

Lecture Text

David A. Garvin: Creating Competitive Advantage Through Organizational Learning

(edited for clarity)

Introduction

Let me say a few words about what we plan to do today. I'm here to talk to you about putting the learning organization to work, which is really a fancy way of saying, how do you help your organizations improve? How do you get better at whatever it is that's important to you?

I'm going to fairly quickly set the stage for you. And you should have, at your desks, copies of all of the slides and all of the transparencies. Paul Valéry is a French symbolist poet. And, he, in many ways, captures the challenge that many of you face in your organizations, whether they've been around for eons or whether they're relatively new organizations.

Valéry says that "The trouble with our times is that the future is not what it used to be." And just think about your own experiences. Particularly in the post-war period, it was relatively easy to be successful, with a vast pent-up demand for products and services, very well-defined industries and industry boundaries, and a pretty clear competitive game based on long-established business models.

An unpredictable future

The world today is very different. There are collapsing industry boundaries. It used to be true that, if you were in commercial banking, your competitors were commercial banks. Today, there's a world of financial services. It used to be that if you were in software, you competed in software. Today, in software, you compete with Disney. Disney makes products of that sort. So, the boundaries are no longer clear. At the same time, there's vast excess capacity in lots of businesses, everything from the traditional—petro chemicals—to the untraditional—Internet portals or dot-com retailers. So, again, you have to be much more efficient than you used to have to be.

Many of the industries that people went into right out of business school were regulated, for instance, electrical power. The world of electrical power has been deregulated. You saw what happened in California. There was a very different set of challenges. The result is that the future, which was relatively stable and relatively secure, is now quite different: very tumultuous, very uncertain, highly unpredictable.

Now there are a variety of ways organizations respond. The one that I am going to argue is really the competitive edge and where I want to spend most of the time here is learning. Now think about it. Products can often be copied. Even if you get a patent, it's only good for a certain number of years. Services can even more easily be copied. Even processes of various kinds can be copied. For those of you who have large manufacturing-based businesses, you often use enterprise integration software, these complex packages. They're now available on the open market. You buy them from SAP, or others, as the case may be.

And even some processes that people think are relatively proprietary in their advantages can be copied. Think, for a moment, of Six Sigma. Six Sigma is actually an idea that's been spread from company to company to company. It originally came, people thought, from GE. In fact, the original design came from Motorola many, many years ago, who passed it on to Allied Signal who passed it on to GE. Now there are books on the subject, so you can get that advantage.

The solution: learning

The only advantage that's left is learning. Ray Stata was, for many years, the CEO of Analog Devices, a very, very successful semiconductor manufacturer. He says that "The rate at which individuals and organizations learn may well become the only sustainable competitive advantage."

I want to give you a few concrete examples to make the point, because there is a link here with bottom-line performance, and it's important to establish that link. Learning is not something you do just because it's fun to do, but also because it has performance implications.

The first example comes from George Fisher. George Fisher was, for many years, president of Motorola. He later went on to head Kodak. When he was running Motorola, he was once asked to explain the power of learning and improvement. And he said, "Well I'll give you a contrasting set of examples. I want you to think of semiconductors. Then I want you to think of the auto industry." He said, "If the auto industry had learned, had improved, at the same rate as semiconductors, a Cadillac, giant American car that twenty years ago cost \$10,000, would today cost eighty cents." Eighty cents. He said, "That, by the way, is the good news. The bad news is it would also be two inches long." That's the rate of learning. That can't be duplicated any other way.

Now that's a bit of a flip example, so I want to give you a fairly concrete one. It turns out that in the 1970s, virtually every company that bought a continuous caster—a gigantic two- and three-story machine that pours molten steel; there were about thirty of them bought around the decade, around the world—bought one manufactured and sold by a single Swiss company. Same technology. Little difference in features and bells and whistles, but not very much.

And a bright academic had the idea of looking at these machines, and trying to understand the differences in the start-up periods. The start-up period is the time between your installation and first pouring of molten steel and the time when you reach standard or full production. Virtually identical machines. Those start-up periods range from seven-and-a-half months to six years. The average was about 24 months. It was about two years.

Just think about those numbers for a minute. You can't explain that difference by technology. Essentially, it's the same piece of equipment. You can't explain that difference by the sources of supply. The same Swiss company. You can't even explain it by on-site engineering support because it's the same engineers from the Swiss company helping out. The only way you can explain the difference is by different learning rates at the companies that were putting in the equipment.

Now, think about anything that you're putting in your organization: a new machine tool, a new piece of equipment, even a new process for an advertising campaign, for creative development. If you can do that more rapidly and more efficiently than your competitors

can, then you have an edge. And if you can keep doing that with new equipment, new processes, even new people whom you're bringing on stream, you have an economic advantage that you can trace to learning. That's the power of learning.

Now I want to give you one more example. This one is a little hard to follow because it's a complicated chart. So, let me explain to you what you're seeing, and you also have a copy in your notes. First of all, all of this information is from one company. It's from Analog Devices, that semiconductor manufacturer. And it's all from one year, as you see at the top, August 1987 to July 1988. A, B, C, D, E, F, and G are simply different divisions making roughly similar products—they're all semiconductors—at this company.

