Chapter 20
Creating a Culture of Innovation

Arthur B. Markman

Abstract Innovation—the process of generating and implementing practical new ideas—can be difficult for organizations to do successfully. To make innovation a part of an organization, it is often necessary to change the culture in ways that bring more innovative processes into the workplace. In this chapter, I explore key factors that have to be part of a culture of innovation, including the need to favour innovation over efficiency and to tolerate failure. I also explore the importance of having an ecosystem to support the development of ideas. I illustrate these concepts with an example from the US military.

Keywords Innovation · Innovation ecosystem · Organizational culture
Failure

20.1 Introduction

In modern developed economies, innovation is a central part of a strategy for growth. It is hard for the developed world to compete with emerging economies on labor costs, and manufacturing jobs are often located where labor costs are cheapest. Advances in supply-chain management over the past 25 years have made it difficult for firms to compete by offering lower prices for goods that are commodities.

Although the word *innovation* is used in many ways, there is a general consensus that it refers to a process of developing, refining, and commercializing new ideas [1, 2]. Those ideas can lead to new products, but they can also lead to new services, new intellectual property, or new processes within an industry. As a result, innovation can lead to new markets (when a new product is developed) or to new levels of efficiency within an existing market (when a process is refined).

A.B. Markman (✉)
Department of Psychology, University of Texas,
108 E. Dean Keeton St., A8000, Austin, TX 78712, USA
e-mail: markman@utexas.edu

© Springer International Publishing AG 2018

C. Faucher (ed.), *Advances in Culturally-Aware Intelligent Systems and in Cross-Cultural Psychological Studies*, Intelligent Systems Reference Library 134, https://doi.org/10.1007/978-3-319-67024-9_20

Although firms of all sizes talk about the importance of innovation for their success, most are not able to be innovative on a consistent basis. Instead, many new ideas are developed initially within research and development labs owned by companies or organized at universities or government institutions. At companies that thrive on discovery (such as drug companies), there is also an elaborate infrastructure for testing new products and bringing them to market. For other research facilities, new ideas are often commercialized through entrepreneurial ventures that spin off to commercialize a product [3].

This state of affairs leads to a paradox. Firms value innovation and aspire to commercialize new ideas. Yet, they lack an infrastructure that supports innovation. More importantly, the culture at most firms also biases against the development of new ideas.

In this chapter, I begin by thinking about what it would mean for an organization to have a culture. Then, I explore three barriers to a culture of innovation within a firm. First, most organizations are structured around efficiency, which makes innovation difficult. Second, they often punish failure rather than learning from it. Third, large organizations are structured in ways that block the flow of information and ideas. As a result, it is crucial for organizations that want to become more innovative to create structures that allow different groups within the firm to work together. Each of these themes will be addressed in this chapter.

20.2 What Can a Culture of Innovation Be?

The word culture is used in many ways across disciplines. Cultural anthropologists focus on the rituals, social structures, and shared stories of members of groups. They describe cultures and also find dimensions that characterize differences among cultures [4]. These dimensions can be used to predict differences in behavior between members of different cultures [5].

Cultural psychologists then look at how aspects of culture shape the behavior of the members of that culture. For example, Nisbett and his colleagues have demonstrated broad cultural differences in performance on a variety of cognitive tasks between members of Western cultures (which tend to emphasize the importance of individuals) and members of East Asian cultures (which tend to emphasize the importance of the collective) [6, 7].

Cultural psychology also aims to understand the variables that culture can influence that may give rise to these differences in behavior. For example, studies in my lab explored the relationship between fear of isolation and cognitive performance [8]. These studies suggest that members of East Asian cultures have a higher chronic fear of isolation (on average) than members of Western cultures. That difference is correlated with differences in performance on a variety of cognitive tasks like the ones studied by Nisbett and his colleagues. Importantly, inducing a higher level of fear of isolation in members of a Western culture makes their cognitive performance closer to that observed in members of East Asian cultures.

The main idea here is that cultural factors may affect a motivational variable (in this case fear of isolation) that in turn affects cognition. Once culture is seen as an organizational force that influences core motivational factors, we can begin to look for other motivational influences that organizations have on the behavior of the members of that organization. The organizations that create these motivational orientations may be the global culture in which a person is raised, but they may also be the narrower set of social norms imposed by other organizations including religious institutions and workplaces.

