is necessary but insufficient to produce learning from failure. Effective analysis of failure requires both time and space for analysis and skill in managing the conflicting perspectives that may emerge. In terms of time and space, some organizations, like the military, set aside time for After Action Reviews, while hospitals use M&M conferences to analyze failures.

In addition to putting such structures in place, leaders need to involve people with diverse perspectives and skills in order to generate deeper learning. This produces the tension

and conflict that both essential for learning and yet can interfere with the ability to keep a dialogue learning-oriented. Decades of research by organizational learning pioneer Chris Argyris has demonstrated that people in disagreement rarely ask each other the kind of sincere questions that are necessary for them to learn from each other.[45] People also try to force their views on the other party rather than educating the other party by providing the underlying reasoning behind these views.

For example, during the teleconference the night before the space shuttle *Challenger* was launched, both engineers and administrators proved unable have a discussion in which they were able to understand each other's concerns. Rather than try to explain what they saw in their (incomplete) data to educate the administrators and fill in the gaps in their understanding, the engineers made abstract statements such as "It is away from goodness to make any other recommendation" and "It's clear, it's absolutely clear." In turn, the administrators did not thoughtfully communicate their own concerns and questions, but rather contributed to an increasingly polarized discussion, in which the engineers' competencies were impugned. Eventually individuals with the most power made the decision.[46]

Thus, we recommend either developing or hiring skilled facilitators who can ensure that learning-oriented discussions take place when analyzing organizational failures. Managers can go through training to learn to test assumptions, inquire into others' views, and present their own views (no matter how seemingly correct or thorough) as incomplete – or partial – accounts of reality. These interpersonal skills can be learned, albeit slowly and with considerable effort, as action research has demonstrated.[47] When managers have these skills, they are able to model this behavior and to provide active coaching to others to help them be more effective in generating learning from the heated discussions that are often produced when analyzing failure.

Recommendation 6: Overcoming social system barriers to experimentation

As long as incentives are inconsistent with espoused values advocating learning from failure, true experimentation will be rare.[48] Managers thus need to align incentives and to offer resources to promote and facilitate effective experimentation. Those who experiment intelligently themselves and publicize both failures and successes demonstrate both value of these activities and help others see that the ideal of learning from failure in this organization is more than talk. In addition, coaching and clear direction may be useful in helping subordinates to understand what types of experiments should be designed. Finally, to develop the ability to manage all these processes, managers may need to work on their own psychological and emotional capabilities to enable them to shift how they think about failure.

Reframing failure

These six individual recommendations are best implemented as an integrated set of practices accompanied by an encompassing shift in managerial mindset. Table 2 summarizes this shift.

Insert Table 2 about here

First, failure must be viewed not as a problematic aberration that should never occur but rather than as an inevitable aspect of operating in a complex and changing world. This is of course not to say leaders should encourage people to make mistakes but rather to acknowledge that failures are inevitable and hence the best thing to do is to learn as much as possible – especially from small ones, so as to make larger ones less likely; beliefs about effective performance should reflect this. This implies holding people accountable not for avoiding failure but for failing intelligently and for how much they learn from their failures.

Organizational scholar Sim Sitkin identifies five characteristics of intelligent failures: "(1) they result from thoughtfully planned actions that (2) have uncertain outcomes and (3) are of modest scale, (4) are executed and responded to with alacrity and (5) take place in domains that are familiar enough to permit effective learning."[49]

Examples of unintelligent failure include making the same mistake over and over again, failing due to carelessness, or conducting a poorly designed experiment that would not produce helpful learning. Finally, managers need to create an environment in which they and their employees are open to putting aside their self-protective defenses and responding instead with curiosity and a desire to learn from failure.

Conclusion: Putting Failure to Work for Organizational Learning

This article noted that few organizations make effective use of failures for learning, due to formidable and deep-rooted barriers. We showed that properties of technical systems combine with properties of social systems in most organizations to make failures' lessons especially difficult to glean. At the same time, we highlighted noteworthy exceptions – organizations that have done a superb job of making failures visible, analyzing them systematically, or even knowingly encouraging failures as part of thoughtful experimentation.

Organizational learning from failure is thus not impossible but rather counter normative and often counter-intuitive. We suggest that making this process more likely requires breaking it down into essential activities – identifying failure, analyzing failure, and experimenting – in which individuals and groups can engage. By reviewing examples from a variety of organizations and industries where failures are being mined and put to good use through these activities, we sought to demystify the potentially abstract ideal of learning from failure. We offered six actionable recommendations and argued that these recommendations are best implemented by reframing managerial thinking rather than by treating them as a checklist of separate actions.

In conclusion, leaders can draw on this conceptual foundation as they seize opportunities, craft skills, and build routines, structures, and incentives that help their organizations enact these learning processes. At the same time, we do not underestimate the challenge of tackling the psychological and interpersonal barriers to this organizational learning process. As human beings, we are socialized to distance ourselves from failures. Reframing failure from something associated with shame and weakness to something associated with risk, uncertainty, and improvement is a critical first step in the learning journey.