Each dot, which is a data point, represents the percentage of product lines delivered late. So lower is better. You want fewer product lines delivered late. Okay. Let's do a quick quiz. This is the only test that you're going to face today. Which division president do you want to be? A, B, C, D, E, F, and G.

___: C.

DAVID: C? C clearly has the best performance. Now, here's the punch line. Look where C started. C was among the two poorest performers in this company at the beginning of the year. G was the other. But notice its learning rate—and what you see at the bottom is a measure of the learning rate, how long it takes to make a 50 percent improvement. Takes division C four months. Takes division E more than five years. And division D will shortly have a new division president. They're going in the wrong direction.

Generalize this to your own businesses. Pick any measure you care about: customer satisfaction, productivity, cost, quality, anything you care about. If you're C, and your competitors are A, B, D, E, F, or G, you can get ahead, even if you're starting behind. Even if you're at a disadvantage, you can get ahead and stay ahead because of the power of learning.

So that's the rationale. That's the economic argument that says you pursue learning for very sound business reasons, not just because it's something that stimulates the workforce or is fun to do. Now, any comments or questions before we move on? Okay.

What is a learning organization?

Now, what I want to do now is actually define for you a learning organization. And this is where we're going to pick up knowledge management for the first time. A learning organization is an organization skilled at two things. First, it's very good at creating, acquiring, interpreting, transferring, and retaining knowledge. And it's equally good at purposely, with a reason, with a business reason, modifying its behavior to reflect the new knowledge and insights.

Okay. Let's take this one at a time. Creating knowledge. That's what R&D labs do all the time. That's what truly creative employees do all the time. As you'll see, there are a number of things you can do to encourage people to be more creative on the job. And we'll get to that in a minute. It typically has to do with whether you're providing them outlooks and opportunities, removing certain risks—this is what I meant by changing incentives—or simply stressing the systems in ways that force people to think more creatively. And we'll get to all those examples in a little while.

The second part is that you can actually acquire knowledge. This is what people do when they license from other organizations. It turns out, though, that there's another way you can acquire information. And that's through mergers and acquisitions. Cisco was previously a darling; now it's in the dumps. But while it was growing, Cisco's strategy was to acquire new knowledge, particularly technological knowledge, by acquiring small, growing start-ups. And then they integrated those people, their systems and their thinking, into the Cisco organization and culture. That was R&D via acquisition.

The third thing you've got to do, as a learning organization—and this one is particularly hard because it's often not understood—is interpret what you get. You've got to make sense out of it. And here, I want to give you a very precise example.

This is the Intel® Pentium® Processor, and they had a problem when they first launched that chip. You may remember, when they first launched the Pentium® Processor, it had a floating point calculation problem once every nine billion calculations. Now, I don't know how many use a calculator that intensely. But, it doesn't come up in my day-to-day affairs as a real issue.

And Intel, frankly, was convinced it was not a real issue. But they diligently gathered customer feedback. And then, for four months, they failed to act. And it's fascinating. If you read Andrew Grove's biography—Andrew Grove was the CEO at the time—he has this wonderful description of the period. He says, "It's like sailing a boat when the wind changes. You know you have to tack. You know you have to change direction. But it's not clear which way you're supposed to turn."

Now what had happened to Intel® was that Intel® had largely been an industrial goods company. They sold their chips to computer manufacturers. But they had advertised intensely, for months and months, around the Intel Inside® advertising campaign, without them realizing that they had shifted over to become a consumer goods company, held to very different standards. It took them four months to figure that out. The problem was not acquiring the information. They had all the data. It wasn't having meetings. They were meeting weekly and, eventually, daily about the problem. It was making sense out of what it meant. So a learning organization is able to shift its mental frameworks, shift the filters it uses, as the organization itself evolves and changes.

Now last two. Transferring and retaining. This is where knowledge management comes in. And it comes in because there's a big difference between having an organization that learns and having individuals that learn. Simple example. You have two crackerjack salespeople who have managed to crack the biggest accounts. If they walk out the door, and they take their knowledge with them, you are not a learning organization. You have had individuals who are learned and knowledgeable, but you don't have it embedded in the organization fabric.

Now there are a variety of ways to do it. One is knowledge management systems. The other is to build it into the fabric of the culture or to develop well-grooved processes for getting things done. I'll give you a very specific example.

I've spent a fair amount of time in the last year studying Enron. Enron is an energy giant that is renowned for innovative skills. They are continually building new businesses. Enron is always worried about losing key people. But as it turns out, that's not where the innovative capability lies. If you talk to new hires, within about six months they can tell you exactly

what steps you have to follow to get a new business approved. Those steps are not written down. But everybody knows the sequences of steps. They have become a learning organization when it comes to innovative capability.

So you have to retain it. You also have to transfer it. The worst thing in the world is to have a terrific skill in Factory A, and have Factory B, a hundred miles away, not have that same skill, even though they use the same equipment and the same technology. Again, you will have local learning, but not organizational learning. So that's the first half of the definition. Unfortunately, that's the easy half.

