

homerdixon.com

Thomas Homer-Dixon – Complexity Science and Public Policy

Complexity Science and Public Policy

This article is based on the Canada School of Public Service's 2010 John L. Manion Lecture, delivered on May 5, 2010 at the National Arts Centre in Ottawa, and entitled Complexity, Crisis and Change: Implications for the Federal Public Service. The opinions expressed are those of the author and do not necessarily reflect the views of the Canada School of Public Service.

I am absolutely delighted to be back in Ottawa. In the late 1970s and early 1980s I lived here for five years, and it's always a joy to return to this city, especially in the spring when the tulips are out.

It's also an enormous honour to be asked to present the Manion lecture. This lecture provides an opportunity to engage with new ideas – and to test new ideas – before a highly experienced audience. This evening I'm going to talk about the science of complexity and how it might be applied to public policy.

I have been working in complexity science for about 15 years. Until recently, the community of complexity researchers was quite isolated. We consisted of a few clusters of researchers here and there around the world, but for the most part we hadn't integrated ourselves into a larger worldwide community. Also, our work hadn't received much attention, partly because we often couldn't show how complexity science might be used to make the world a better place. Theories of complexity can be very abstract. Their ideas and concepts don't cohere well; in fact, the body of thinking we call complexity science is largely fragmented into bits and pieces.

But now, exciting projects like the New Synthesis Roundtable led by Mme Jocelyne Bourgon here in Ottawa are taking ideas from complexity science and applying them to the real world – to see how they might help us develop new approaches to public policy and better address the extraordinarily hard problems our societies face today.

As complexity theory has started to receive more attention, it has also started to accumulate critics. There are probably some of you out there in the audience this

evening who are grumbling, “This is all just another fad.” I can understand why you might think so. Some people, especially rebellious graduate students I find, are similarly starting to talk about the “cult of complexity.”

My job this evening is to argue that complexity science isn’t a fad. I will offer a brief survey of some core concepts and ideas, and I will make a strong case that the tools and ideas of complex systems theory can give us significant purchase on the new and strange world we’re living in today. Most importantly, they can help us develop new strategies for generating solutions and prospering in this world.

So let’s begin.

We live in a world of complex systems

We need to start thinking about the world in a new way, because in some fundamental and essential respects our world has changed its character. We need to shift from seeing the world as composed largely of simple machines to seeing it as composed mainly of complex systems. Seeing the world as composed mainly of simple machines might have been appropriate several decades ago: we commonly thought of our economy, the natural resource systems we were exploiting, and our societies in general as machines that were analogous, essentially, to a windup clock. Each could be analyzed into parts, with the relations between those parts precisely understood, and each was believed to be nothing more than the sum total of its parts. As a result, we believed we could predict and often precisely manage the behaviour of these systems.

But now, increasingly, we live in a world of complex systems, and we have to cope with the vicissitudes of these systems all the time. Earth’s climate is clearly complex. Ecological systems are complex, and we’ve often managed them miserably when we’ve assumed they worked like simple machines – take a look, for example, at what we did to the east-coast fishery. Our economy, especially the global economy, is a complex system. Our energy systems, such as our electrical grids, are increasingly behaving like complex systems. Food systems, information infrastructures, and our societies as a whole all exhibit characteristics of complex systems.

No longer is it appropriate for us to think about the world as equivalent to, or an analog of, a mechanical clock, which one can dismantle and understand completely and which is, ultimately, no more than the sum of its parts. Instead we have to think in a new way.

To do so, we must first ask: What is complexity and what features distinguish

complex systems from other kinds of systems? Surprisingly, even some of the world's leading complexity thinkers have trouble answering this question. They tend to provide a checklist of properties common to complex systems, as I will in a moment. But to a certain extent, understanding complexity requires that one work with complex systems for an extended time and then study in depth the literature on complexity. In this way, one eventually develops an intuition for what complexity is.

Since we don't have such time this evening, I'll instead identify some properties that I regard, for the most part, as necessary features of complex systems.

Most complex systems have many components. They also have a high degree of connectivity between their components – an issue I'll discuss extensively later in my presentation, because we need to unpack it completely. Additionally, complex systems are thermodynamically open. By this I mean that they're very difficult to bound: we can't draw a line around them and say certain things are inside the system while everything else is outside. As a result, in terms of their causal relationships with the surrounding world, complex systems tend to bleed out – or ramify or concatenate out – into the larger systems around them. And ultimately the boundary that we draw demarcating what is inside and what is outside is largely arbitrary.

Flowing across this boundary are information, matter, and most importantly energy. The flow of high-quality energy into complex systems allows them to sustain their complexity. In thermodynamic terms, these systems maintain themselves far from equilibrium. If we take away this energy, they start to degrade. The complexity disappears, they become simple, and they fall apart. Keep your mind focused on that point, because I am going to return to it in a few minutes.

The behaviour of complex systems is also non-linear. By this complexity specialists mean something very specific: in a nonlinear system small changes can have big effects, while sometimes big changes in the system don't have much effect at all. These systems exhibit, therefore, a fundamental disproportionality between cause and effect. In contrast, in a simple machine small changes generally have small effects, while big changes have big effects. This difference in the nature of causality is one of the fundamental ways of discriminating between a simple machine and a complex system.