Table 1: A Framework for Enabling Organizational Learning from Failure

	Key Processes in Organizational Learning From Failure		
	Identifying failures	**Analyzing failures**	**Experimentation**
Barriers embedded in Technical Systems	Complex systems make many small failure ambiguous	A lack of skills or techniques to extract lessons from failures.	Lack of knowledge of experimental design
Recommendations	*R1:* Training in systems analysis and judicious use of technical expertise from multiple disciplines	*R2:* Training in skills and techniques for systematic analysis of complex data	*R3*: Training in experimental design for effective experiments
Barriers embedded in Social Systems	Threat to self-esteem inhibits recognition of one's own failures, and corporate cultures that "shoot the messenger" limit reporting of failures.	Ineffective group process limits effectiveness of discussions in which failure analysis occurs	Organizations may penalize failed experiments inhibiting willingness to incur failure for the sake of learning
Recommendations	*R4*: Develop psychological safety in work groups and teams and celebrate and publicize failures as a means of learning	*R5* Develop forums for analyzing failure and provide training in skills for listening and inquiry (double loop learning)	*R6* Provide resources and reward systems to promote experimentation

Table 2: A Shift in Managerial Mindset to Promote Learning from Failure

	Traditional Frame	**Learning oriented reframe**
Expectations about failure	Failure is not acceptable	Failure is a natural byproduct of a healthy process of experimentation and learning
Beliefs about effective performance	Failure is avoided	Organizations are most effective when they learn from *intelligent failure* and communicate the lessons broadly in the organization.
Psychological and interpersonal responses to failure	Self-protective	Characterized by curiosity and humor that make it possible to learn from failure

Endnotes

[1] See Diane Vaughan, *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago, IL: University of Chicago Press, 1996.
[2] See S. Sitkin Learning through failure: The strategy of small losses. Research in Organizational Behavior 14:231-266.
[3] See especially A Tucker, A.L., and Edmondson, A.C. (2003). "Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change." *California Management Review*, 45,2: 55-72.
[4] See Cannon, M. and Edmondson, A. (2001). "Confronting failure: Antecedents and consequences of shared beliefs about failure in organizational work groups." *Journal of Organizational Behavior*, 22, 161-177.
[5] For example, see both Leonard-Barton, D. (1995). *Wellsprings of knowledge: Building and sustaining the sources of innovation*. Boston, Harvard Business School Press and Sitkin, S. B. (1992). "Learning through failure: the strategy of small losses." *Research in Organizational Behavior*. L. L. Cummings and B. Staw. Greenwich, CT, JAI Press. 14: 231-266.
[6] See especially A. C. Edmondson (2002). "The local and variegated nature of learning in organizations." *Organization Science 13, 2:* 128-146. Also see Rice, A. K. (1958). *Productivity and social organization: The Ahmedabad experiment*. London: Tavistock.
[7] See Trist, E.A., and Bamforth, K.W. (1951). Some social and psychological consequences of the Longwall method of coal-getting," in D. S. Pugh (Ed.), *Organization theory*. London: Penguin, (pp. 393-419).
[8] See F. Lee, Edmondson, A., Thomke, S, and Worline, M. forthcoming. The mixed effects of inconsistency on experimentation in organizations *Organization Science*.
[9] M.D. Cannon, Motivation and learning: A paradox for performance, paper presented at the 53rd annual Academy of Management Meeting, Atlanta, GA, (1993).
[10] See C. Argyris. *Overcoming organizational defenses* (1990).
[11] S. Finkelstein, *Why Smart Executives Fail and What You Can Learn from Their Mistakes*, New York, Portfolio (2003).
[12] See R.F. Baumeister, Self-Esteem: The puzzle of low self-regard. Plenum, New York (1993).
[13] See S.E. Taylor, and J. Brown, Illusions and well being: A social psychological perspective on mental health. *Psychological Bulletin, 103*, (1988). 193-210. Also see M.D. Cannon, The impact of positive illusions on performance, paper presented at the 59th annual Academy of Management Meeting. Chicago, IL, (1999).
[14] R. H. Hayes, S. C. Wheelwright, K. B. Clark (1988). *Dynamic Manufacturing: Creating the Learning Organization*. New York, Free Press.
[15] See S. Finkelstein, (2003). At Ref. Above. P. 282-3.
[16] See Make no mistake, Inc. Magazine, June (1989), p. 105.
[17] M. Moss (2002, June 27), Spotting breast cancer, doctors are weak link, *The New York Times [late ed.]*, pp. A1. See also, M. Moss (2002, June 28), Mammogram team learns from its errors, *The New York Times [late ed.]*, pp. A1.
[18] A. J. DiBella, J. M. Gould, and E. C. Nevis, *Understanding Organizations as Learning Systems* (n.d.), Society for Organizational Learning, retrieved July 8, 2002, from http://www.sol-ne.org/res/wp/learning-sys.html
[19] See S. W. Brown and S. S. Tax, Recovering and Learning from Service Failures, *Sloan Management Review* 40:1, 75-89 (1998).
[20] See S. W. Brown and S. S. Tax (1998). At Ref. above.
[21] See National Aeronautics and Space Administration report on the Space Shuttle Columbia at http://www.nasa.gov/home/hqnews/2003/oct/HQ_N03109_new_CAIB_volumes.html.
[22] A. Edmondson, *Organizing to learn*. Harvard Business School Note, 9-604-031. (2003).
[23] A. Edmondson, M.A. Roberto, and A. Tucker, *Children's Hospital and Clinics,* Harvard Business School Case, 9-302-050. (2002).
[24] See M.H. Bazerman. Judgment in Managerial Decision Making, fifth edition, John Wiley & Sons, New York, (2002). Also see S.T. Fiske and S.E. Taylor, Social Cognition, Random House, New York, 1984.