The really hard part is changing your behavior to reflect the new knowledge and insights. Now there are two elements here. And this gets to Steven's question. The first is simple recognition of the fact that there is a disconnect between what you know and what you do. And here, I'll give you a very personal example.

About a dozen years ago, I had done a lot of research on quality management. And at the time, I was teaching in our Advanced Management Program, in the AMP. Now, you have to understand, the AMP don't have quite the hallowed view of your faculty that you MBAs do—they're just a bit less respectful than you folks are.

So I was in the middle of one of my lectures and all of a sudden a student shouted—he actually didn't raise his hand; he just shouted it out—"Excuse me, David. Excuse me. But, do you folks at Harvard Business School actually practice quality management?" Exactly. I hemmed and hawed, and finally said, "Hell no! We just teach it." And that was the shock of recognition.

We had met all of the first tests. I created new knowledge. I had done ten years of research. I had interpreted lots of raw data from companies from the shop floor. I had transferred it. I had written articles, books, cases. And we had retained it. Because I had trained additional faculty. And we had it in our syllabus. But we hadn't consciously changed our own behavior in the way we ran our programs. And shortly thereafter, we actually conducted a training program for ourselves, including faculty administrative staff, to begin to develop more of a customer service orientation. That's when we made the shift.

So the first barrier that needs to be overcome is that people simply have to see that there's a disconnect between what they know and what they do. Now the other barriers we'll get to in a little while, because they have to do with incentives and a willingness to take risks, which are huge barriers in most organizations. Okay. So that's the definition.

Now I want to spend most of our time talking about practice. We've spent some time understanding why learning is important. Spent some time understanding a definition of a learning organization. And you could apply that definition to your own organizations. Very easy to apply.

Six critical tasks

Now let's talk about what these organizations do. Okay. They do six things. First they collect intelligence about the environment, typically about customers, about competitors, and about trends: social trends, economic trends, technological trends, regulatory trends. Second, they learn from the best practices of other companies. It's just a fancy way of saying they benchmark.

Third, they learned from their own experiences and past history. This is one that few organizations do. Typical management philosophy is, if it's a victory, declare it, pat yourselves on the back, and move on. If it's a failure, bury it, declare victory, and move on. Unfortunately, you don't learn from that process. What you need to do is review both your successes and your failures, and also the techniques for this, which we'll go over in a few minutes, so you can do it more effectively.

You experiment with new approaches. It's a fancy way of saying you try new things. You encourage systematic problem solving. Key word there is systematic. And finally, you transfer knowledge throughout the organization, both through systems, but also through personal contacts.

Three important conclusions about this list. The first is that all six of those items are processes. What that means is that you can systematically understand the steps. And there are better and worse ways to do it, just as there are with almost any process.

Second, all of your organizations are already doing these things. You wouldn't be in business if they weren't. All I can do is give you some best-practice examples that suggest that some of these things are things you probably can do better. And third, I am going to give you technique. But the reality is that learning is as much art as science. It's as much temperament as technique. All I can do is tell you, even the great artists work on their technique. But recognize that there are some limitations to what we're going to be talking about here. Okay. Now let's actually start running through these in some detail.

Intelligence Gathering

First: intelligence gathering. To begin, it's worth recognizing that there are different ways that you collect intelligence. Search is what you do when you go on the Internet, and somebody says, "Look up the financials for company XYZ." The information already exists. The only challenge is finding it, and then verifying it to make sure it's true.

The second approach is inquiry. Inquiry is just a fancy way of saying question and answer. This is what most market research is all about. You go out, and you ask somebody something because the information is in their heads. As it turns out, that always doesn't work. And the reason it doesn't always work is because, often, people can't or won't tell you what really is going on.

Then what you need to do is to rely on direct observation. And I'll give you a very concrete example. This is one of my favorites. This comes from Best Western Hotels. Best Western was trying to understand how decisions were made by seniors, by couples fifty-five, sixty, sixty-five, and older, when they were driving across country, how they chose to get off the highway and choose a hotel. So they were interviewing about thirty couples who were making this trip. But they also got these couples to agree to allow them to put a video camera in the back seat of their cars to film the conversations.

So, typically it goes like this. They'd sit down with the couples at the end of the day. And they'd say, "So how did you two decide when you were going to pull off the highway?" And the husband would instantly pipe up, "Well I make that decision." Okay. Say, "Fine. That's great. Now how did you decide which hotel to choose?" And the husband immediately piped up, "Well I make that decision, too."

Now, in this case, they actually had an observational record. They had this tape. So then they played the tapes. And, as you might expect, in about 70 percent of the cases the wife made the decision about when we get off the highway and which hotel we choose to go to. This is not uncommon. And it's true in more traditional business settings as well.

Xerox, for instance, at one point, was trying to figure out how office clerks use copying machines. So they asked them. And they gave answers that bore a striking resemblance to the manual of procedures. No great surprise. But then they actually brought in anthropologists who, through a combination of direct observation and videotape, actually observed the way these people worked.