This analysis suggests that a core active ingredient in cultures is a set of influences that organizations have on the motivational states of their members. In this paper, I focus on ways that organizations can affect aspects of motivation that influence innovation.

20.3 The Tradeoff Between Efficiency and Innovation

Without realizing it, most companies are structured based on a manufacturing model, even if their core business is intellectual property or service. The core of manufacturing is efficiency. Companies typically focus themselves on efficiency in a number of ways.

First, they try to minimize personnel costs. In many companies, hiring a new employee requires justifying why this position is needed. In some cases, the new hire will bring new skills that enhance the functionality of the organization. More often, though, the existing workload has become too much for the existing employee base to handle, and so the justification of the new hire rests on increasing the capacity of the organization to do business.

In many ways, this mode of hiring parallels the way a manufacturing plant is run. Each machine in the plant should be used to its capacity. Many plants run three shifts a day in order to allow the machinery to be used full-time. New machines are only purchased when they create a necessary new functionality for the plant, or when the current machinery is being fully utilized.

In order to make these determinations, factories and human resource (HR) departments measure employee productivity. In service companies, HR aims to maximize billable hours. In other firms, yearly evaluations lay out a set of goals (or *contributions* to use Peter Drucker's [9] term), and successful evaluation requires reaching these goals.

This strategy is effective for job functions that are focused on execution (which is really the knowledge- and service-work equivalent of manufacturing). When fast-food chains measure time-per-order and call-centers seek to minimize the length of service-calls, they are maximizing the efficiency of execution.

Execution isn't just perfected by focusing on minimizing time, of course. The aviation industry in collaboration with government agencies like the FAA focus on minimizing error [10, 11]. They catalog nearly every error made during routine work and then use those errors to fix procedures that may be leading to those

problems. They also analyze all malfunctions and catastrophes to draw lessons to make future air travel safer.

Because errors, time to completion, and billable hours are easy to measure, they become important yardsticks for improving efficiency and evaluating employees. Unfortunately, these evaluations also bias most companies against innovation.

In particular, innovation does not have a straightforward time-course the way manufacturing and execution do. Research and development to create new knowledge is slow and expensive. In addition, there is no clear way to measure how close a project is to a new discovery. Many R&D projects can feel like they are stumbling around creating as many problems as they solve until suddenly the path to a solution is clear.

The same things holds true for other kinds of idea-generation techniques. Often, new ideas emerge as a result of analogies between a current problem and an existing solution in another domain [12, 13]. That means that for employees to have a chance to consistently generate new ideas, they have to continually expand their knowledge base. Furthermore, it is hard to target in advance which knowledge will lead to successful solutions to difficult problems. As a result, employees in creative roles must be given latitude to explore new domains whose applicability to current projects may not be clear. Only in retrospect will it be obvious which aspects of their knowledge were useful.

Even after promising ideas are generated, they must be further refined and ultimately commercialized. At each stage of this process, new problems may emerge that can set back the time-line of a project or lead it to be abandoned. Thus, unlike manufacturing and execution, it is difficult to manage innovative projects efficiently.

On the one hand, firms need to institute procedures that allow them to make decisions about whether to continue or to abandon projects [14]. There is certainly collected wisdom about innovation that can guide firms through the process.

At the same time, firms that want to become more innovative must commit to the expense of overstaffing. The most popular example of this kind of overstaffing is Google's "20% Time," in which (at least early in the company's development) all employees were encouraged to take about one day a week to pursue projects of interest to them, regardless of their other work priorities [15]. Although the need to execute has cut into the amount of flexible time that many employees at Google have [16], this idea reflects a recognition that company-wide innovation means that more employees need to be hired overall than would be required simply to execute existing company priorities.

Thus, the HR process for creating new positions must take into account that a certain amount of over-capacity in each division of an organization is required if the firm is committed to sustained innovation. Yearly evaluations of employees must focus not just on measures of efficiency, but also on engagement with activities that are reliably associated with the creation and commercialization of new ideas.

Finally, within any organization, it is important to set up reward structures that motivate people to act in ways that promote innovation. In every organizational

culture, there is what leaders say, what people do, and what the organization rewards. These factors are listed in increasing order of importance.