Finally, we have the characteristic of emergence. Emergence is probably the property that comes closest to being sufficient for complexity: if you see it, you're very likely dealing with complexity. We have emergence when a system as a whole exhibits novel properties that we can't understand – and maybe can't even predict

– simply by reference to the properties of the system’s individual components. It’s as if, when we finish putting all the pieces of a mechanical clock together, it sprouts a couple of legs, looks at us, says “Hi, I’m out of here,” and walks out of the room. We’d say “Wow, where did that come from?”

Once we have a list of common characteristics of complex systems, we can then think productively about how we might measure complexity. One method involves developing a computer program or algorithm that accurately describes or predicts the system’s behaviour under different circumstances. The longer the computer program or algorithm, the more complex the system it describes. Specialists call this metric “algorithmic complexity.” Experts have proposed a variety of other metrics of complexity. For our present purposes, the most important point is that by a lot of metrics our world is unquestionably becoming more complex. It’s becoming more connected, we see larger flows of energy into our socio-ecological systems, it’s exhibiting greater non-linearity, and it’s exhibiting lots of emergent surprises – more and more, it seems, all the time.

In this regard, I want to highlight a particularly interesting characteristic of our modern societies that has led to a surge in the number of system components: the rapid dispersion of power. Enormous increases in technological power have fundamentally changed the distribution of political power within our societies. A standard laptop computer today has about as much computational power as was available to the entire American defence department in the 1960s, and in those days a computer of such power would have filled a building about the size of the one we’re sitting in this evening. Today, this power is compressed into a little four-litre box, and these boxes are available to hundreds of millions of people around the planet. These people have at their fingertips, as a result, staggering computational, analytical, information-gathering, and communication capability. For all intents and purposes, this capability has translated into political power – a disaggregation or flattening of the social and political hierarchy – because of the diffusion throughout our societies of the capacity of groups and individuals to express forcefully their political and economic interest. In a sense, this diffusion of power has led to a proliferation of agents, which is equivalent to a rapid increase in the number of components within our societies, and in consequence a rapid increase in societal complexity.

This last point leads us naturally to the question: What, in general, causes complexity to increase?

Sources of complexity

The most straightforward answer is that human beings introduce complexity into

their social, economic, and technological systems to solve their problems. The scholar Joseph Tainter has made this point very effectively. He suggests that over time societies encounter problems, and they tend to respond to these problems by creating more complex technologies and institutions.

This is a useful response to the question, but inevitably we can go deeper. In 1994 the economist W. Brian Arthur, one of the world's most insightful complexity theorists, wrote an article that I regard as one of the foundation pieces of complexity science. He suggested there are really three deep sources of complexity. The first is growth in co-evolutionary diversity. This process applies equally to societies, economies, and ecological and technological systems. Ecological systems offer, perhaps, the clearest illustration. Arthur says each ecological system has a number of niches or ecological roles that may or may not be filled by various species. Niches filled by one or more species are separated by vacant niches. These vacant niches offer resources of various kinds – material, food, energy – and as a result new species evolve to fill those niches. When a new species fills a niche, it automatically creates more niches, which provide further opportunities for the evolution of yet more species. In this way over time, complexity begets complexity.

Arthur shows that this important and interesting process operates within human societies. For instance, it applies to the evolution of technologies like computer systems: we start with relatively simple computers and associated components; then new technologies, such as software packages, and hardware, such as printers and backup systems, are developed to fill the gaps between those entities, thus creating further gaps that yet newer technologies can fill.

A second process Arthur identifies is structural deepening. It's a very different phenomenon: if growth in co-evolutionary diversity happens at the level of the whole system, structural deepening happens at the level of the individual component or unit within the system. As these components (such as species in an ecological system or firms in an economy) compete with each other, they tend to become more complex in order to break through performance barriers. This idea is similar to Joseph Tainter's: as a species, firm, or organization confronts problems in its environment, it responds by becoming more complex.

We can see structural deepening at work in many of our technologies. Compare for instance an automobile engine back in the 1960s with one produced today. The modern engine runs much more cleanly, it's far more efficient, and it has other attributes that make it a great improvement over the earlier version. But back in the 1960s, you might have been able to fix the engine yourself. I would challenge you to do so now. As the world gets more complex – as it structurally deepens – we have become more reliant on specialists to take care of us and to provide essential

services.

Finally, Arthur talks about the phenomenon of capturing software, in which larger systems appropriate or capture the grammar that governs the operation of smaller or subordinate systems. Arthur points to the way societies have captured the software – or the fundamental physical grammar – of electricity and have then used electricity in all kinds of marvellous ways to improve people's lives. But in the process, we have made our world much more complex.

Complexity depends on high-quality energy

I want to turn now to the relationship between energy and complexity, a topic I've already mentioned. I've noted that over the course of their history human beings have dealt with their problems by developing more complex institutions and technologies. Joseph Tainter has additionally emphasized that as we develop more complex institutions and technologies, our requirement for high-quality energy to build and sustain these institutions and technologies generally rises. Today's modern cities for instance, exhibit extraordinary complexity from the point of view of, say, somebody in the 19th century, and the energy inputs needed to sustain these urban systems are in many respects beyond belief. (I use the term "high-quality" to refer to the thermodynamic quality of the energy in question. Some forms of energy like natural gas and electricity are very useful for doing work – essentially they can be used for a lot of different purposes – while others, like the ambient heat in our natural surroundings, are not much good for anything. A modern society can't sustain its complexity with low-quality energy; it needs copious quantities of high-quality energy.)

