

The Structure Of Organized Change:

A conversation with Kevin Kelly

by Joe Flower

This article appeared in *The Healthcare Forum Journal*, vol. 38, no. 1, January/February 1995.

International Copyright 1995 Joe Flower All Rights Reserved

Please see our free [downloading policy](#).

1. [Organization as organism](#)
 2. [Turbulence](#)
 3. [Basins of attraction](#)
 4. [The landscape of adaptation](#)
 5. [The mind of an organization](#)
 6. [Complexity out of simplicity](#)
 7. [The limits of adaptability](#)
 8. [Doing things right or doing the right thing](#)
 9. [Making healthcare transparent](#)
 10. [Top-down change](#)
 11. [Managing bottom-up change](#)
 12. [Exploitation vs. exploration](#)
 13. [Big is different](#)
 14. [Connection and disconnection](#)
 15. [Complexity](#)
 16. [Out of control](#)
-

Introduction:

A mile back from the edge of the continent, where the line of redwoods runs out, blackberry thickets skirt the edges of the bluff brown California hills. Kevin Kelly writes in a one-room

shack set a little ways off from his house. There is a beehive under the window. Cybernetic bees swarm across the cover of his new book, *Out of Control: The Rise of Neo-Biological Civilization* (Addison-Wesley 1994). They are a symbol of what he describes: organization as organism, the mind as a society, society as a mind, networks as distributed intelligence, the interlacing of the made and the born. It's a profound new lens with which to look at the world, one that arises out of chaos theory, systems thinking, experiments in artificial life, the rapid growth of machine intelligence, and the birth of cyberspace. Kelly is executive editor of *Wired*, the chief print chronicle of that birth (where I am a contributing writer). He was formerly editor and publisher of *Whole Earth Review*, as well as one of the founders of the cyberspace node called the WELL (Whole Earth `Lectronic Link), plus the Hackers Conference, and Cyberthon, the first virtual-reality jamboree. The book has created a stir, with George Gilder calling it "a sweeping and imaginative tour de force" and Stewart Brand praising it as "a breakthrough concept." We invited him here because his view of complex adaptive systems seemed particularly powerful and useful for healthcare, a system of extraordinary complexity in a period of great turbulence. Sitting with him in a quiet corner of his house under the redwoods, we started with one question: what is an organization?

Kelly:

Organization as organism

An organization is a set of relationships that are persistent over time. One of the functions of an organization, of any organism, is to anticipate the future, so that those relationships can persist over time.

An organization's reason for being, like that of any organism, is to help the parts that are in relationship to each other, to be able to deal with change in the environment. That means that they are trying to anticipate the future in some way or another. This is true of the separate pieces of a cell, of the cells in a body, of the bodies in a society, of the societies in an economy. Each system is trying to anticipate change in the environment.

The most interesting thing about change in the environment is that for the most part the environment isn't changing. The most certain thing you can say about the environment tomorrow is that it probably is going to be just like today, for the most part.

That may not sound as profound as it is, but it's actually a very important thing to have in mind when you are trying to structure something to adapt. The way that organizations and organisms anticipate the future is by taking signals from the past, most the time. Knowing what has just happened in the past, an organism can bring that data forward to modify its

future actions.

That is the definition of a feedback loop: a circuit that is taking information from the past and bringing it forward. We have common mechanical examples all around us, such as thermostats, or flush toilets. You have a signal that something has just happened, and you're bringing that signal from the past back into the present. The mechanism uses that signal to attempt to anticipate the future: a change in the temperature of the room, or the rising level of water in the toilet tank.

These are simple mechanisms that do not entail intelligence, or consciousness. They are very elementary, yet these feedback mechanisms are basic to complex systems.

Turbulence

There has been a sea change in our understanding of organisms. It has become evident that the primary lesson of the study of evolution is that all evolution is coevolution: every organism is evolving in tandem with the organisms around it. Each organism's environment, for the most part, consists of other organisms. In a broad systems sense, an organism's environment is indistinguishable from the organism itself. Obviously we can make a distinction that the lizard is not the same thing as the desert that it is in. But at the same time we can step back and say that the lizard and the desert form one system which acts in a certain way. They are one organism, the lizard and the desert.