Regardless of what leaders in the organization claim that they want to do, individuals are sensitive to what they see other people doing. Considerable work on goal contagion demonstrates that people automatically adopt the goals of people they observe. Seeing people in the workplace engaged in activities associated with innovation promotes that same behavior in others.

Rewards at work are particularly important for guiding behavior. Rewards can be pay and promotion, but they can also be praise and opportunity. People like to be recognized for their good work, and so praise influences behavior. In addition, studies of well-being in the workplace suggest that meaningful work is particularly important to people. Providing opportunities for people who engage in desirable behaviors to have more meaningful work opportunities is an excellent way to reward this desirable behavior.

One barrier to rewarding behavior associated with innovation is that existing corporate cultures that are focused on efficiency can create anxiety about work that is not progressing toward a clearly defined goal. One reason why the actions of leaders in the organization is so important is that many employees find it stressful to take time out of the workday to read and improve skills that are not directly related to specific ongoing tasks.

As this section suggests, many organizations (implicitly or explicitly) have adopted a culture of efficiency and execution. This orientation is incompatible with innovation. Without addressing this tradeoff directly, firms will have a hard time motivating employees to engage in more behaviors that lead to innovations.

20.4 Orientations to Failure

James Dyson's invention of a bagless vacuum is often held up as a great example of innovation [13]. Discussions of his achievement often focus on the application of industrial cyclone technology normally used in settings like sawmills to home vacuum cleaners. These discussions do not typically talk about the long years of effort that Dyson put in after having that motivating insight until he perfected a model that functioned well and could be brought to market.

An important lesson from this example is that there are many small failures that are part of even the most successful story of innovation. And most innovative ventures do not succeed, which means that even after running the innovation process flawlessly, there are many factors that can cause the project to fail. Competing projects may cause the market to shift in unexpected ways. An economic downturn can cause potential customers to delay purchases longer than the company can wait for sales.

Because innovation has a number of inherent risks, a company that wants to promote more innovation needs to be willing to tolerate failure. Unfortunately, companies often punish failure in a number of ways.

First, the kinds of employees that companies are likely to hire are often afraid of failure. Success in the modern education environment is defined as getting good grades. Good grades result from high test scores, and that means that the best students are the ones who make the fewest mistakes. Thus, most successful students try to avoid failure rather than embracing it.

Even when particular individuals are comfortable with failure, many firms are not. For example Saxenian [17] explored factors that led Silicon Valley to become the dominant center for high-tech industry in the United States by late 1980's. In the 1940's, it looked like the East Coast of the United States (and particularly the area around Boston, MA) was the best candidate region to excel in high tech. The US Government invested heavily in the research universities in this area (like MIT) and companies (like IBM and Raytheon) during and after World War II.

These large companies had a big technical advantage over West Coast firms in the 1960s, but these larger firms tended to stifle innovation in two ways. First, executive pay was typically tied to success of business units. If an innovative project failed, then profits from that unit lagged. Thus, management often avoided supporting innovative projects.

Second, mid-level managers who were given responsibility for innovative projects were often punished for leading projects that did not succeed. They were passed over for promotion and marginalized within the business.

In contrast to this treatment of failure, firms on the West Coast (many of which were smaller entrepreneurial ventures) were more tolerant of failures. People who worked for startups that failed were seen as having valuable experience that would make them more effective in subsequent companies. There was a general recognition that most of the factors that lead to success or failure of a particular venture are beyond the control of that firm, and so hard-working people with experience are valuable, regardless of whether past ventures succeeded.

This analysis suggests that many companies need to shift their culture to be more tolerant of failures and to treat them as learning experiences. A shorthand recommendation is to suggest that firms should punish negligence rather than failure.

One reason that failure is so important is that it provides valuable information. Scientists focus on disconfirming hypotheses, because a large amount of supporting evidence for a particular theory does not mean that evidence disconfirming it will not ultimately be found [18]. Similarly, it can be difficult to assess which collection of factors led a venture to be successful, but the factors that lead to a failure can be easier to diagnose.