Now the problem, of course, is that humankind is going through a fundamental energy transition. We're facing supply constraints for one of humankind's best energy sources – oil. I want to emphasize the significance of this change in our circumstances. Conventional oil provides 40 percent of the world's commercial energy and around 95 percent of the world's transportation energy. It's literally the stuff the planet's economy runs on, and in thermodynamic terms it's very special. Three tablespoons of oil contain as much free energy as would be expended by an adult male labourer in a day. Every time we fill up a standard North American car, we put the equivalent of two years of manual labour into the gas tank. For the last century or so, cheap oil has translated into essentially dozens of nearly free slaves working for each one of us.

The oil age began in 1858 in Oil Springs, Ontario with the first commercial discovery of oil, and it will end around the middle of this century – lasting about two centuries all told. This statement doesn't mean that we are going to run out of

oil by 2050; rather, it means we're going to switch to something else, because oil will become much more expensive than it is now.

By “expensive” I mean energetically expensive. Even now, drillers are going further into more hostile natural environments to drill deeper for generally smaller pools of lower-quality oil. They're working harder for every extra barrel. The trend is long term, inexorable, and striking. In the 1930s in Texas, drillers got back about 100 barrels of oil for every barrel of energy they invested to drill down into the ground and to pump oil out. Today, the “energy return on investment” (as specialists call it) for conventional oil in North America is 17:1. For the tar sands in Alberta it's around 4:1; so producers get back about four times the energy they invest. For corn-based ethanol the figure is about 1:1, which means producers put in about as much energy as they get back. Corn-based ethanol is a great subsidy for farmers but a terrible energy technology.

Taking the average energy return on investment of all energy sources in our economy, as we slide down that slope from 100:1 to 17:1 to 4:1 to 1:1, we're inexorably using a larger and larger fraction of the wealth and capital in our economy simply to produce energy, and we have less left over for everything else we need to do – like solving our increasingly difficult problems. Steadily more expensive energy will have all kinds of effects on our societies, but most fundamentally it will make it progressively harder for us to sustain our societies' complexity.

I'll come back to this issue later in my presentation, but first I'd like to address the question: Is rising social, economic and technological complexity a good thing or a bad thing? I would say that it depends on the state of evolution of the complex system in question. Before I elaborate further, I'll answer the question in brief.

The good and bad sides of complexity

Greater complexity is often a good thing. We wouldn't be living as we do now if it weren't for complexity. We are vastly healthier, we live longer, and we have enormously more opportunities, options, and potential in our lives as a result of the complexity we have introduced into our technologies and institutions. As Tainter argues, complexity helps us solve our problems.

Often, too, complexity is a source of innovation because it allows things that would not otherwise be combined to be brought together in unexpected ways. The complexity theorist Stuart Kauffman calls these combinations “autocatalytic sets.” Complex societies are like a big stew: we throw in all kinds of different things, mix

them together for a while, and then see what happens. Richard Florida's theory of innovation in urban areas picks up on this idea: large, diverse, and tolerant cities are engines of innovation, because they allow for countless novel and unexpected combinations of people, ideas, cultures, practices and resources.

Finally, complexity provides us with greater capacity to adapt to change, at least under certain circumstances. To the extent that complexity boosts diversity in a societal system, we have available a wider repertoire of routines, practices, and ideas for adaptation and survival when our external environment changes and new challenges arise. Some people, firms, organizations, groups, or cultures will do well and some of them won't, but diversity raises the likelihood that at least some components of the social system will prosper in the face of change.

Similarly, if a system has distributed capability and redundancy – as many complex systems do – then if one component of the system is knocked out because of an accident in a technological system, a pathogen in an ecological system, or a fire in a forest, other components can step in to prevent cascading damage to the larger system.

In all the above respects, complexity is a good thing. But inevitably there is another side to the story, and increasingly I think we're seeing the bad side of complexity. First of all, complexity often causes opacity; in other words, complexity prevents us from effectively seeing what's going on inside a system. So many things are happening between the system's densely connected components that it becomes opaque. Complexity also contributes to deep uncertainty. While opacity is a variable that operates in a slice of time – say, the present – uncertainty arises when you try to project the behaviour of a system forward into the future. The further we try to predict into the future, the fewer clues we have about what the system is going to do and how it's going to behave.

As our world has become more complex, we have, in fact, moved from a world of risk to a world of uncertainty. In a world of risk, we have data at hand that allow us to estimate the probabilities that any given system we are working with will evolve along certain pathways, and we can also estimate the likely costs and benefits associated with evolving along one of those pathways or another. In a world of uncertainty, we simply don't have a clue what is going to happen. We don't have the data to estimate the relative probabilities that the system will evolve along one pathway or another; in fact we don't even know what the possible pathways are. And we certainly can't estimate the costs and benefits that will accrue to us along different pathways.

This is a world filled with “unknown unknowns.” I find it interesting that members

of the military who have seen combat are deeply familiar with this concept. Indeed, it's in such common use in the US military that people abbreviate it to "unk unks." From their hard personal experience, soldiers know that surprises happen on the battlefield. Surprises come out of the blue. In his renowned treatise *On War*, Carl von Clausewitz, the 19th century Prussian military theorist, wrote about "friction" on the battlefield and the "fog of war." Military people throughout history have known that they can't plan and predict everything. They have known that in a world of uncertainty and unknown unknowns, we are ignorant of our own ignorance; often, we don't even know what questions to ask.