When a system is in turbulence, the turbulence is not just out there in the environment, but is a part of the organization or organism that you are looking at. The organization and the environment are in concert. You can't easily divorce the organization and say, "Look at all that turbulence out there. How do we respond to it?" The organization is already responding to it, no matter what you do, and is indeed, in some measure, causing the disturbance.

Basins of attraction

We are learning something about how to understand turbulence, how to perceive the order in it. Generally we find that there is more order there than we thought.

For instance: there is a sense in which evolution is driven by random mutations. The current understanding was that it was impossible to predict how something would evolve because it was a very turbulent environment full of things interacting with each other. But in fact, when you try to model that on a computer you find that because of the very structure of matter and of the chemical bonds that are the basis of every organism, evolution is not random at all. It will tend to follow certain paths. Those paths are what in complex systems theory are called basins of attraction.

When you are run a model of a complex system, you cannot determine, from where you start, exactly where you will end up. They call that sensitivity to initial conditions. On the other hand, you can say that it is likely that you will end up in one of, say, five different basins. That's a great advance.

In biological evolution, these basins of attraction say to us, "No, you can't predict what kind of organism will evolve from an amoeba. But one of the basins of attraction is a form with bilateral symmetry and tubular guts, mouth at one end, anus at the other, homologous appendages arranged down the sides. If you could run the tape of evolution over and over again, you might never again get humans or giraffes, but you're always going to find organisms evolving tubular guts and bilateral appendages. It's built into the nature of the underlying forces and materials."

Basins of attraction, of self organization, show up as well in our complex social environment, in human organizations. Here again, while we cannot predict the result of any given input, we can say that it will likely fall within one of several areas. All imaginable futures are not equally possible.

This is actually a very important principle that science is learning about large systems like evolution and that futurists are learning about anticipating human society: just because a future scenario is plausible doesn't mean we can get there from here. The dynamics of that plausible future may be possible, but the pathway from here to there may be too arduous. It may be impossible, it may require certain things that we don't have at hand. It may require a political will or a belief system that we no longer have. It may require a sacrifice of some sort that we are not willing to make.

The landscape of adaptation

When evolutionary scientists speak about adaptation, one of the metaphors they use is that of a landscape, in which altitude represents how well-adapted the organism is to its environment. An organism may climb a hill of adaptation, becoming more and more adapted to its environment, only to discover that the hill is isolated, is not connected to the mountain of adaptation it can see in the distance. In order to get to that much more adapted state, it has to climb down the hill. It has to go through stages in which its adaptations are not useful, or are even mal-adaptive, in order to reach the more adaptive state. So it's stuck on its little hill. It has reached what is called a "local optimum," a false peak, and it can't get any better without losing what it has.

Organizations do this. You reach a point where you have to sort of de-evolve, you have to let go, you have to cross this valley, this desert. And sometimes there is not enough will, or resources, or sense of direction to be able to cross this desert to get to the other peak with is actually higher.

This happens to individuals within organizations, as well. In medicine, for instance, there are entire careers built on certain specialties that are in the process of being rendered obsolete by technological change, or de-emphasized because of the shift toward primary care, or because of other changes. The doctors in those specialties are not necessarily interested in, or skilled at, doing any other kind of medicine. In order to change, they have to back down on the evolutionary scale, perhaps go down an income level, perhaps go down a level of power in the organization from what they were before -- because their training and experience optimized them for that one particular thing.

The mind of an organization

We tend to think of the mind of an organization residing in the CEO and the organization's top managers, perhaps with the help of outside consultants that they call in. But that is not really how an organization thinks. We are infected by our own misunderstanding of how our own minds work.

Over the past few decades, people have worked very hard to build robots with artificial intelligence. One of the surprising discoveries that came out of that intense experience is that trying to make a central brain run things does not work. If you try to make a robot that walks, and you give it a brain that has some sort of eyes to see with, and give that brain the job of notifying the legs when to move, it will never fail to flop over. Using a centralized brain for the task of trying to anticipate the future and deal with change with just doesn't work.

Some researchers found they made more headway when they started from the bottom up, instead of working from the top down. They decided to build intelligent robot that was only as smart as an ant. They had observed that ants walk really well. The little tiny ant's brain did that job a lot better than any robot. So the researchers wondered how they were doing it. And they discovered something very interesting: when it comes to walking, most of the ant's thinking and decision-making is not in its brain at all. It's distributed. It's in its legs.

