



ORGANIZATIONAL CULTURE

The Hard Truth About Innovative Cultures

by [Gary P. Pisano](#)

FROM THE JANUARY-FEBRUARY 2019 ISSUE

A culture conducive to innovation is not only good for a company's bottom line. It also is something that both leaders and employees value in their organizations. In seminars at companies across the globe, I have informally surveyed hundreds of managers about whether they want to work in an organization where innovative behaviors are the norm. I cannot think of a single instance when someone has said "No, I don't." Who can blame them: Innovative cultures are generally depicted as pretty fun. When I asked the same managers to describe such cultures, they readily provided a list of characteristics identical to those extolled by management books: tolerance for failure, willingness to experiment, psychological safety, highly collaborative, and nonhierarchical. And research supports the idea that these behaviors translate into better innovative performance.

But despite the fact that innovative cultures are desirable and that most leaders claim to understand what they entail, they are hard to create and sustain. This is puzzling. How can practices apparently so universally loved—even fun—be so tricky to implement?

The reason, I believe, is that innovative cultures are misunderstood. The easy-to-like behaviors that get so much attention are only one side of the coin. They must be counterbalanced by some tougher and frankly less fun behaviors. A tolerance for failure requires an intolerance for incompetence. A willingness to experiment requires rigorous discipline. Psychological safety requires comfort with brutal candor. Collaboration must be balanced with individual accountability. And flatness requires strong leadership. Innovative cultures are paradoxical. Unless the tensions created by this paradox are carefully managed, attempts to create an innovative culture will fail.

1. Tolerance for Failure but No Tolerance for Incompetence

Given that innovation involves the exploration of uncertain and unknown terrain, it is not surprising that a tolerance for failure is an important characteristic of innovative cultures. Some of the most highly touted innovators have had their share of failures. Remember Apple's MobileMe, Google Glass, and the Amazon Fire Phone?

And yet for all their focus on tolerance for failure, innovative organizations are intolerant of incompetence. They set exceptionally high performance standards for their people. They recruit the best talent they can. Exploring risky ideas that ultimately fail is fine, but mediocre technical skills, sloppy thinking, bad work habits, and poor management are not. People who don't meet expectations are either let go or moved into roles that better fit their abilities. Steve Jobs was notorious for firing anyone he deemed not up to the task. At Amazon, employees are ranked on a forced curve, and the bottom part of the distribution is culled. Google is known to have a very employee-friendly culture, but it's also one of the hardest places on earth to get a job (each year the company gets more than 2 million applications for about 5,000 positions). It, too, has a rigorous performance management system that moves people into new roles if they are not excelling in their existing ones. At Pixar, movie directors who cannot get projects on track are replaced.

It sounds obvious that companies should set high quality standards for their employees, but unfortunately all too many organizations fall short in this regard. Consider a pharmaceutical company I recently worked with. I learned that one of its R&D groups had not discovered a new drug

candidate in more than a decade. Despite the poor performance, senior leaders had made no real changes in the group's management or personnel. In fact, under the company's egalitarian compensation system, the scientists in the group had been receiving approximately the same salaries and bonuses as scientists in much more productive R&D units. One senior leader confided to me that short of ethics violations, the company rarely terminated anyone in R&D for subpar performance. When I asked why, he said, "Our culture is like a family. Firing people is not something we're comfortable with."

The truth is that a tolerance for failure requires having extremely competent people. Attempts to create novel technological or business models are fraught with uncertainty. You often don't know what you don't know, and you have to learn as you go. "Failures" under these circumstances provide valuable lessons about paths forward. But failure can also result from poorly thought-out designs, flawed analyses, lack of transparency, and bad management. Google can encourage risk taking and failure because it can be confident that most Google employees are very competent.

Creating a culture that simultaneously values learning through failure and outstanding performance is difficult in organizations with a history of neither. A good start is for senior leadership to articulate clearly the difference between productive and unproductive failures: Productive failures yield valuable information relative to their cost. A failure should be celebrated only if it results in learning. (The cliché "celebrating failure" misses the point—we should be celebrating learning, not failure.) A simple prototype that fails to perform as expected because of a previously unknown technical issue is a failure worth celebrating if that new knowledge can be applied to future designs. Launching a badly engineered product after spending \$500 million developing it is just an expensive flop.

Building a culture of competence requires clearly articulating expected standards of performance. If such standards are not well understood, difficult personnel decisions can seem capricious or, worse, be misconstrued as punishment for a failure. Senior leaders and managers throughout the organization should communicate expectations clearly and regularly. Hiring standards may need to be raised, even if that temporarily slows the growth of the company.