It turns out that their work bears no relation to stated procedures. People do end runs. They use the machines at the wrong time. They make fifty copies when they're supposed to make ten. Exactly what we do. Then they understood how they needed to design for real use.

So, the first message in terms of guidelines is to try to collect data from diverse sources. Try to triangulate, to combine evidence that you search for with evidence that you get from question and answer with direct observation. Second, cross-check the findings to ensure that they are valid. It's very dangerous to rely on a single source. And here, this is my excuse to tell you one of my very favorite stories about the dangers of relying on a single source. This is a Harvard story. But it's a Harvard story from across the river.

For many years, at the College, they used to give examinations in a place called Memorial Hall. Memorial Hall has since become a cafeteria. But at the time, it was this gigantic room, about five, ten times the size of this room. And they would literally give examinations to about a thousand students at a time. Students would be lined up at tables, very long tables, by course. So the first table might be Economics 10. The second table might be English 20. Third table would be Physics 55, and so on.

You take your exam. The undergraduate exams are four hours long. At the end, you would come up to the front of the room, and you would pile up your exams. A pile of blue books would be about this high, in front of a proctor. The proctor had the single master list which listed every student by name and by course. And he would pick up the exams one by one, turn them over, and check off the students. He's an inspector.

Now, this being Harvard College, as you might expect this is a terribly important job. The fellow who's been doing it at the College has been doing it for twenty-five years. It's one of our doctoral students who never quite made it to the real world. And he's nicknamed Mr. Test. And apparently, Mr. Test was proctoring one of these exams a few years ago. And he goes through the process. At the end, he announces, "Time's up." And a thousand students come to the front of the room and they pile up these exams. And he begins this laborious process, one by one, of checking off the exams.

About twenty minutes go by. And he looks up. And he suddenly notices that way in the back of the room there's a student still writing. This guy is still working on his exam. This goes on for about five more minutes. Then the student says to himself, "I think I'm done now." Walks to the front and says to the proctor, "I'd like to turn in my exam now." The proctor is absolutely incredulous. He says, "I can't accept that. It's twenty-five minutes late. Look at this pile I have." And understand, this is a Harvard undergraduate.

The student rears up. He says, "Do you know who I am?" Now this is Mr. Test. He's been around the block. He's not very impressed. He says, "No." Student says, "Good." Takes his exam. Stuffs it in the middle of the pile, and runs out. I tell you that so you never again rely on a single, unverified source to collect your findings.

Now, some other things. Obviously, you need to avoid biased, slanted questions. They come in various forms. The most dangerous ones are where you lead the witness. Don't go in and say, "Gee, do you prefer our potato chips for their taste or for some other reason?" You're not going to get very good responses. Open-ended questions, which don't dictate a response, are far more powerful.

Even more important is what appears to be this innocuous statement: keep an open mind. It turns out that we actually have very good studies of how managers use market research. And they use it exactly the wrong way. People typically use marketing research to confirm things they believe they already know. Well if you're really using it for intelligence gathering, what you should be doing is exactly the opposite, attending to the unexpected, attending to surprises. It turns out that when most managers get surprises in marketing research, they discount them. They push them under the rug because that's not what they want to hear. You're not going to get very good learning from a process like that.

And finally, recognize that you can do this process in lots of different ways. Xerox—this is way back during its glory years—actually did this every three to five years, trying to look ten years ahead. So Xerox '92 was done in '82. Xerox '95 was done in '85. Very elaborate, multi-day planning exercises to project where the world was going. They don't have to be that complicated, particularly in smaller companies.

DuPont, a large company, does something that any of you could do tomorrow. Chad Holliday, who is the CEO of DuPont, every two weeks has a telephone call with managers around the globe. The purpose of that call is not to make decisions. It's to share information on markets, on competitors, on customers. So he very much leads it like a seminar. What are you hearing in Latin America? Are there any trends there that you folks over in Asia are also picking up as well? What about competitive action? What's Dow doing in these markets? Are you seeing that as well? And it's a desire to collect a shared view of the world.

GE, the Major Appliances Division, considers it so important that they actually conduct this call weekly with about a dozen or two dozen salespeople. And the reason is that, if you conduct it every week, the most you can get behind is one week. A very powerful form of intelligence gathering. And it's very cheap. Doesn't have to be expensive. Okay. Any comments or questions right now? Just in order to take a pause, just to make sure we're covering what you need to hear. Okay.

Learning from other companies

Now, learning from other companies. As I said, this is a fancy way of saying benchmarking. These are some quotes from a leading book on benchmarking. "Benchmarking is the search for industry best practices that will lead to superior performance. Those practices are to be pursued regardless of where they exist, in one's own company, industry, or outside one's industry."

The key word there is practices. All too much benchmarking is an attempt to look at performance, to look at the numbers somebody else has. But not enough effort is made to figure out how they got those results. The experts tell me that 90 percent of the value of benchmarking comes from understanding somebody else's practices. Only 10 percent comes from understanding their performance numbers. And you need those numbers to know who is superior. But that's just the starting point.