Complexity out of simplicity

The way to build a complex system that works is to build it from very simple systems that work. You start with a very simple module and you make it bug free, you make it perfect. Then you build on top of it without changing that initial module. The way you build a system that walks is not with a global intelligence that has is all figured out, but with a distributed intelligence making very simple decisions based on very limited criteria, such as: "I'll move this leg that I am in charge of if the leg behind me has moved, and if there is nothing in front of me." Each leg moves in response to cues from the other legs, or from the eye which says there is nothing in front of it. Each leg sort of walks on its own. It doesn't know where it's going. It has no idea what else is happening in the body beyond the cues it needs to make its

decisions. There are sensors for reflexes that can discourage or to turn off a leg, inhibit its action. But this little insect leg will continue to move forward as long as it's allowed to.

What you have is a very complex circuit of small, simple reflexes that are inhibiting, suppressing, or launching other reflexes -- and out of a collection of simple reflexes you get complex behavior.

This is actually how our brains work. A brain is a society of very small, simple modules that cannot be said to be thinking, that are not smart in themselves. But when you have a network of them together, out of that arises a kind of smartness.

Organizations work very much like that. The idea that there is somebody actually directing them, that there can be a "I" at the center, an autonomous person or directorate making everything happen, is an illusion. The "I" of an organization is an emergent phenomenon, greater than the sum of its parts, which arises out of the whole thing.

My favorite example is the bee hive. People think that the queen bee is directing the hive. But she's just following it. Where does the "I" of the bee hive reside? It resides in the hive as a whole. The hive makes decisions without any of the individual bees even being aware of making a decision, or even that a decision is being made. And yet the distinctive personality of the bee hive emerges.

Organizations are a lot like that. Certainly there are people that are in charge, people whose job it is to try to look for and embody the "I" of the organization, people who make decisions and sign the documents. But if you try to isolate who is really thinking for the organization, it's a futile quest. The organization as a whole does the thinking.

An organization's intelligence is distributed to the point of being ubiquitous. It's distributed to every component of the organization, and there is no place that you can put your finger on it and say, "Here's the mind," just as intelligence is ubiquitous and holographic in the brain.

The limits of adaptability

It's generally much easier to kill an organization than to change it substantially. Organisms by their design are not made to adapt too far. They have only a limited ability to adapt beyond a certain point. And beyond that point it's much easier to kill them off and start a new one than it is to change them.

Managers tend to treat organizations as if they are infinitely plastic. They hire and fire, merge, downsize, terminate programs, add capacities. But there are limits to the shifts that organizations can absorb.

In some cases there are quantum levels of change: an organization may be able to be what it is now, or be something else that its managers can envision, but not necessarily be the things in between those two states. Or they may be able to go quickly through those transitional states, but they can't be stable in them.

Species go extinct because there are historical constraints built into a given body or a given design. If organizations were not embodied in physical buildings and real people in real locations, they could become anything they want at any time. But when you are embodied in a location, in a physical plant, in a set of people, and in a common history, that constrains your evolution and your ability to evolve in certain directions.

If you are the executive or the board of an organization and you can see, looking at your environment, that you really need to have some capacity that is significantly different from the capacities that you have now, you are more likely to succeed if you start a wholly new enterprise than if you try to shift what you are doing now.

It's psychologically easier to think of it as siring offspring that retain some of the genes, some of the stuff that you have learned, and growing something new. But it's becoming ever more important to be able to let go and to kill things.

Doing things right or doing the right thing

Peter Drucker synthesized a powerful thought for our era. He pointed out that in the past, in the industrial era, we focused great effort on allowing people to do their jobs better. We saw learning how to do the job better as the principle driver of quality, and ultimately of wealth. In this new era of network economics, where information plays a greater and greater competitive role, the question becomes not, "How can I do this job better?" but "What should I do? What is the right thing to do?"

When you start asking those questions you find that just doing what you have been doing, but doing it better, is not necessarily the answer. You often have to stop doing what you're doing, and go do something new. This process of letting go, of stopping what we are doing entirely and doing something else in its place, has not really made it into our vocabulary of corporate management.