Not only are complex systems opaque and uncertain, they also exhibit threshold behaviour. By threshold behaviour I mean a sharp, sudden move or "flip" to a new state. This new state may or may not be a new equilibrium – that is, it may or may not be stable. In the wake of the collapse of Lehman Brothers in September 2008, the world economy certainly flipped somewhere – it clearly exhibited threshold behaviour – but it's not at all clear that it flipped to any kind of equilibrium, because the crisis continues to unfold today. As we have seen with the world economy, electrical grids, and large fisheries, complex systems exhibit a capacity for sudden, dramatic change.

Not all instances of threshold change are bad. The fall of the Berlin Wall and the subsequent collapse of the Soviet Union were, I would argue, indisputably good things. But to the extent that the sudden change is a surprise, so we're not ready for it, and to the extent that our existing regime of beliefs, values, rules, institutions, and patterns of behaviour are tightly coupled to the former situation, and we don't have any clear plans to adapt to the new situation, then threshold change is basically a bad thing.

Complexity can also cause managerial overload. This is basically an issue of information flow. I imagine I'm ringing bells in your heads when I say that today our cognitive capacity is too often exceeded by too many things happening at too high a rate. With email, BlackBerries, iPhones, and the like, we're all at the convergence point of multiple streams of information, and we're all juggling five, ten or more tasks or crises simultaneously.

In the last 30 years or so, with the development of fibre optic cables and advanced information switching systems, humankind has increased its ability to move information by hundreds of millions of times. But our ability to process that information in our brains has stayed the same. So waves of information pile up at the doorstep of our cerebral cortex. The proliferation of urgent demands produces decidedly sub-optimal responses like multitasking and superficial information processing, and it sharply increases stress. And if this stress exceeds the coping

capacity of a person, organization, or society, it can ultimately lead to systemic breakdown.

Additionally, complexity is a bad thing when it boosts the vulnerability of systems to unexpected interactions and cascading failures. These outcomes result from a combination of dense connectivity and tight coupling between system components. Dense connectivity and tight coupling are often conflated, but they are really distinct phenomena. The former means the system has lots of links between its components. In our modern societies, new information technologies have boosted enormously the number of links between people, organizations and technologies. Tight coupling, on the other hand, means that two events in a given system are separated by a very small physical space or a very short interval of time.

When we link together tightly lots of previously unconnected things, we sharply raise the probability of unexpected interactions. A couple of decades ago, in his marvellous book *Normal Accidents*, the Yale sociologist Charles Perrow detailed the dangers of unexpected interactions within increasingly densely and tightly coupled systems. Today, Perrow's warnings seem prescient, especially since humankind is now connecting together entire systems that were previously largely independent. For example, the spike in energy prices in the summer of 2008 showed that the world energy system is not only tightly linked to the world economy (many economists believe that the 2008 energy shock was the precipitating cause of the US recession and, ultimately, the current world economic crisis), but also now to the world food system. Higher oil prices stimulated a rush to biofuel production, which caused huge tracts of land to be switched from food to biofuels; this change in turn caused a surge in basic food prices around the world. Such consequences are exceedingly hard to predict in advance. Once again, we're in a world of unknown unknowns.

Dense connectivity and tight coupling also raise the probability of cascading failures. Think of a row of dominoes falling over: the dominoes are close enough together that tipping the first one tips all the rest in succession. Cascading failures occur more often now in our modern systems because the sharply higher speed and volume of movement of energy, material and information between components of our economies, societies and technologies has dramatically tightened the physical and temporal proximity of events in these systems.



I use the analogy of a system of cars tailgating each other at high speed on a freeway. The cars are traveling fast and close together, so they cover the distance between themselves in an instant. Then, if one driver is not really paying attention, perhaps because he or she is entering a text message into a BlackBerry while switching lanes (at this point in the presentation I always see a lot of people turn their heads down, because they know who they are), a sideswipe happens and in a flash dozens of cars are piled in a heap.

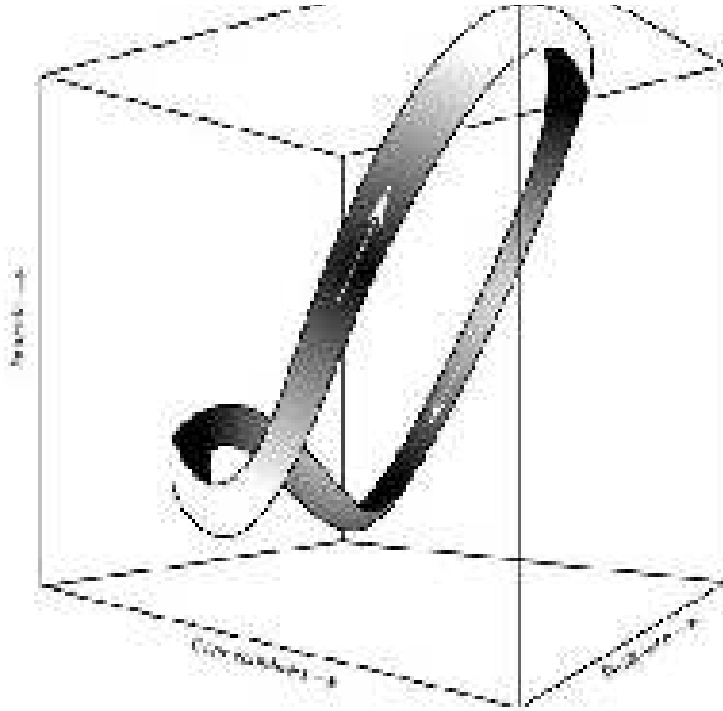
This image looks a bit like the American economy about a year ago and maybe the global economy in a few weeks or months. I would argue that the resemblance is more than superficial.