As an example, when working for Apple, Ron Johnson spearheaded the project to develop the Apple Store. The modern design and smart helpful employees were an immediate success and continue to be a popular destination for technology shoppers. Based on this success, Johnson was hired to be the CEO of JC Penney, the mall department store known for inexpensive clothing and housewares. To turn around the struggling retailer, Johnson updated the design of JC Penney stores in an effort to appeal to a younger more stylish crowd [19]. Unfortunately, this strategy was a failure. It alienated existing customers, but failed to attract patrons of chains like H&M that offer trendy designs at low cost.

This example suggests that Johnson had difficulty determining the factors that led to the success of the Apple Store. Consequently, he repeated the formula that worked at Apple without modifying it to suit differences between the customers Apple hoped to attract and the customers that JC Penney was likely to be able to attract. In addition, The Apple Store was a new venture, and so Johnson did not have to worry about an existing customer base. In contrast, JC Penney attracted shoppers who were quite different than those that the redesigned store was target to appeal to.

Because failure creates learning opportunities, many firms need to change their orientation toward failure. The guiding principle is to punish negligence, but not failure. In practice, that means that yearly employee evaluations need to incorporate both employees' accomplishments as well as the process that they go through to achieve those results. This focus on process ensures that members of a firm are working toward positive innovation outcomes. It also recognizes that when employees are working on innovative projects, there are many sources of failure that are beyond the control of an individual or team.

As discussed in the previous section, the execution orientation of many firms leads them to focus on minimizing errors. Thus, tolerance for failure contradicts a central tenet of many corporate cultures. That is why it is crucial to build a focus on innovative activities rather than innovative outcomes into evaluations. People who have been socialized into organizations are likely to feel uncomfortable rewarding someone whose project has not succeeded. Thus, the criteria for evaluation need to be laid out explicitly to help people get beyond their habitual reaction.

20.5 Creating an Ecosystem that Supports Innovation

A critical part of successful innovation in organizations is having a flow of information that allows groups to share ideas and allows the organization to use all of its resources to enable a new idea to succeed. Unfortunately, many large organizations are structured in a way that limits the spread of information.

When organizations are small (particularly when they have less than 25 people), the social structure in the organization need not be that complex. Everyone in the organization is likely to have a clear sense of the roles and knowledge of their colleagues. As a result, information in small organization passes freely. This is one reason why startup companies are often successful as small entrepreneurial ventures.

As organizations get larger, the structure begins to grow. There is some evidence suggesting that when organizations get larger than about 150 people, it is not possible for everyone to really know everyone else [20]. Practically speaking, however, the management structure of an organization begins to get more complex above about 25 people. At that point, groups of people are required to perform specialized functions (such as technical support, customer support, or accounting).

As those groups begin to form, they now contain substructures within the whole organization. The basic psychology of group processes leads these clusters within an organization to form silos that distinguish between those inside the group and those outside [21]. The members within a group work well together, but they are less prone to share information with people outside it. This happens both because members of a group work in proximity to each other and far from others, and because there is a higher level of trust within the group than across groups. As a result, these groups naturally form silos that keep information contained within them and limit the flow of information outward.

Thus, even though members of different silos within an organization are part of the same company overall, the greater group cohesion within a silo than across silos hampers innovation. Members of different silos may not reach out for information or help to the right people. Indeed, as the organization grows later, members of one silo may be unaware of the skills and knowledge contained in different silos.

A related problem involves bureaucratic structures that are developed as organizations get larger. Because the members of one silo are unlikely to be well-acquainted with members of other silos, systems of checks and balances are often put in place to ensure that employees in the organization act honestly and that power is shared. These procedures take the place of the personal relationships that are the basis of trust in smaller organizations.

These structures also limit the flow of information. Members of particular groups often expect information to flow upward to the management of a silo and for knowledge to be transferred across groups by management. Other organizations create procedures for sharing knowledge that involve the creation of reports that are stored in a central database. These reports are meant to be used by teams that want to make use of lessons learned by other segments of the organization.

In practice, these strategies also limit the flow of information that is crucial to innovation. Innovation is most effective when the person who needs the information communicates directly with the person who has that knowledge. Hierarchical structures that pass information place a number of people in between the two people who ultimately need to connect. Those layers of bureaucracy decrease the likelihood that people who need to be in contact will actually connect.