Second, you can apply benchmarking "to all facets of the business: basic products and services, processes that go into manufacturing those products, and the methods that support getting them to customers and meeting their needs." You can benchmark anything. Any activity that you do, especially routine activities where you have a well-defined process, you can compare to somebody else. In fact, my favorite example is this one: at one point Milliken, the textile company, benchmarked Xerox's approach to benchmarking. It's a process. And as a process, you can study it, and you can compare it.

And then, finally, here's the connection to learning. "In the formal sense, benchmarking is no more than an ongoing investigation and learning experience that ensures that best industry practices are uncovered, analyzed, adopted, and implemented." It's worth pointing out that there are two quite different approaches to benchmarking, both powerful, both equally effective, but they serve very different ends.

The one that most people are familiar with is benchmarking as copying. You go to another company. You study their billing practices. You study their accounting systems. You study their IT processes. And your goal is to take those and lift them up, transfer them over to your organization with little or no change. You want to copy it as closely as possible.

There's a second approach to benchmarking that you might call benchmarking as a catalyst, as a vehicle for stimulating creative thinking. The best example I know of here comes from Mobil Oil. Mobil Oil, at one point, was trying to improve the speed and quality of service at their service stations, at the gas stations. They put together two teams: a Speed Team, to study speed, and a Smiles Team, to study customer satisfaction and customer service.

What's revealing is whom they chose to benchmark. The Speed Team benchmarked Roger Penske's Indianapolis 500 race team. Now, I don't know how many of you have driven your car into a service station lately. But I have never been able to get in and out in 5.5 seconds, which is what a pit crew does at the Indianapolis 500. You can't do it.

And, in fact, that was never the goal of the benchmarking team. It was to come up with ideas that might stimulate their thinking. So for those of you who know Mobil Oil, they now have something called the Speed Pass. A Speed Pass is essentially a pre-paid credit card. It gets you in and out rapidly. That idea came out of this process of benchmarking. At one point, they even experimented with a separate lane, a speed lane, for road warriors. That came out of the idea that the pit crew works off the race track. So, you can use it to stimulate creativity.

In the case of the Smiles Team, they benchmarked the local Ritz Carlton Hotel. Obviously, they're not going to get that level of service quality at a gas station. But they learned some things about training, and bringing in new employees, that they could level down a little bit, and then apply in their own organizations. So that's benchmarking.

Now, here's where we pick up what Steven and Ed were talking about, about barriers and resistance. And it has to do with how you get people to change. And, in particular, how do

you get them to learn from their past experience, both good and bad, and change their behavior?

Learning from past experience

Now first of all, very few companies take the time to review their own past experience, to assess it systematically, and to record it in a form that's useful and accessible to other parts of the organization. I want to highlight two examples here because they are very relevant to these questions. The first is AT&T and the second is the U.S. Army.

What AT&T did, very much akin to the kinds of issues you're facing, is that they found, within Bell Laboratories, that there were vastly different levels of productivity among software engineers. Some were just better than others at getting the job done.

First, they identified those engineers. And it was interesting how they had to do it. They asked management, who are the most productive engineers? And they asked the engineers, who are the most productive engineers? It turns out that there was only a 50 percent overlap between the people nominated by the two groups. They took that 50 percent.

And then they did something very clever. They interviewed a select set of that 50 percent and a control group of average performers. And they interviewed them around education, around training, around experience, around intelligence and, most important, around how they got the job done.

What they found was that the superior engineers were using different work practices. In particular, they were building a network. And they were using the network to pull in information and send information out. They then set up a very simple training program, about an hour a week for several weeks, where the superior engineers designed the program and taught the other members of the staff.

Within about six months, they had improved the productivity of the people in that class 200 to 300 percent, compared with a control group of engineers who didn't go through the program. You could do that with any kind of superior employees, with salespeople. It's reflecting on what is done well and what is done less well by the members of the organization, and using that experience.

Now the other example comes from the U.S. Army. And the U.S. Army has a process called After-Action Reviews. It's a way of getting people out of the rut, out of the old way of doing things, and into new ways of doing things. So what they will do is, after any activity, any mission of any sort—originally, this was done in training, but, after Desert Storm, and also in Bosnia and Haiti, it became a regular part of every mission—they will debrief within a few hours around four simple questions.

First, what did we set out to do? What was our objective? And this can be everybody from the foot soldier to generals, because they run these reviews across all levels. Second question, what actually happened? Third question, obviously, why was there a difference? And the fourth question, what do we do next time, which is broken up into two parts. What activities do we sustain or continue to do? What activities do we improve or try to do differently?

Now just think for a minute. This is actually an extraordinarily thoughtful design. First question, what did we set out to do? In business, very seldom do people ask that question. The assumption is that everybody knows what we set out to do. The reality is typically quite different. People don't have a clear understanding of the objectives. It's true at all levels. They call this, in the Marines, Commander's Intent. And it's the ability for people to act on the fly when things happen in a way that corresponds to what their leader really wanted them to do.

And invariably, they find, when they stop time, that the commander wished they had gone this way, and they went that way. And he said, "Wasn't that clear from my instructions?" And the answer almost always is, "No, not at all. You didn't give us enough in the way of details."