Last but not least, complexity is sometimes a bad thing because it increases brittleness. To explain why, I need to outline ideas developed by one of the world's most brilliant ecologists, a Canadian, C.S. or "Buzz" Holling. Holling's ideas on system brittleness – more specifically on system resilience – fall under the general rubric of Panarchy Theory. I will spend some time this evening unpacking this theory, because it's staggeringly powerful. Conceptually, this will be the most difficult part of my talk.

Evolution of complex adaptive systems (Panarchy Theory)

Panarchy theory represents the evolution of complex adaptive systems (that is, systems that adjust or adapt to their external environment as that environment changes) in three-dimensional space. This space is defined by the variables, potential, connectivity, and resilience, as you can see in the accompanying figure.





By potential, Holling and his colleagues (now spread around the world in a loose-knit organization called the Resilience Alliance) mean the possibility for novelty within a system. A rough analog would be the system's information content. By connectivity, they mean something very much like the concept of connectivity I've used throughout this presentation. Finally, for Holling and his colleagues, resilience is the capability to withstand shock without catastrophic failure. As one of the New Synthesis documents says, resilient systems are able "to adapt and adjust to unforeseen events, to absorb change, and to learn from adversity."

The accompanying figure is my own interpretation of this "adaptive cycle." Specifically, to make it easier for a lay audience to grasp, I've reversed the orientation of the resilience variable, which has caused some change in the shape of the three-dimensional loop, but nothing that interferes with the model's underlying message.

To illustrate the adaptive cycle, let's take a simple example from Holling. He began his work studying forests, in particular the spruce forests in New Brunswick, because he was interested in understanding outbreaks of spruce budworm. In terms of its potential, connectivity, and resilience, a young forest starts at the rear of the cube, in the far bottom corner. It has relatively low the entire forest genome. The species and organisms are dispersed loosely across the landscape and thus are relatively loosely connected. But precisely because of this loose connectivity, the forest exhibits high resilience.

As the forest grows and moves towards a climax state at the cube's front, centre, and top, it climbs what panarchy theorists call the adaptive cycle's "front loop." It becomes more and more connected, because more species move in, and they develop more relationships among them in terms of flows of material, energy, and fundamental elements (such as carbon, sulphur, nitrogen). As the forest climbs the front loop, potential for novelty also rises: mutations in the forest's genetic material proliferate; these mutations may not be expressed, but they are available as possibilities of future novelty (which is, I believe, one of Holling's most penetrating insights). Interestingly, the whole system eventually becomes less resilient too, for reasons I'll explain shortly.

This model has enormous power to explain the evolution of other kinds of complex adaptive systems, including economies, firms, organizations, institutions, technological systems, and even whole societies. I would suggest, in fact, that it captures many characteristics of today's global socio-ecological system. We now have to think of humankind's global society and economy as intimately linked with an ecological system that provides the food, energy, and resources it needs to sustain itself. This global system in its entirety is now very tightly coupled, with both enormous information content and potential for novelty, but nonetheless declining resilience.

What happens at the top of the front loop is a very important part of the adaptive cycle's story. Eventually, because of the combination of loss of resilience and some proximate trigger – in the case of a forest, perhaps a drought or the outbreak of fire or disease – the system breaks down. At this moment the time-frame shifts: while things have progressed slowly as the forest climbed the front loop – that is, change has been relatively incremental – the breakdown process that begins at the top of the loop (called the omega phase) happens quickly. The system disaggregates or decouples, and connectivity is lost, which allows for the reorganization of the system's remaining components into new forms. This change in turn allows for the adaptation of the system to a new environment or circumstance.

So breakdown is a vital part of adaptation, an idea that's by no means foreign to us. As you know, Joseph Schumpeter, the great Austrian economist of the middle 20th century, introduced the idea of "creative destruction." He argued that modern capitalist economies are extraordinarily innovative precisely because their components constantly go through cycles of breakdown and rejuvenation. When a firm goes bankrupt, its resources, including its human and financial capital, are liberated and reorganized within the economy, aiding the economy's overall adaptation.

But while we might accept this idea – more or less – within modern capitalist economies, we haven't accepted it at all within our social or political systems. Instead, when it comes to our societies and political processes, we try to extend the front loop indefinitely; we try to make sure breakdown never happens. Holling and his colleagues say that such practices simply increase the probability of an even more serious crisis – a more catastrophic breakdown – in the future.

In working with the idea of the adaptive cycle, I have concluded that it's important to add an amendment to Holling's general idea: as a system moves up the front loop, stresses of various forms build. These stresses accumulate because the system learns to displace a lot of its problems to its external environment – quite simply, it pushes them beyond its boundaries. The system might become increasingly competent at managing everything within its loose boundaries, but it pushes away things it can't manage well.

Humankind has done something like this with the consequences of its massive energy consumption: we have pushed untold quantities of carbon dioxide into the larger climate system. Now this perturbation of Earth's climate is rebounding to stress our economies and societies. The same type of phenomenon is visible in our national and global economies: as these economies have grown in recent decades, they have accumulated enormous debts to sustain demand and employment. These debts have essentially externalized to the future the present costs of consumption. Once again, though, the chickens have come home to roost: accumulating debt has recently become a huge stress – in the present – on our economies and societies.