Creating databases of reports is rarely successful as a means of capturing organizational knowledge to support innovation [22]. First, these reports are often written in a cursory way, because people are anxious to move on to the next task rather than capturing what was learned in the task just completed. As a result, the reports themselves do not always provide details of the problem that was solved or the method used to solve it.

Even when the documents are well-written, they need to be found by people who need them for the documents to be useful. In order to retrieve a document in a database, you need to search on terms that are actually in that document. Unfortunately, because many innovations involve knowledge that comes from across disciplines, the search terms are rarely those that were used in previous documents that may have been relevant to solving a new problem [23].

The fundamental difficulty is that search is literal. That is, when people are seeking information in most databases, the documents need to contain specific words that are part of the query. Someone solving a problem about traffic flow is likely to search using terms from that domain like traffic, roads, cars, and driving. That search will yield documents that are also about traffic on the roads.

What we want instead is analogical search [24, 25]. The idea is to find other domains that have the same structure as the problem you are trying to solve. For example, there is a broad analogy between road systems and the circulatory system that might provide productive avenues for exploring a problem.

These analogies preserve relationships among elements in the domain rather than the objects in the domains [26–29]. In the comparison between roads and the circulatory system, there is a parallel between capillaries and arteries, which differ in the amount of blood that can flow through them and side streets and boulevards, which differ in the traffic they can handle. Blood does not look like cars, and veins and arteries do not look like roads. They are analogous just because of these parallel relationships.

To find an analogous domain like this, you could try to use more abstract search terms that might potentially describe both domains. There are two problems with this approach. First, analogies are not really about abstraction, but rather about similarities in relationships. Second, the natural way to write about anything is to use terms relevant to the domain in which you are writing. Thus, the people writing the documents in the database will use specific terms and not the abstract ones. All of these factors clarify why databases of reports are mostly useful for helping people at a company to re-use solutions when the same problem arises again.

Thus, to promote sharing of knowledge, organizations need to find ways to encourage more interactions among people from different silos.

20.5.1 Studying Business Ecosystems

In order to explore ways to encourage more interactions within firms, my colleagues and I looked to an analogous case [30]. We examined how networked technology incubators function.

Technology incubators are organizations that help early-stage technology companies to develop and commercialize technological innovations. In the 1980s, the dominant role of the incubator was to provide inexpensive office space for fledgling companies. Bringing several companies together in the same space was also valuable, because the members of those companies could share expertise.

Over time, the office space became less important than the social interactions. To explore these interactions in more detail, we examined the Austin Technology Incubator (ATI) in great detail. The goal of ATI is to take new companies and help them to get their first round of investor funding. Research assistants followed directors of the incubator, interviewed members of companies and other members of the technology community, and analyzed emails and invitations to incubator

events. For comparison, the team also examined less successful incubators in Portugal.

Two important features emerged from these analyses. First, successful incubators focus on companies that have a common goal. When a new venture begins, its owners may want to develop a company that they grow and nurture for life. These *lifestyle companies* often grow slowly and aim to be small to moderate-sized businesses. Other teams want to create companies that will ultimately lead to large exits by IPO or sale that allow the team to reap a large reward. These companies seek significant investment from venture capital and angel investors in order to grow quickly. The cost of taking on these investments is that they have to grow significantly in order for the investors to see a reasonable return.

Successful incubators focus exclusively on one type of company. For example, all of the companies at ATI ultimately plan to have an IPO or to be sold to a larger organization. Unsuccessful incubators bring in many different types of companies with a diverse set of goals.

The reason why a common set of goals is crucial is because the primary value of incubators is that they create an ecosystem for the survival of new companies that functions like a metaphorical *coral reef*. The idea is that the incubator protects new companies. It attracts a broad community of individuals who can help new companies to thrive. Potential investors meet the team. Technology experts help the company to develop its product. Business experts refine the company's business model. Students from nearby universities engage with companies and get experience with entrepreneurship.

The interactions between companies and this community are not structured. Instead, the incubator holds a variety of events that support happenstance interactions among community members and companies in ways that benefit the startups as well as the surrounding business community.

If the companies in an incubator have divergent goals, then community members do not have enough positive interactions when they attend incubator events, and so they stop coming. When there is a critical mass of companies with a common goal, then there is a high probability that a community member who engages with incubated companies will benefit.