Managers are especially uncomfortable about firing or moving people when their “incompetence” is no fault of their own. Shifting technologies or business models can render a person who’s very competent in one context incompetent in another. Consider how digitization has impacted the value of different skills in many industries. That sales representative whose deft interpersonal skills made him a superstar may no longer be as valuable to the organization as the introverted software engineer who develops the algorithms used to predict which customers are most likely to buy the company’s products. In some cases, people can be retrained to develop new competences. But that’s not always possible when really specialized skills (say, a PhD in applied math) are needed to do a job. Keeping people who have been rendered obsolete may be compassionate, but it’s dangerous for the organization.

Maintaining a healthy balance between tolerating productive failures and rooting out incompetence is not easy. A 2015 *New York Times* article about Amazon illustrates the difficulty. The piece, which was based on interviews with more than 100 current and former employees, labeled Amazon’s culture as “bruising” and recounted stories of employees crying at their desks amid enormous performance pressures. One reason striking a balance is so hard is that the causes of failure are not always clear. Did a product design turn out to be flawed because of an engineer’s bad judgment or because it encountered a problem that even the most talented engineer would have missed? And in the event of bad technical or business judgments, what are the appropriate consequences? Everyone makes mistakes, but at what point does forgiveness slide into permissiveness? And at what point does setting high performance standards devolve into being cruel or failing to treat employees—regardless of their performance—with respect and dignity?

2. Willingness to Experiment but Highly Disciplined

Organizations that embrace experimentation are comfortable with uncertainty and ambiguity. They do not pretend to know all the answers up front or to be able to analyze their way to insight. They experiment to learn rather than to produce an immediately marketable product or service.

A willingness to experiment, though, does not mean working like some third-rate abstract painter who randomly throws paint at a canvas. Without discipline, almost anything can be justified as an experiment. Discipline-oriented cultures select experiments carefully on the basis of their potential learning value, and they design them rigorously to yield as much information as possible relative to the costs. They establish clear criteria at the outset for deciding whether to move forward with,

modify, or kill an idea. And they face the facts generated by experiments. This may mean admitting that an initial hypothesis was wrong and that a project that once seemed promising must be killed or significantly redirected. Being more disciplined about killing losing projects makes it less risky to try new things.

A good example of a culture that combines a willingness to experiment with strict discipline is Flagship Pioneering, a Cambridge, Massachusetts, company whose business model is creating new ventures based on pioneering science. Flagship generally does not solicit business plans from independent entrepreneurs but instead uses internal teams of scientists to discover new-venture opportunities. The company has a formal exploration process whereby small teams of scientists, under the direction of one of the company's partners, undertake research on a problem of major social or economic importance—nutrition, for example. During these explorations, teams read the literature on the topic and engage the company's broad network of external scientific advisers to conceive new scientific insights. Explorations are initially unconstrained. All ideas—however seemingly unreasonable or far-fetched—are entertained. According to founder and CEO Noubar Afeyan, “Early in our explorations, we don't ask, ‘Is this true?’ or ‘Is there data to support this idea?’ We do not look for academic papers that provide proof that something is true. Instead, we ask ourselves, ‘What if this were true?’ or ‘If only this were true, would it be valuable?’” Out of this process, teams are expected to formulate testable venture hypotheses.

A willingness to experiment does not mean randomly throwing paint at a canvas.

Experimentation is central to Flagship's exploration process because it is how ideas are culled, reformulated, and evolved. But experimentation at Flagship differs in fundamental ways from what I often see at other companies. First, Flagship does not run experiments to validate initial ideas. Instead, teams are expected to design “killer experiments” that maximize the probability of exposing an idea's flaws. Second, unlike many established companies that heavily fund new ventures in the mistaken belief that more resources translate into more speed and more creativity, Flagship normally designs its killer experiments to cost less than \$1 million and take less than six months. Such a lean approach to testing not only enables the firm to cycle through more ideas more quickly; it also makes it psychologically easier to walk away from projects that are going nowhere. It

forces teams to focus narrowly on the most critical technical uncertainties and gives them faster feedback. The philosophy is to learn what you have gotten wrong early and then move quickly in more-promising directions.

Third, experimental data at Flagship is sacred. If an experiment yields negative data about a hypothesis, teams are expected to either kill or reformulate their ideas accordingly. In many organizations, getting an unexpected result is “bad news.” Teams often feel the need to spin the data—describing the result as an aberration of some sort—to keep their programs alive. At Flagship, ignoring experimental data is unacceptable.