So, while everything may seem to be relatively stable as a system moves up the front loop of the adaptive cycle, underlying stresses – what I've come to call “tectonic stresses” – are often worsening.

Causes of declining resilience in complex adaptive systems

And why does resilience fall as a system approaches the top of the front loop? It appears that three phenomena common to all complex systems are at work. The first is a steady loss of capacity to exploit the system's potential for novelty. A climax forest, for instance, has clusters of species (often including very large organisms) that absorb the majority of matter and energy coming into the forest from the external environment. As a result, very little residual matter and energy is available to support the expression of other possibilities – to support the expression of novelty. Many of the mutations that might have slowly accumulated within the forest's genetic information don't have a chance to express themselves.

Canadian society today offers an interesting analogue: health care. This component of our social and economic system is gobbling up an ever-larger fraction of our total resources, leaving fewer resources to support experimentation, creativity, and novelty elsewhere in our society.

The second cause of falling resilience is the declining redundancy of critical components. As a forest approaches its climax stage, redundant components are pruned away. Early in the front loop, a forest might have, say, a dozen nitrogen fixing species, each of which takes nitrogen out of the atmosphere and converts it into a form usable by plants. At its climax stage (at the top of the front loop), the forest has likely pruned away much of this redundancy, so that it has only one or two nitrogen fixers left. As a result, it becomes vulnerable to loss of those particular species and, potentially, susceptible to collapse.

The similarity to processes in our world economy is striking, although the data are somewhat anecdotal. As the world economy has become more integrated, we have seen a steady concentration of production in a relatively small number of firms – analogues of a forest's nitrogen fixers. Two companies make all large jet liners, three companies make all jet engines, four companies make 95 percent of the world's microprocessors, three companies sell 60 percent of all tires, two manufacturers press 66 percent of the world's glass bottles, and one company in Germany produces the machines that make 80 percent of the world's spark plugs. I think it's safe to say that redundancy has been pruned from the global economy in the same way that Holling observes in ecological systems.

Third and finally, as a system moves up the front loop, rising connectivity increases the risk of cascading failure, which in turns lowers resilience.

For these three reasons, resilience eventually falls as complex adaptive systems mature. But in our contemporary world, we have something else happening too. As I've already noted this evening, our global economic, social, and technological systems need almost inconceivable amounts of energy to maintain their complexity, and the steady supply of this energy is now in question. Our global systems are under rising stress at the same time they're moving steadily farther from thermodynamic equilibrium. It's as if we're pushing a marble up the side of a bowl: we have to expend steadily more energy to keep the marble up the side of the bowl, and if that energy suddenly isn't available, the marble will roll back down to the bowl's bottom, which is equivalent to a dramatic loss of complexity.

That's my brief synopsis of Panarchy Theory. I find the parallel between these ideas and what we're seeing in our world quite astonishing. I believe Panarchy

Theory provides us with tools to understand our situation and think more creatively about the challenges we face.

For instance, earlier I remarked that whether we regard complexity as a good or bad thing depends to an extent on the stage of evolution of the system in question. Now I can explain what I meant in more detail. To an entrepreneurial actor dealing with a system early in its front loop of development – a period in which rising potential and connectivity are producing novel combinations and exciting innovations – complexity might look like a good thing. On the other hand, to a manager trying to keep a system running at the top of the front loop with its staggering connectivity and declining resilience, anticipating a breakdown because the system has become critically fragile, complexity might look like a really bad thing. It's actually quite difficult to say finally, once and for all, whether complexity is good or bad. We have to say that it depends – on the interests of people involved with the system in question and on the system's stage of evolution.

My interpretation of Panarchy Theory also suggests that we can expect significant breakdowns in major global systems. That statement sounds apocalyptic, and I have received a lot of grief over the years for making such statements. But I receive less grief now than I did ten years ago – which is maybe why I am speaking to you now.

Effective government in a world of complex adaptive systems

At this point, you might ask: In a world of rising complexity, uncertainty, and potential for systemic breakdown, how can we possibly govern?

The challenge, I believe, is difficult, but not insurmountable. There are many things we can do to govern our societies and the world more effectively. First of all, we need to be able to identify when we're dealing with a complex system or problem. I don't mean to suggest this evening that we should jettison all our previous paradigms of system management. Sometimes thinking of the world as a simple machine – or of a particular problem as the consequence of a system that operates like a simple machine – is entirely appropriate. Sometimes a Newtonian, reductionist, push-pull model of the world should guide our problem solving. But we must learn how to discriminate between simple and complex problems, which means we must have the intuition to recognize complexity when we encounter it.

When it comes to dealing with complex systems that are critically important to our well being, one of our first aims should be to increase as much as possible their resilience. As I've explained, systems that are low in resilience – that are brittle –

are likely to suffer from cascading failure when hit by a shock. Such failures can overwhelm our personal, organizational and societal coping capacity, so that we can't seize the opportunities for deep and beneficial change that might accompany a shock. Boosting the resilience of our critical complex systems helps ensure that we have enough residual coping capacity to exploit the potential for change offered by crisis.

Resilient systems almost always use distributed problem solving to explore the landscape of possible solutions to their problems (what complex systems theorists call the "fitness landscape"). We can infer, therefore, that if we're to address effectively the complex problems our societies face, we need to flatten and decentralize our decision-making hierarchies and move our capacity to address our problems outwards and downwards to as many agents and units in our society as possible. It turns out that the dispersion of political power throughout our societies that we've seen recently is good, because this dispersion, if properly exploited, can aid distributed problem solving.