Second question, what actually happened? Almost never do you see an organization actually ask this question. The assumption is that we all have the facts. The reality is that very seldom do you have shared facts. You need to document it, have a third party observer. Now this gets to the question of how you use outsiders. One of the most powerful things you can do is bring in a third party to help in the observation process.

The Army actually uses people called Observer Controllers. On critical missions, they will accompany the team. And their primary role is to help answer the question, what actually happened? So when there's a disagreement about facts, somebody says, "Wait a minute. We took the hill at 05:00." And somebody else says, "No, no, no. We took it at 06:00." He can say, "Excuse me. You took it at 05:30. I was watching my watch." No more do we have disputes over the facts.

Third question is pretty obvious. Why the gap? Fourth question is very, very clever. What do we do next time? Suppose the only part of that we answered is what do we improve? Nobody wants to engage in an exercise where all we do is find out what we did wrong. But if we're also going to ask, what do we sustain or continue to do, it's a discussion about what went right as well as what went wrong. Those are very powerful processes.

Now, the solution here, though, is not just a process solution. It's a mindset solution. It has to do with the approach and the culture that you build in the organization. "At the core of effective organizational learning is a mindset. And it's the mindset that distinguishes between productive failure and unproductive success." Productive failure is one where you learn from it. "You get insight and understanding in addition to the commonly held wisdom of the organization. And unproductive success occurs when something goes well, but nobody knows how or why." And I want to develop this some because it's directly pertinent to how you get people who are set in their ways to change.

First, let me say what I'm not saying. I always have this fear, at this point, that somebody comes back from the reunion and says, "Well, I heard this speech by David Garvin. And he told me to screw up a few." That's not what this says. What this says is actually very well represented by an old proverb. The proverb goes, "Good decisions come from wisdom, knowledge, and experience." And wisdom, knowledge, and experience come from bad decisions. We learn from our mistakes.

Now, here's the important implication. You have to create an environment that is accepting and tolerant of certain kinds of well-thought-through, well-intentioned mistakes. People make mistakes, particularly if you're asking them to try something new. If you immediately

shoot the bringer of bad news who was trying experiments, you're not going to get many experiments.

And there's a very important distinction. There is nothing wrong with making a mistake. There is something wrong with repeated mistakes because it shows a failure to learn. So in the Army, somebody once asked me, "How do they get people to talk?" The way they get people to talk in these reviews is to hold anything you say in an After-Action Review separate and confidential, and apart from performance evaluations.

So somebody once asked, "Yeah. But, how do you weed out the bad performers?" And I actually had a guy in one of my classes who was an Observer Controller, who played this role in the Army. And he said, "I want to make this distinction. We do not punish people for mistakes. We do weed out people who are unwilling or unable to learn. That's the difference." And that's a very hard culture to create.

Now, some of the things you can do. One is to make it clear when we're in an experimenting mode. And make it clear that, when we're experimenting, there will be a higher acceptance for deviations, and things not gone right. Second, consider working actively to change the incentive system.

And I'll give you a very formal example and a cultural example. The formal example comes from a steel company, Allegheny Ludlum Steel. Allegheny Ludlum Steel has people who run as much as twenty or thirty shop floor experiments at a time, routinely. They run hundreds over the course of the year.

Allegheny Ludlum also has the tightest control system I have ever seen. As an Operating Manager, you are evaluated on your variances against standard for cost, quality, productivity, for every grade of steel on every piece of equipment. Why would you possibly want to experiment? Because if you sign an experiment and get it approved—it requires three signoffs—it goes off budget. It's paid for from a separate account. And it doesn't count in your variances. Okay. We just removed the risk.

Now, what's the opportunity? Why do you do it? Because at the beginning of the year, every single one of the top 103 managers of Allegheny Ludlum elects certain improvement targets that they hope to meet in their area that are tied to bonuses at the end of the year. Now we have the system set up. Affirmative reasons, you want to improve. You want to run experiments because you're going to get paid if you improve. And the risk, the penalty, has been minimized and eliminated because it's not against variances. But there's still a check and balance because somebody has got to sign off on the experiments to approve them. That's one of the ways you get people to become more creative, by removing risks.

I'm going to give you one more example because I think this is so important. Has to do with the issue of supervisors. It's a very successful company, again a steel company, called Chaparral Steel. And they want people to become more comfortable in pushing their frontiers, and trying new things. Particularly, they want to develop new supervisors.

Chaparral has a policy called Vicing. Suppose the supervisor on your shift is ill. The next ranking person—it's a very flat organization; it's actually quite a bit lower experience—temporarily plays the role of supervisor. But the supervisor from the previous shift stays over on that shift as a temporary vice supervisor. What did you just do? You just created a safety net. You put the people in the role. You give them the opportunity to learn. But you

give them somebody to fall back on just in case. A very powerful way of building and learning, but also reducing the risk to the individual, as well as to the organization.

So that's successes and failures.

Experimentation

I'll say a little bit about experimentation. Experimentation requires this continual flow of new ideas, whether you're trying to discover something new, invent something new, or just come up with a new way of doing things. We talked a little bit about the incentive system that favors risk-taking. But how do you get new ideas? I'll give you two examples.