Finally, Flagship’s venture team members themselves have a strong incentive to be disciplined about their programs. They gain no financial benefit from sticking with a loser program. In fact, just the opposite is true. Continuing to pursue a failed program means forgoing the opportunity to join a winning one. Again, compare this model with what is common in many companies: Having your program canceled is terrible news for you personally. It could mean loss of status or perhaps even your job. Keeping your program alive is good for your career. At Flagship, starting a successful venture, not keeping your program alive, is good for your career. (Disclosure: I serve on the board of a Flagship company, but the information in this example comes from a Harvard Business School case I researched and coauthored.)

Disciplined experimentation is a balancing act. As a leader, you want to encourage people to entertain “unreasonable ideas” and give them time to formulate their hypotheses. Demanding data to confirm or kill a hypothesis too quickly can squash the intellectual play that is necessary for creativity. Of course, not even the best-designed and well-executed experiments always yield black-and-white results. Scientific and business judgments are required to figure out which ideas to move forward, which to reformulate, and which to kill. But senior leaders need to model discipline by, for example, terminating projects they personally championed or demonstrating a willingness to change their minds in the face of the data from an experiment.

3. Psychologically Safe but Brutally Candid

Psychological safety is an organizational climate in which individuals feel they can speak truthfully and openly about problems without fear of reprisal. Decades of research on this concept by Harvard Business School professor Amy Edmondson indicate that psychologically safe environments not

only help organizations avoid catastrophic errors but also support learning and innovation. For instance, when Edmondson, health care expert Richard Bohmer, and I conducted research on the adoption of a novel minimally invasive surgical technology by cardiac surgical teams, we found that teams with nurses who felt safe speaking up about problems mastered the new technology faster. If people are afraid to criticize, openly challenge superiors' views, debate the ideas of others, and raise counterperspectives, innovation can be crushed.

We all love the freedom to speak our minds without fear—we all want to be heard—but psychological safety is a two-way street. If it is safe for me to criticize your ideas, it must also be safe for you to criticize mine—whether you're higher or lower in the organization than I am. Unvarnished candor is critical to innovation because it is the means by which ideas evolve and improve. Having observed or participated in numerous R&D project team meetings, project review sessions, and board of directors meetings, I can attest that comfort with candor varies dramatically. In some organizations, people are very comfortable confronting one another about their ideas, methods, and results. Criticism is sharp. People are expected to be able to defend their proposals with data or logic.

In other places, the climate is more polite. Disagreements are restrained. Words are carefully parsed. Critiques are muffled (at least in the open). To challenge too strongly is to risk looking like you're not a team player. One manager at a large company where I worked as a consultant captured the essence of the culture when she said, "Our problem is that we are an incredibly nice organization."

When it comes to innovation, the candid organization will outperform the nice one every time. The latter confuses politeness and niceness with respect. There is nothing inconsistent about being frank and respectful. In fact, I would argue that providing and accepting frank criticism is one of the hallmarks of respect. Accepting a devastating critique of your idea is possible only if you respect the opinion of the person providing that feedback.

Still, that important caveat aside, "brutally honest" organizations are not necessarily the most comfortable environments in which to work. To outsiders and newcomers, the people may appear aggressive or hard-edged. No one minces words about design philosophies, strategy, assumptions, or perceptions of the market. Everything anyone says is scrutinized (regardless of the person's title).

Building a culture of candid debate is challenging in organizations where people tend to shy away from confrontation or where such debate is viewed as violating norms of civility. Senior leaders need to set the tone through their own behavior. They must be willing (and able) to constructively critique others' ideas without being abrasive. One way to encourage this type of culture is for them to demand criticism of their own ideas and proposals. A good blueprint for this can be found in General Dwight D. Eisenhower's battle-plan briefing to top officers of the Allied forces three weeks before the invasion of Normandy. As recounted in *Eisenhower*, a biography by Geoffrey Perret, the general started the meeting by saying, "I consider it the duty of anyone who sees a flaw in this plan not to hesitate to say so. I have no sympathy with anyone, whatever his station, who will not brook criticism. We are here to get the best possible results."

Eisenhower was not just inviting criticism or asking for input. He was literally demanding it and invoking another sacred aspect of military culture: duty. How often do you demand criticism of your ideas from your direct reports?

4. Collaboration but with Individual Accountability

Well-functioning innovation systems need information, input, and significant integration of effort from a diverse array of contributors. People who work in a collaborative culture view seeking help from colleagues as natural, regardless of whether providing such help is within their colleagues' formal job descriptions. They have a sense of collective responsibility.