[25] Cannon and Edmondson, At. Ref Above. And A. Edmondson and B. Moingeon (1996). "When to learn how and when to learn why." In B. Moingeon and A. Edmondson (Eds). *Organizational learning and competitive advantage*. London: Sage.
[26] Tucker, A.L., and Edmondson, A.C. (2003). "Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change." *California Management Review*, 45,2: 55-72.
[27] For research on the lack of inquiry in group discussions, see Argyris, C. (1990). *Overcoming organizational defenses.* Needham, MA; Garvin, D. A., and M. A. Roberto, (2001) "What You Don't Know About Making Decisions," *Harvard Business Review* 79, 8: 108-116; for research on premature convergence on a solution or decision, see Janis, I.L. & Mann, L. (1977) *Decision-Making*. New York: The Free Press, and Langer, E. (1989). *Mindfulness*. Reading, MA, Addison-Wesley, and others.
[28] See A. Edmondson, M.A. Roberto, and A. Tucker, (2002). At Ref. Above.
[29] F. F. Reichheld with T. Teal, The Loyalty Effect: The Hidden Force Behind Growth, Profits, and Lasting Value, Harvard Business School Press, Boston (1996). Pages 194-195.
[30] F. F. Reichheld with T. Teal, (1996). At Ref. Above, Pages 207-208.
[31] P.F. Drucker, Innovation and Entrepreneurship: Practice and Principles, Harper & Row, New York, (1985). p. 43.
[32] S. Thomke, *Experimentation matters*. Harvard Business School Publishing (2003).
[33] Maidique, M. A. and Zirger, B. A Study of Success and Failure in Product Innovation: The Case of the U.S. Electronics Industry, *IEEE Transactions on Engineering Management*, 31 (4), 192-204 (1984).
[34] P.C. Wason, On the failure to eliminate hypotheses in a conceptual task. Quarterly Journal of Experimental Psychology 20, 273-283. (1960).
[35] T. Kelley with J. Littman, *The Art of Innovation: Lessons in Creativity from IDEO, America's Leading Design Firm*, Currency Books, New York, (2001). P. 232.
[36]See A. C, Edmondson and L. Feldman. "*Understand and Innovate" at IDEO Boston.* Harvard Business School Publishing, Case # 9-604-005. (2004).
[37] J. Pfeffer and R.I. Sutton, *The Knowing Doing Gap: How Smart Companies Turn Knowledge into Action*, Harvard Business School Press, Boston, (2000). P. 129.
[38] S. Thomke and A. Nimgade Bank of America. Harvard Business School Case, 9-603-022. (2002).
[39] S. Thomke and A. Nimgade (2002). At Ref. Above, p. 7.
[40] See F. Lee, A. Edmondson, S. Thomke, and M. Worline. The mixed effects of inconsistency on experimentation in organizations. *Organization Science*. Forthcoming.
[41] See P. Senge, *The firth discipline: The art and practice of the learning organization*. New York: Doubleday (1990) and A. Edmondson, R. Bohmer and G. Pisano, Disrupted routines: Team leaning and new technology implementation in hospitals. *Administrative Science Quarterly,* 46: 685-716. (2001).
[42] See A. C. Edmondson, et al. (2001) At Ref. Above.
[43] See A. C. Edmondson, Speaking up in the operating room: How team leaders promote learning in interdisciplinary action teams, *Journal of Management Studies* 40:6, 1419-1452 (2003).
[44] A. C. Edmondson. (2002). At Ref. Above.
[45] See C. Argyris *Strategy, change, and defensive routines*. (1985). New York: Harper Business.
[46] See A. C. Edmondson, Group process in the Challenger launch decision (B), Harvard Business School Case, N9-603-070. (2003).
[47] See Action Design website www.actiondesign.com
[48] See F. Lee et al, forthcoming.
[49] See S. Sitkin At Ref. above p. 243.