It is valuable to understand how these entrepreneurial ecosystems function, because it allows directors of incubators to create more effective ways to nurture a startup community in a region. However, we were also interested in this ecosystem, because a similar structure might be valuable for large organizations to help them break down silo walls.

20.5.2 The Innovation Ecosystem

The same principles for success of networked incubators can be applied to innovation to create an *innovation reef*. The core idea is to create events for people from

different silos of a company to come together on a regular basis to exchange ideas and to talk about projects they are working on.

A reef cannot be legislated by the management of an organization. Instead, the individual events need to add value for the participants so that a critical mass of employees engage in them and so that they spur interactions.

It is possible to use some technology to support these interactions. For example, company-wide wikis and bulletin boards where people can ask questions are an excellent support for a reef.

However, there has to be a number of live events for people to attend. The active ingredient in the innovation ecosystem is serendipitous discussions that happen surrounding events rather than targeted queries. So, events within the reef have to include sufficient time for participants to engage in small group discussions. Of course, in order to provide time for people to attend reef events, there has to be enough over-capacity (as discussed earlier) to allow people the flexibility to be part of the reef while still handling their day-to-day responsibilities.

In order to jump start social interactions across silos, it can be helpful for managers to schedule "field trips" in which one group gets together with another for an extended lunch in which key group members give talks to describe what they have been working on and to raise key problems they are addressing. These events help to develop personal relationships among individuals across groups that can be maintained at later events.

In addition, for the innovation ecosystem to succeed, there has to be some clear mechanism for groups that form around innovative projects to get funding to continue their projects. It is also helpful if there is a standardized way for members of groups that are working on innovative projects to apply for more time to focus on those projects when they look promising.

This type of innovation ecosystem contrasts with a more typical way that large companies have tried to spur innovation. Taking a cue from successful design firms like IDEO [31], many companies created in-house design teams that would be staffed with experts in design thinking and idea generation. These teams would often have a fancy off-site location where intact groups from the company would be brought to engage in creativity and innovation exercises.

Unfortunately, many of these centers proved hard to sustain [32]. Sustaining innovative projects requires more than just good design thinking. It requires extensive collaboration across business units throughout the development of a project. Consequently, a more distributed approach to innovation that involves members from different research teams is a more sustainable model for supporting innovation in large organizations.

20.5.3 Putting the Reef into Practice

Large organizations are starting to implement this kind of innovation reef. In this section, I focus on two examples, one of which has been more successful to date than the other.

The more successful reef has been developed by USAA, the large insurance and financial services firm [33]. USAA started as an automobile insurance company that catered to members of the United States military, veterans, and their families. Over the years, the company has diversified into a broad-based financial services firm that serves this client base.

In an effort to become more innovative, USAA partnered with the University of Texas to develop training for a broad base of their employees to learn an end-to-end innovation strategy that encompasses idea development, evaluation of new technologies, and implementation of innovative projects. The trainees for this program are selected explicitly from across a variety of segments of the company and from different job functions.

There are three aims to this training. First, it distributes expertise in innovation across business units within the organization. Second, it creates networking opportunities for individuals from different business units who can find common ground and establish relationships with people from other silos. Third, it rewards individuals who are interested in innovation with additional training and opportunities that allow them to advance their career.

A great example of the success of this program is the Pole Cam [33]. USAA provides homeowner's insurance to members. When a policy holder sustains roof damage, an adjuster needs to assess that damage. This job is dangerous, particularly in the aftermath of a significant storm. A cross-disciplinary team at USAA convened to develop a way to minimize the number of roof inspections that required adjusters to get up on the roof to inspect it.

Their first solution involved a portable remote-controlled drone with a camera attached that would fly and hover over damaged roofs. Unfortunately, while the drone worked effectively in preliminary tests indoors, it was easily blown off course by gusts of wind, and so that idea was scrapped. Ultimately, the team developed an extendible fiberglass poll with a digital camera on it controlled by a tablet at the base. This lower-tech solution has been deployed in the field, and has drastically reduced the number of times adjusters need to climb on damaged roofs.

The USAA reef is succeeding, because there is a commitment to continuing to train interested employees in methods of innovation and to provide these employees with time to focus on innovative projects. In this way, the reef is well-aligned with the organization's reward structure. Finally, as the Pole Cam example demonstrates, the teams are willing to learn from their failures and to continue developing innovative projects with the full support of management.