One of them is through stressing the system. My favorite example here comes from a very small Japanese soy sauce manufacturer. They have something they call the cutting-in-half game. Take a production line. Production line has thirty people. The supervisor will say at one point, "Let's see tomorrow if we can run that production line with fifteen people." Literally cuts the number in half.

And they struggle. And they try. And they run experiments. And the next day they go back to thirty. And they periodically make the cut. Now why is he doing that? If you ask him to run the line with twenty nine or twenty eight, they will just work faster. You cut it to fifteen, and you have to come up with wholly new ideas. On one of their lines, on a bottling line, through doing this repeatedly, they actually managed to cut the line from thirty people to eighteen.

It was very interesting what they did. It turned out it was thirty people. They are line people. The eighteen are fifteen line people and three roving troubleshooters. They had to invent a new job. Once they figured out the new job, they figured out a way to run the line with fewer people. That's one way.

The other route is through opportunity, to provide people a way of trying new things. Make it fun. Make it relatively risk-free. My favorite example here comes from Disney. Disney, in the feature animation area, the ones that make animated films, learned something called a Gong Show. Now, for those of you who don't know American variety shows, in this one often you get people up on stage, not very talented. They would do their song and dance. They'd sing. They'd perform. But just in case they weren't very good, there would be a gigantic metal gong, basically a big cymbal off to the side. And after a few minutes, if people weren't doing very well, somebody would get up and hit the gong. And that was your signal that it was time to get off stage because you weren't doing very well.

Well Disney does this with new ideas. Every now and then, it's typically once or twice a year, they will have a Gong Show. And what happens in the Gong Show is that anybody, at any level in the organization, can come and make a pitch for a new idea for a new animated film. The maximum time limit is five minutes. As you might expect, very few people get a full five minutes because there's this gong on the side.

Here's the punch line. Who attends the meeting? Three people. Chairman and CEO. That's Michael Eisner. Vice chairman. For many years, that was Roy Disney. And the president of the Feature Animations Unit. Over time, this process produced, among other things, the idea for the film "Hercules." "Hercules" made Disney \$100 million. The person who proposes the idea gets the standard fee for a pitch, for an idea, which is \$20,000. That's a pretty

good return on investment. All it does is give people a vehicle for coming forward with their ideas. A much more creative vehicle than the standard suggestion box system. This is much more visible, and a lot more fun.

Now, you can carry out this kind of experimentation process in lots of different settings. One of my favorites is one that was undertaken at Timken. And I particularly like it because it's so cheap, and yet it has a terrific set of ideas embodied in it. Timken was about to design a wholly new factory, with a completely new kind of technology. And there were big questions about where to put the equipment, where to put the machines. They were worried about clearances, but they were also worried about the cost of the factory, which is going to run tens of millions of dollars.

So, they built something called Cardboard City in a big warehouse. They built full-scale mockups, full-size mockups on all the pieces of equipment. They made them out of cardboard and 2x4s. And they put them on wheels. Then they put together three teams, and they made it a competition. Each of the teams had operators and engineers. And they said to them, "Put the equipment"—remember this is cardboard—"Put the equipment where you think it should be." They've got a few rules about machine times, machine sizes, clearances, but it's up to you.

And then they brought in real operators who simulated production. So, I have these videotapes—they're incredible videos—of these big cardboard contraptions. And these guys are walking along as if they were holding bearings, which are actually pieces of paper—you know, going to the machine, holding them to the machine, it's a six second cycle, count 1-2-3-4-5-6, and then they pass other operators, and they hand off to them.

Well, it turned out that not one of the designs was wholly superior. Because they then put all the teams together and they debrief. And then they said, "Okay, why don't we form a super team." And they formed a team of all members. And then they selected those to take the lead from each of the teams one more time to figure out where to put the equipment. Picked the best ideas from each of the other teams.

Today the factory is in Ashborough, North Carolina. Every single piece of equipment is exactly where the last of the teams put it. Cardboard City cost a grand total of \$2,000. It saved them millions and millions of dollars in square footage. Great. Run a competition around ideas, and then pull together the best of the ideas. Have multiple simulations that are cheap and moveable, so you can simulate the reality. It doesn't have to be expensive to run a creative valid experiment.

Systematic problem solving

I'm not going to say very much about systematic problem solving because it is something that's been around for years. It's really about data, not assumptions, about testing what you know, and about distinguishing what I'll call hard facts from gut facts. Gut facts are those that we assume to be true, because we've been in this business so long it's got to be right. Hard facts are those that you validated and tested by actually collecting the raw information.

Transferring knowledge

I want to spend a fair amount of time on this last one, transferring knowledge, because you can transfer knowledge in lots of ways other than just knowledge management systems. Five major ways. The most traditional are written, oral, and visual reports. Visual reports are just videotapes, and typically they summarize their findings. But there's another thing that we can do that's very powerful. And that's to create checklists, particularly for operating level people. But, it goes beyond operating level people.