This kind of change is disruptive. It's heart-rending. It's not easy. We see this often not with entire organizations but with products. Organizations get invested into a particular product. And sometimes the best thing is to stop making that product, even though it's profitable, because it has optimized at a local peak. This is where it becomes hard. It's like saying, "This is a nice peak, but we are going to get stuck here. We've got to let go. We've got to kill it and climb down this hill, so that we can cross the valley and move on to larger peaks."

We don't have the organizational and managerial tools for giving things up very easily.

For instance, as health care moves from just fixing people who are sick toward more prevention and community work, it is finding out something important: the kind of organizational and individual skills that you need to fix people are very different from the kind of skills that you need to get people to change the behaviors that make them sick in the first place. This is a very different business.

Making healthcare transparent

Here again we cannot separate the organism from its environment. We lay people have to shift our understanding of the process: that doctors and hospitals are not just for fixing me, but are partners with me in anticipating change. For that to happen, healthcare has to become more transparent.

Healthcare has been something that happens in hospitals, hidden, removed from view. The growing ability of information databases to have all my medical records together, in a format that I can understand, with backup materials to deepen that understanding, will help me deepen my sense of my own medical history.

Technological advances could allow us to see more clearly into our own lives. They will lend greater transparency, and a longer view, both to the people and to the organizations serving the people, so that we together can see what our environment is.

For instance, two weeks ago our three-year-old daughter fell out this window here, onto the concrete one floor down. She's okay, but it was pretty traumatic and scary. The hospital did a great job, but I still intend to write them to get copies of the CAT scan and X-rays that they did, just for our own education, for my daughter to see herself in a different way. Why shouldn't I be able to put that onto our computer monitor, learn from it and show her what happens? If we can bring people that information, and make it evident to them, medicine will not be such a black art.

That's what personal computers did. The great advance of personal computers was not the computing power per se but the fact that it brought it right to your face, that you had control over it, that were confronted with it and could steer it. As new technologies enrich our environment, the health care system need not seem as remote from our own lives as it does right now.

Top-down change

Complexity that works is built up out of modules that work perfectly, layered one over the

other. Through the '80s and into the '90s health care organizations, much like many other American organizations, have been going through a sort of new management "fad of the week." But trying to change a complex system from above carries with it great difficulties.

Think for a moment about the difference between a free market economy and a Soviet-style command economy. A free market economy allows you to have mini-revolutions on a daily basis, rather than have a big one once every fifty years.

So how do you get a large organization to respond to the environment? You can do it in two ways. One is you can wait and then change the entire organization all at once and hope that it works -- and if it doesn't work you're dead. Or you can allow all the little parts of the organization to adapt on a regular basis, so that you don't have your change happen all at once.

Changing things from the top down works when things are stable. It's a very efficient way to do it. But in a turbulent environment the change is so widespread that it just routes around any kind of central authority. So it is best to manage the bottom-up change rather than try to institute it from the top down.

Managing bottom-up change

Managing bottom-up change is its own art. You can start by honoring errors. You have to have a certain number of mistakes and small failures or you're not trying hard enough.

There is an art to this. There are some processes that don't allow for experimentation, creativity, or error at all. We don't want the people who run nuclear power plants to get creative and make up their own rules. Some things are "mission critical." Those are usually simple, proven processes on which you can build complexity. That's probably not where you are going to find a creative advantage.

Almost all innovation in all a system happens at fringes. So maximize fringes. The nature of an innovation is that it will arise at a fringe where it can afford to become prevalent enough to establish its usefulness without being overwhelmed by the inertia of the orthodox system.

Maximize fringes. Have more skunk works. Emphasize subsidiaries that are some distance from the main office. Give people some slight autonomy. Have a few hidden budgets. The fringes are where innovation and change originate. If you try and do it in the central area, if you have a "Department of Change" it will get overwhelmed by the orthodoxy that is surrounding it. You have got to give it some room out there where the innovation can blossom and take root for a little while, so that it can prove itself and move back into the center.

Exploitation vs. exploration

In a complex adaptive system there is always a trade-off, which is played out at every level -- the trade-off of exploitation versus exploration. The system continually has to make this choice: it can either continue to exploit a known process and make it more productive, or it can explore a new process at the cost of being less efficient.