In short, the general public must be involved in problem solving. Innovation and adaptation should be encouraged across our population as a whole. Governance – as opposed to government – involves the collaborative engagement of the public in addressing common problems. And this engagement should involve lots of what Buzz Holling has wonderfully called "safe-fail experiments." Such experiments are generally small; if they don't work, they don't produce cascading failures that wipe out significant chunks of larger systems that are vital to our lives. I will return to these ideas shortly.

If we want to boost our resilience and prepare for crisis, we also need to generate scenarios for breakdown. We need to look into the abyss a bit – to think about how breakdown might happen and what its consequences could be. Doing so will help us make more "robust" plans for a highly uncertain and non-linear future. Robust decision making involves developing plans that should work under a wide range of future scenarios. The plans aren't tightly tailored to, or specified for, a particular possible future; instead, we intend that they'll produce satisfactory outcomes across many different futures. They are, in short, robust across a number of possibilities.

We're in a world of unknown unknowns and have only the most threadbare understanding of what might happen – even just a couple of years in the future. But we can still use our imaginations effectively. Whatever happens in the future probably won't map precisely onto any one scenario we develop now, but events might ultimately resemble a combination of two or three of our scenarios. We need capacity to respond that can be applied across many different scenarios.

Although it's not part of the conventional understanding of robust decision making, I would argue that this approach should also include preparing ourselves to exploit the opportunities created by crisis and breakdown, as I mentioned before. We can't always prevent breakdown, nor should we want to. People responsible for managing our public affairs, especially those in our public services, don't want to acknowledge this reality, because they believe their job is to make sure that breakdown and crisis never happen. Alas, significant, even severe, breakdown is going to be part of our future. Instead of denying this fact or desperately trying to figure out how to keep breakdown from ever occurring, our public managers should think about what our societies can do at moments of crisis to produce deep and beneficial change.

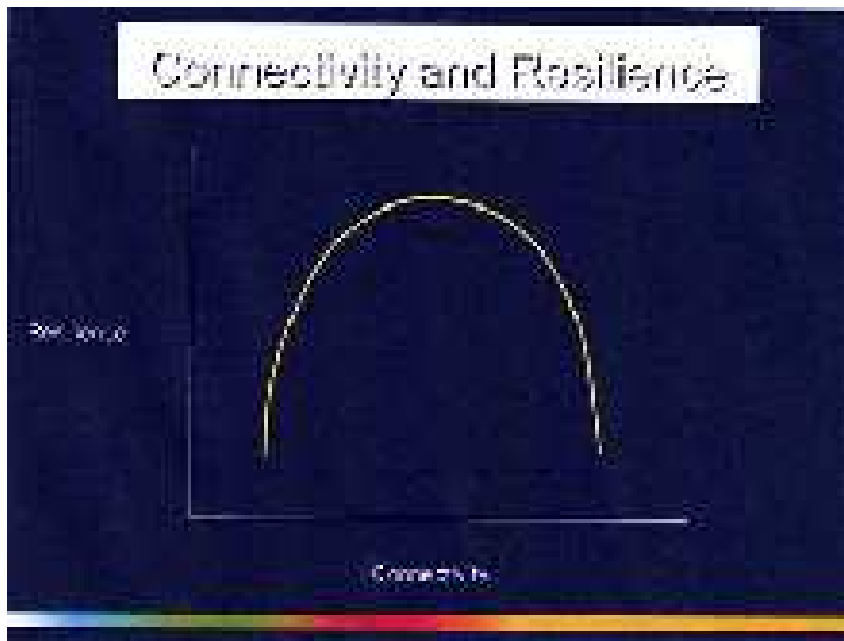
These are moments of high contingency and fluidity, when people are scared, worried, and looking for answers, and when conventional wisdom and conventional policies have lost credibility. We're going to be much better off if we think now about what we're going to do then, than if we produce ad hoc responses only when the crisis is upon us.

Governing to increase resilience

I'm going to propose a few possible scenarios. (This is where I indulge the apocalyptic side of my temperament.) Only a decade ago, these scenarios would have seemed entirely implausible; today they seem, unfortunately, much more realistic. I'm going to focus particularly on three circumstances in which a proximate shock leads to a cascading failure in a tightly coupled, brittle system.

First, Israel and Iran go to war. Israel's fighter bombers attack Iran's nuclear facilities, crossing Saudi airspace to get there and back. Iran responds by launching missiles at Saudi oil installations and by blocking the Straits of Hormuz – immediately taking 17 million barrels of oil a day (about 20 percent of global consumption) off the world market. In Canada, because of gaps in our domestic pipeline network that prevent Alberta oil from being shipped east, much of Ontario and Quebec experience an absolute shortage of fuel. Within two weeks of the beginning of the crisis, 30 to 50 percent of the gasoline stations in central Canada close – curtailing food shipments, emergency services, and all economic activity.

What is your response going to be? I don't just mean your emergency response – your coping response – but your larger, longerterm response. How are you going to use the crisis as an opportunity to begin the hard process of reconfiguring Canada's energy supply system to make it more resilient?



Here's a second scenario: terrorists launch a major radiological attack in Washington, D.C., American officials believe the attackers have come from Canada, so they close the US-Canada border – not just for a few days but for weeks. What is your response going to be? How are you going to use the incident as a chance to reconfigure the Canadian economy so that it's more resilient and more self-sufficient in a future where trade and intercourse could suddenly be curtailed again?