A second example of a reef that was less successful involved United States Special Operations Command (SOCOM). SOCOM, which operates out of MacDill Air Force Base in Tampa, Florida, is the command center for US Special Forces (including the Army Rangers and the Navy Seals). Following the terrorist attacks in the US on September, 11, 2001, Special Operations Command was given a mandate to coordinate defense against terror attacks. As a result, the command structure of SOCOM swelled and came to include military personnel, civilian contractors, and government employees.

In 2013, following a conference on innovation held at the request of Adm. William McRaven (ret.), SOCOM set up their own reef (as evidenced by Twitter account @SOFReef) that aimed to bring together individuals from across the organization interested in innovation. They also developed a number of on-line tools to support communication about problems requiring innovative solutions. One of these, called SOFBox (which had a hashtag #sofbox on Twitter) aimed to be a central bulletin board for the exchange of ideas.

Unfortunately, the reef and the tools like SOFBox were not as successful as the reef at USAA. Indeed, the SOFReef twitter account has not had a new tweet since 2013. A key reason why the reef structure has had a harder time getting established at SOCOM is because of the three types of employees who work there. The HR rules for military personnel, government employees, and civilian contractors are all different. Consequently, it is difficult to create a uniform way to involve people across these groups in innovative projects. The rules governing each of these types of employees has the unintended consequence of maintaining silo walls despite a desire by the upper-level command structure of SOCOM to improve innovation. However, efforts to work with these constraints are under way through the SOCOM iLab (http://www.socom.mil/sofilab/default.aspx).

Looking across these examples, it is clear that a reef can be created that leads to successful innovations. However, it cannot be mandated top-down. Instead, it is important for management to provide support and flexibility to individuals who engage with each other across silos and to tolerate failure. When an organization cannot be flexible in the way that it uses employees, then it is hard to sustain a reef, despite the best of intentions.

20.6 Summary and Future Directions

Most organizational cultures have a variety of elements that tend to perpetuate existing behaviors. In particular, existing reward structures make it easiest for people to continue engaging in behaviors that have brought them success in the past. HR practices that focus on efficiency make it difficult for people to be flexible in their pursuit of innovative ideas. Fear of failure biases individuals away from innovative projects.

A central difficulty that large organizations have when trying to innovate is that it is difficult for them to share information across the silos that inevitably develop when an organization grows larger than about 100 people. To help information flow across silo walls, it is valuable to create an innovation reef that creates events that allow unstructured interactions among people from different groups to share information and to coalesce around new ideas. Organizations interested in creating a reef need to have some flexibility in their HR practices to allow employees to have the time to devote to new projects.

Future work must continue to explore methods for improving information flow in organizations, particularly those that are geographically dispersed. Multi-national

companies have a hard time getting employees in distant offices to work together. For one, differences in time zone mean that employees in different countries may only overlap in their workdays for a few hours each day. Even when these individuals synchronize their work schedules, the distance means that conversations must be scheduled and mediated by technology that still does not provide a deep sense of copresence.

In addition, more detailed case studies of organizations that have changed their culture to become more innovative will support the development of new tools. Ultimately, the reef is just one ecosystem metaphor that is useful for thinking about how to improve innovation in organizations.