But too often, collaboration gets confused with consensus. And consensus is poison for rapid decision making and navigating the complex problems associated with transformational innovation. Ultimately, someone has to make a decision and be accountable for it. An accountability culture is one where individuals are expected to make decisions and own the consequences.

There is nothing inherently inconsistent about a culture that is both collaborative and accountability-focused. Committees might review decisions or teams might provide input, but at the end of the day, specific individuals are charged with making critical design choices—deciding which features go and stay, which suppliers to use, which channel strategy makes most sense, which marketing plan is best, and so on. Pixar has created several ways to provide feedback to its

movie directors, but as Ed Catmull, its cofounder and president, describes in his book *Creativity, Inc.*, the director chooses which feedback to take and which to ignore and is held accountable for the contents of the movie.

Accountability and collaboration can be complementary, and accountability can drive collaboration. Consider an organization where you personally will be held accountable for specific decisions. There is no hiding. You own the decisions you make, for better or worse. The last thing you would do is shut yourself off from feedback or from enlisting the cooperation and collaboration of people inside and outside the organization who can help you.

A good example of how accountability can drive collaborative behavior is Amazon. In researching a case for Harvard Business School, I learned that when Andy Jassy became head of Amazon's then-fledgling cloud computer business, in 2003, his biggest challenge was figuring out what services to build (hardly an easy task given that cloud services were a completely new space for Amazon—and the world). Jassy immediately sought help from Amazon's technology teams, its business and technical leaders, and external developers. Their feedback about requirements, problems, and needs was critical to the early success of what eventually became Amazon Web Services—today a profitable \$12 billion business run by Jassy. For Jassy, collaboration was essential to the success of a program for which he was personally accountable.

Leaders can encourage accountability by publicly holding themselves accountable, even when that creates personal risks. Some years ago, when Paul Stoffels headed R&D at Johnson & Johnson's pharmaceutical division, his group experienced a failure in a major late-stage clinical program. (Disclosure: I have consulted for various divisions of Johnson & Johnson). As Stoffels recounted at a meeting of J&J managers that I attended, senior leadership and the board demanded to know who was at fault when the program had its setback. "I am accountable," Stoffels replied. "If I let this go beyond me, and I point to people who took the risk to start and manage the program, then we create a risk-averse organization and are worse off. This stops with me." Stoffels, now chief scientific officer for J&J, shares this story frequently with employees throughout the corporation. He finishes with a simple promise: "You take the risk; I will take the blame." And then he urges his audience to cascade this principle down the organization.

5. Flat but Strong Leadership

An organizational chart gives you a pretty good idea of the structural flatness of a company but reveals little about its cultural flatness—how people behave and interact regardless of official position. In culturally flat organizations, people are given wide latitude to take actions, make decisions, and voice their opinions. Deference is granted on the basis of competence, not title. Culturally flat organizations can typically respond more quickly to rapidly changing circumstances because decision making is decentralized and closer to the sources of relevant information. They tend to generate a richer diversity of ideas than hierarchical ones, because they tap the knowledge, expertise, and perspectives of a broader community of contributors.

Lack of hierarchy, though, does not mean lack of leadership. Paradoxically, flat organizations require stronger leadership than hierarchical ones. Flat organizations often devolve into chaos when leadership fails to set clear strategic priorities and directions. Amazon and Google are very flat organizations in which decision making and accountability are pushed down and employees at all levels enjoy a high degree of autonomy to pursue innovative ideas. Yet both companies have incredibly strong and visionary leaders who communicate goals and articulate key principles about how their respective organizations should operate.

Here again, the balance between flatness and strong leadership requires a deft hand by management. Flatness does not mean that senior leaders distance themselves from operational details or projects. In fact, flatness allows leaders to be closer to the action. The late Sergio Marchionne, who led the resurrection of first Fiat and then Chrysler (and was the architect of their merger) commented to me during an interview for a Harvard Business School case I wrote: “At both companies, I used the same core principles for the turnaround. First, I flattened the organization. I had to reduce the distance between me and the people making decisions. [At one point, Marchionne had 46 direct reports between the two organizations.] If there is a problem, I want to know directly from the person involved, not their boss.”

At both Fiat and Chrysler, Marchionne moved his office to the engineering floor so that he could be closer to product planning and development programs. He was famous both for being detail oriented and for pushing decision making down to lower levels in the organization. (With so many direct reports, it was nearly impossible for him not to!)