In healthcare, for example, suppose you have a new drug. If you are the patient, you don't want to get an experimental drug. You want the healthcare system to optimize a known drug and make it better. The exploration for new drugs involves a certain level of experimental loss -- there is no guarantee that the new drug is as effective and problem-free as the old one. But unless you explore, you'll never have new drugs, you'll only have better old ones.

With the increasing use of outcomes research and outcomes management, healthcare is shifting the emphasis slightly from being a very exploratory environment, to being more exploitive. Much of outcomes research is a systematic attempt to exploit what is known and make it better.

The necessary exploratory component that a system must have to stay adapted to its environment is very small, something in the neighborhood of one percent. In an organization that might mean one percent of your people, or one percent of your revenue -- one percent of your actions are sent exploring in a very open way.

That's a small amount, but it is necessary. If outcomes research were used to so regularize medical practice that there was no variation whatsoever, you could lose that wild, random data. You would have no comparable data with which to continue to improve the system.

Big is different

Healthcare is going through a period of amalgamating into larger and larger systems. Managers of those systems will experience an interesting quality of systems: big is different. More is different. There are qualitative differences in systems as they become larger. The dynamics go from linear to non-linear. They become much more unpredictable. In a very simple system you may be able to describe all the states it could be in; but a complex system is in a sense unknowable. You cannot even describe all the states it could be in.

That threshold into complexity is very different for different systems. But if you are inside the system you can often tell when it crosses the threshold. Things are suddenly not acting at all like they were before.

There will be a different kind of bigness to deal with, a complexity that is dispersed geographically, temporally, and organizationally.

This calls for an organizational model of loose affiliation rather than tight control, with the

hierarchy determined not so much by rank as by time and size: the higher levels are those that are concerned with longer periods of time over greater parts of the organization. This is much more realistic than thinking of a hierarchy as the person at the top giving orders that the people at the bottom must follow.

Connection and disconnection

We are in a phase now, generally, of opening up wider and deeper and more connections to one another. And we still value those connections. But when everything is connected to everything else, connection becomes cheap, and what has value is disconnection. We are approaching -- some people are already there -- the stage where information and connection become overwhelming, and the technologies of disconnection become more valuable. We will develop more and more sophisticated and powerful screens and filters for our information and connections. One of the lessons from computer modeling of very complex systems is that if you connect everything to everything you get gridlock and paralysis. In our human institutions we are at the cusp of wiring everything to everything with email and conferencing systems, and we don't yet have a very good idea of how to manage those connections and discover what we really need to connect and what just leads to paralysis.

Complexity

What we find in nature are self sustainable, self repairing, self replicating systems. Isn't that what we would want for the things that we make? What we would love to create is an organization that went on for a hundred years, got better, continued to grow, repaired itself, exploited itself, and overall governed itself. In order to do that we humans, we creators, have to let go of the thing. We have to surrender some of our control and let the system run itself.

When we surrender some of our control we get all these cool things like artificial evolution, adaptive materials, autonomous agents, smart buildings, and so forth. It's not that we entirely accept whatever behavior they have, and say that we'll just tolerate whatever happens. It's more along the lines of raising a child: we train the system to a certain range of behaviors that we find most useful. But then we let it go, because we don't want to have to be babysitting it the whole time.

Out of control

So the title of my book, out of control, to me means something like co-control or para control. We don't drive systems, we shepherd them. The sheep are doing their own thing, eating the grass, finding their own water, producing the wool. We have some guard dogs that are keeping them in line. The shepherd keeps the flock in the right general area, and harvests the

results. This is the kind of systems, and the kind of management of systems, towards which we are headed -- pieces of software, for instance, with billions of lines of code, that adapt themselves to their environments, automatically repairing the broken pieces of code, deciding when it needs updates, communicating with the other pieces of software in the world and finding out how they are doing, learning about the world. And in that sense it will be literally out of our control.

Many of our systems are approaching that level of complexity. That means that there are going to be surprises. That means that we may not understand exactly what is happening.

Everything that we are making, we are making more and more complex. As the complexities of the things that we make increase, they become somewhat biological in nature. In order for us to manage them, we need to import some of the principals that nature uses to manage those natural ecosystems -- trees, bushes, birds, meadows, and jungles.