References

1. Chesbrough, H.W.: Open Innovation: The New Imperative for Creating and Profiting from Technology. Harvard Business Review Press, Cambridge, MA (2003)
2. Markman, A.B., Wood, K.L. (eds.): Tools for Innovation. Oxford University Press, New York (2009)
3. Gibson, D.V., Rogers, E.M.: R&D collaboration on trial. Harvard Business School Press, Boston, MA (1994)
4. Triandis, H.C.: Individualism and collectivism: past, present, and future, In: Matsumoto, D. (ed.) The Handbook of Culture and Psychology, pp. 35–50. Oxford University Press, New York (2001)
5. Hofstede, G., Van Hofstede, G., Minkov, M.: Cultures and Organizations, 3rd edn. New York: McGraw-Hill (2010)
6. Nisbett, R.E., et al.: Culture and systems of thought: Holistic versus analytic cognition. Psychol. Rev. **108**(2), 291–310 (2001)
7. Peng, K.P., Nisbett, R.E.: Culture, dialectics, and reasoning about contradiction. Am. Psychol. **54**(9), 741–754 (1999)
8. Kim, K., Markman, A.B.: Differences in fear of isolation as an explanation of cultural differences: Evidence from memory and reasoning. J. Exp. Soc. Psychol. **42**, 350–364 (2006)
9. Drucker, P.F.: The Practice of Management. HarperCollins Publishers, New York (1954)
10. Helmreich, R.L.: On error management: lessons from aviation. BMJ **320**(7237), 781–785 (2000)
11. Harper, M.L., Helmreich, R.L.: Identifying barriers to the success of a reporting system. In: Henriksen, K., Battles, J.B., Marks, E.S. (eds.) Advances in patient safety: From Research to Implementation (vol. 3: Implementation issues). Agency for Healthcare Research and Quality, Rockville, MD (2005)
12. Basalla, G.: The Evolution of Technology. Cambridge University Press, Cambridge, UK (1988)
13. Markman, A.: Smart Thinking. Perigee Books, New York (2012)
14. Cornwell, B.: Quicklook Commercialization Assessments. R&D Enterprise: Asia Pacific, **1**(1) (1998)
15. Walker, A.: Creativity loves constraints: The paradox of Google's twenty percent time. Ephemera **11**(4), 369–386 (2011)
16. Ross, A.: Why did Google Abandon 20% Time for Innovation. 4/4/16] (2015)
17. Saxenian, A.: Regional Advantage. Harvard University Press, Cambridge, MA (1996)
18. Platt, J.R.: Strong Inference. Science **146**, 347–352 (1964)
19. Lublin, J.S., Mattioli, D., Co, J.P.: Penney CEO Out, Old Boss Back In, in Wall Street Journal (2013)

20. Dunbar, R.I.M.: Neocortex size as a constraint on group size in primates. J. Hum. Evol. **22**(6), 469–493 (1992)
21. Gawronski, B., Bodenhausen, G., Banse, R.: We are, therefore they aren't: Ingroup construal as a standard of comparison for outgroup judgments. J. Exp. Soc. Psychol. **41**, 515–526 (2005)
22. Smith, S.M., et al.: The development and evaluation of tools for creativity. In: Markman, A.B., Wood, K.L. (eds.) Tools for Innovation. Oxford University Press, New York (2009)
23. Linsey, J.S., Wood, K.L., Markman, A.B.: Modality and representation in analogy. Artif. Intell. Eng. Des. Anal. Manuf. **22**(2), 85–100 (2008)
24. Ward, T.B.: ConceptNets for flexible access to knowledge. In: Markman, A.B., Wood, K.L. (eds.) Tools for Innovation. Oxford University Press, New York (2009)
25. Forbus, K.D., Gentner, D., Law, K.: MAC/FAC: a model of similarity-based retrieval. Cognitive Sci. **19**(2), 141–205 (1995)
26. Gentner, D., Markman, A.B.: Structural alignment in analogy and similarity. Am. Psychol. **52**(1), 45–56 (1997)
27. Gentner, D.: Structure-mapping: a theoretical framework for analogy. Cognitive Sci. **7**, 155–170 (1983)
28. Gentner, D.: The mechanisms of analogical learning. In: Vosniadou, S., Ortony, A. (eds.) Similarity and Analogical Reasoning, pp. 199–241. Cambridge University Press, New York (1989)
29. Hummel, J.E., Holyoak, K.J.: Distributed representations of structure: A theory of analogical access and mapping. Psychol. Rev. **104**(3), 427–466 (1997)
30. Pogue, G.P., et al.: Building an innovation coral reef: the austin technology incubator case study. In: Markman, A.B. (ed.) Open Innovation, in press. Oxford University Press, New York
31. Kelley, T., Littman, J.: The art of innovation: Lessons in creativity from IDEO, American's leading design firm. Crown Business, New York (2001)
32. Fisher, W.: Applying the creative problem solving process to open innovation. In: Markman, A.B. (ed.) Open Innovation, in press. Oxford University Press, New York
33. Zintgraff, C., et al.: Better, faster, safer, and cheaper: USAA roof inspections with Pole Cam. In: Markman, A.B. (ed.) Open Innovation, pp. 121–140. Oxford University Press, New York (2016)