Getting the balance right between flatness and strong leadership is hard on top management and on employees throughout the organization. For senior leaders, it requires the capacity to articulate compelling visions and strategies (big-picture stuff) while simultaneously being adept and competent with technical and operational issues. Steve Jobs was a great example of a leader with this capacity. He laid out strong visions for Apple while being maniacally focused on technical and design issues. For employees, flatness requires them to develop their own strong leadership capacities and be comfortable with taking action and being accountable for their decisions.

Leading the Journey

All cultural changes are difficult. Organizational cultures are like social contracts specifying the rules of membership. When leaders set out to change the culture of an organization, they are in a sense breaking a social contract. It should not be surprising, then, that many people inside an organization—particularly those thriving under the existing rules—resist.

Leading the journey of building and sustaining an innovative culture is particularly difficult, for three reasons. First, because innovative cultures require a combination of seemingly contradictory behaviors, they risk creating confusion. A major project fails. Should we celebrate? Should the leader of that program be held accountable? The answer to these questions depends on the circumstances. Was the failure preventable? Were issues known in advance that could have led to different choices? Were team members transparent? Was there valuable learning from the experience? And so on. Without clarity around these nuances, people can easily get confused and even cynical about leadership's intentions.

Second, while certain behaviors required for innovative cultures are relatively easy to embrace, others will be less palatable for some in the organization. Those who think of innovation as a free-for-all will see discipline as an unnecessary constraint on their creativity; those who take comfort in the anonymity of consensus won't welcome a shift toward personal accountability. Some people will adapt readily to the new rules—a few may even surprise you—but others will not thrive.

These cultures are not all fun and games.

Third, because innovative cultures are systems of interdependent behaviors, they cannot be implemented in a piecemeal fashion. Think about how the behaviors complement and reinforce one another. Highly competent people will be more comfortable with decision making and accountability—and their “failures” are likely to yield learning rather than waste. Disciplined experimentation will cost less and yield more useful information—so, again, tolerance for failed experiments becomes prudent rather than shortsighted. Accountability makes it much easier to be flat—and flat organizations create a rapid flow of information, which leads to faster, smarter decision making.

Beyond the usual things that leaders can do to drive cultural change (articulate and communicate values, model target behaviors, and so on), building an innovative culture requires some specific actions. First, leaders must be very transparent with the organization about the harder realities of innovative cultures. These cultures are not all fun and games. Many people will be excited about the prospects of having more freedom to experiment, fail, collaborate, speak up, and make decisions. But they also have to recognize that with these freedoms come some tough responsibilities. It's better to be up-front from the outset than to risk fomenting cynicism later when the rules appear to change midstream.

Second, leaders must recognize that there are no shortcuts in building an innovative culture. Too many leaders think that by breaking the organization into smaller units or creating autonomous “skunk works” they can emulate an innovative start-up culture. This approach rarely works. It confuses scale with culture. Simply breaking a big bureaucratic organization into smaller units does not magically endow them with entrepreneurial spirit. Without strong management efforts to shape values, norms, and behaviors, these offspring units tend to inherit the culture of the parent organization that spawned them. This does not mean that autonomous units or teams can't be used to experiment with a culture or to incubate a new one. They can. But the challenge of building innovative cultures inside these units should not be underestimated. And they will not be for everyone, so you will need to select very carefully who from the parent organization joins them.

Finally, because innovative cultures can be unstable, and tension between the counterbalancing forces can easily be thrown out of whack, leaders need to be vigilant for signs of excess in any area and intervene to restore balance when necessary. Unbridled, a tolerance for failure can encourage slack thinking and excuse making, but too much intolerance for incompetence can create fear of risk

taking. Neither of these extremes is helpful. If taken too far, a willingness to experiment can become permission to take poorly conceived risks, and overly strict discipline can squash good but ill-formed ideas. Collaboration taken too far can bog down decision making, but excessive emphasis on individual accountability can lead to a dysfunctional climate in which everyone jealously protects his or her own interests. There is a difference between being candid and just plain nasty. Leaders need to be on the lookout for excessive tendencies, particularly in themselves. If you want your organization to strike the delicate balance required, then you as a leader must demonstrate the ability to strike that balance yourself.

A version of this article appeared in the January–February 2019 issue (pp.62–71) of *Harvard Business Review*.

Gary P. Pisano is the Harry E. Figgie Jr. Professor of Business Administration and the senior associate dean of faculty development at Harvard Business School. He is the author of *Creative Construction: The DNA of Sustained Innovation*.