I'm a pilot; I fly airplanes. And the very first thing they teach you in flight training is that you always rely on a checklist. Now, it's interesting that whether you have several hundred hours of flight time, as I do, or whether you have tens of thousands of hours of flight time, as commercial pilots do, you still rely on a checklist. Checklists have a very distinctive feature—they don't introduce new information. They're there to ensure discipline and repeated performance. Now, just imagine if they did introduce new information. You're about to get on a plane, you're going to visit some friends, and you walk by the cockpit. You hear this conversation between the pilot and co-pilot who are working on the checklist. "Hmm—item number five—I've never seen item number five before." "You ever seen item number five before? Let's just skip it and move on." A bit unsettling.

A checklist is there to remind you of things you already knew, but might not do because you overlooked them. This is actually something very important. Jack Welch, when he introduced a major change process at GE, developed a seven-step process which he insisted—it was his idea—he insisted that whenever they teach it, they call it The Pilot's Checklist.

And the reason was, he said, "It's not rocket science." The problem is people don't do these things consistently. With a checklist, they do it consistently. That's the way to move knowledge around from person-to-person, department-to-department. Obviously, you can take people on site visits and tours; you can rotate people. You can take people who are unusually skilled, perhaps at marketing, and move them into another division that needs marketing skill and marketing resources. Education and training—obviously, here's where knowledge management programs fit in, whether it's best-practice databases records, or whether it's Yellow Pages and directories that point you to experts. Whichever of these approaches you use, there are very real barriers, but often not the barriers that most people expect. The one most people expect is a resistance to things not invented here. That turns out to be a relatively minor barrier, surprisingly.

There are three big barriers. The first is that even when you have knowledge management systems and best-practices databases, it's very hard to tell which are your most important steps. It's very hard to tell what's really driving success. That has a very strong implication. When you transfer a practice, particularly for the first time, you should copy it exactly. You allow no changes in the first transfer.

It's not a coincidence that the company that does this better than anybody in the world, which is Intel, which manages to transfer enormously complex semi-conductor and wafer fabrication processes, calls the transfer process "copy exactly." In fact, the joke at Intel is that if the engineers of the two facilities are of different heights, the transfer process isn't going to work. That's the first barrier.

The second barrier has to do with differences in knowledge and skills between the best-practice site and the receiving site. Frequently, the reason the receiving site isn't doing this

already is that they don't understand it. They don't yet have the knowledge and skills. Again, very real implication. A very good predictor of success in these transfers is the amount of time that the experts from the best-practice site spend on site hand-holding, working together with those at the receiving site. You're going to have to do real-time problem solving, even if you copy exactly.

And third, often a big problem here is differences in status. Think about it. You have ten facilities. You head one of them, and somebody comes down and says, "You know, those folks in Chicago, they're the best at everything that we do. Why don't you go visit them and learn from them." You know your reaction—immediate defensiveness, "I'll be damned if I'm going to listen to what they have to say," etc. There are ways around that.

Rank Xerox, the British version of Xerox, came up with a wonderful way to do this. They do this on the revenue side. They found out among all the countries in Europe where they operated—there were nine of them—which ones were unusually good at selling different kinds of copiers. So, France was good at four-color copiers, Denmark was good at small copiers. They documented the best practices, one from each of nine countries.

Then they said to the country managers, "You must adopt four. It doesn't matter which ones; you must adopt four." Well, think about what's likely to happen. Every country is going to be a teacher, as well as a student, because everybody is good at something. All of a sudden the status difference disappeared because nobody is good at everything. People are good at different things. That program cost them a million dollars to document. That program generated a \$100 million in incremental revenue the first year, and an additional \$100 million of incremental revenue the second year. Again, a terrific return on investment.

Okay, now, in the interest of time, I'm going to wrap up with one more story, because we've gotten through the six major activities and some messages as well about how you lead the learning of others. But, I want to end, because some of you are probably sitting there and saying, "Boy, this is one big job. Man, he's told me I've got to collect market intelligence. I've got to benchmark. I've got to learn from my own past experience. I've got to experiment. I've got to transfer knowledge effectively. And I've got to do systematic problem solving. Got to be an easier way." I'll tell you one more story, because I believe there is no easier way.

This is a story about a fellow who was out hiking in the American Rockies, in the mountains. And he slips. He's on a trail high up in the mountains, and he slips, he falls off. And he manages to catch on to a tiny little twig. And he shouts, "Help me, help me. Is there anybody up there?" He's dangling off the edge. And all of a sudden the clouds roll in, and there's lightning and there's thunder, and a voice comes down, and it says, "I'm up here." He goes, "Help me, help me. Who's up there?" A voice comes down and it says, "It is I, the Lord."

The fellow goes, "Help me, Lord. Help me." A voice comes down and it says, "Do you believe?" The fellow says, "Lord, this is a very thin twig. I believe." The voice comes down and it says, "Then let go." There's a pause. The fellow looks up and he says, "Is there anybody else up there?"

I'm afraid if you're really serious about building a learning organization, and putting these things to work, there's nobody else up there. There's no easy way. You have to really move simultaneously on all these fronts. Thank you very much. Hope you have a great rest of the day.