And finally, a third scenario – one that's not even on the margins of conversation at the moment, yet is also quite plausible. Because of climate change, China experiences three consecutive years of drought. The result is a 20 percent shortfall in the country's grain production. After China has exhausted its reserves, it enters the international grain market to buy 100 million tonnes of grain. But only 200 million tonnes of grain are available on the international market annually, so the Chinese intervention produces a sudden doubling or tripling of core food prices around the world. The consequences include major violence in developing countries and a significant political crisis in Canada.

What are you going to do? How can Canada reconfigure its food production system so that it's more resilient? Fundamentally, this would involve making it more autonomous, because resilience is largely about boosting autonomy. I don't mean complete autonomy. I'm not talking about autarky, but rather about loosening the coupling between our critical systems, such as our food system, and the rest of the world.

This point brings me directly to the contentious question of how much

connectivity we want in our critical complex systems. I have concluded that resilient systems exhibit what I call “mid-range coupling.” They aren’t too disconnected, and they aren’t too tightly connected; they’re somewhere in the middle, as you can see in the accompanying figure.

There was a common perception in the 1990s that regardless of what variable you’d like to maximize on the left axis of this figure – well-being, prosperity, or resilience as I have here – the relationship between connectivity and that variable was more or less reflected by a line from the figure’s bottom left to its top right. In other words, most people believed that the greater the connectivity within and between our societies and within and between our critical systems, the better off we all were. In the tough intervening years we’ve learned, though, that beyond a certain point – beyond the middle of the range – connectivity actually starts to produce negative consequences of the kinds I have described this evening, including unexpected interactions, rising potential for cascading failure, and declining resilience overall.

When connectivity is low, increasing it can improve things. In a loosely connected agricultural region, for instance, greater internal connectivity allows sub-regions that suddenly can’t grow food to reach out to the rest of the system to get the food they need. But if the overall agricultural system becomes too tightly connected, it will become increasingly vulnerable to cascading failures in which a shock, like the sudden emergence of a pathogen, spreads from its entry point throughout the entire system.

Leadership in a world of rising complexity

I’m going to focus the concluding portion of my presentation on how we lead in a world of rising complexity. This issue has been at the heart of the New Synthesis project.

I have to admit that I’m now treading in somewhat unfamiliar territory. But in reading the documentation for the New Synthesis project, I came to understand that our public service confronts a problem of “entanglement” of principles of compliance with measurements of performance. Increasingly we are using objectified measures of how people perform within our public services as a way of establishing firm control over their actions. This entanglement of compliance and performance appears to be instilling a culture of fear within our public services.

I was quite struck by this quotation from the project’s documentation:

In an environment where ‘what gets measured gets attention’ and where trust is

low, complex services are difficult to manage. Instead of focusing on the whole issue and program, public officials aim to avoid censure by concentrating on those specific aspects that are being measured. When employees are motivated to save face and seek out the maximum score in this way, at the expense of tackling the complex issues in an innovative way, an optimal environment, consisting of supportive behaviour and operating autonomy, which is a key to effectiveness, is lost.

Looking at this situation from the outside, I have been struck by the fact that the culture of compliance now dominating the public service reduces the possibility and potential for experimentation. We need to reform that culture. Our public-service leaders need to be constantly probing the critical systems we depend upon to determine patterns in the changing solution landscape. They can't know exactly what will happen in these systems in the future, so they should engage in interventions to gather information. They can use small safe-fail experiments as probes to help everyone – leaders and the public alike – learn how the landscape is changing.

More generally, leaders should be “gardeners” who create conditions for experimentation and for – as Mel Cappe argued many years ago – creative failure. At the moment, it seems, there's very little possibility for creative failure in our public service. In fact the very idea probably sends shivers up your spine. You might think: “Wow, that would be terrific, but how can we possibly do it?”

Actually, I'm not sure that the public service is the best advocate for experimentation within its own ranks. You'll always appear self-serving, because you'll always appear to be trying to loosen the constraints upon yourselves and your organizations. It's really up to people like me to tell the general public that the popular obsession with governmental efficiency and with ensuring that government be error-free is producing exactly the opposite of what everyone wants. We're getting a timid, risk-averse, conservative and conventional public service with crippled morale, whereas we desperately need a creative, nimble, flexible public service that can help lead a creative, flexible, innovative, and resilient society.

People like me have to make that case to the public, because people like me don't appear to have a particular interest one way or the other.

I'll finish with one last related and important point. The public not only needs to understand the importance of experimentation within the public service; it needs to engage in experimentation itself. To the extent that the public explores the solution landscape through its own innovations and safe-fail experiments, it will

see constant experimentation as a legitimate and even essential part of living in our new world. To the extent that the public understands the importance of – and itself engages in – experimentation, it will be safer for all of you in the public service to encourage experimentation in your organizations.

Ultimately, the public must acknowledge a basic fact of life, something everyone learns the hard way in their personal lives: we learn more from failure than from success, and failure can be the most creative process of all if we take the right lessons from it.

Ultimately, then, we have a critical task of education. All Canadians must understand that we now live in a world that is, in its deepest essence, complex and turbulent. And all Canadians must accept that we can prosper in that world only if all sectors – public, private, and non-governmental – are constantly engaged in collaborative experiments in new ways of living.

Thank you very much.