

THE GREAT TRANSFORMATION –
CLIMATE CHANGE AS CULTURAL CHANGE

Thomas Homer-Dixon
in Essen, Germany
June 8-10, 2009

I'm delighted to be here with you this evening, in part because this is my first visit to the Ruhr. This region has an extraordinary history as a crucible of an industrial revolution that was, of course, powered by coal. And coal is a substance that will likely become increasingly important in coming decades as conventional oil becomes scarcer. That's bad news from the point of view of climate change, because burning coal releases much more carbon dioxide for every unit of energy generated than burning other fossil fuels. So this region has fascinating connections to both our past and our future.

I'm delighted to be here too, because I believe this conference is focusing on exactly the right subject. I have been working on climate change issues for over two decades – one of only a handful of social scientists with such long experience in studying the societal implications of climate change, although many more now populate this discussion. Over these two decades, I have come to realize that the solutions to our climate-change crisis will ultimately reside largely at the level of culture.

A critical conversation about climate change is going on right now through the UNFCCC process; a key stage in this process will be the Copenhagen meeting at the end of this year. This conversation, to the extent that it is prescriptive, generally emphasizes technology and economics. It stresses strategies for dealing with the climate problem that involve technical aspects of, for instance, societies' energy mix and energy efficiency. I don't want to disparage these approaches or suggest that they shouldn't be pursued. But, the fact remains that despite all our efforts we seem to be falling further and further behind.

The climate change problem is getting worse. Leading scientists tell us that carbon sinks around the world appear to be saturating and that the Arctic is changing much faster than was expected. With respect to variables such as sea-ice loss in the Arctic basin, we are at least 20 and perhaps even 50 years ahead of model predictions. And when we look at the proposed technical, institutional, and economic solutions that

are part of the current international climate change dialogue, they seem radically insufficient in light of the problem's rapidly changing character. They don't seem to offer much prospect of getting us close to where we need to be with anything like the speed we need to get there.

Real solutions, I believe, ultimately reside at the level of culture broadly defined – that is, at the level of our deep values and our deep beliefs about how the world around us works. Now, one might say: "Well if that's the case, then solutions will be beyond reach, because these things change only over generations." But I don't think we should listen to that counsel of despair. I'm convinced that in our world today – partly because of the remarkable communication technologies we have available to us and partly because of the extraordinary analytical capability available to people around the world – cultural change can happen far faster than it has ever happened in human history.

So I will spend a few minutes this evening talking about some of the components I believe such cultural change must incorporate. I can't be exhaustive. The topic of this conference covers an enormous range of material and ideas and questions. I'm going to focus on just a few ideas, and I'm going to be quite provocative. I expect you will disagree with some of what I will say. But my intention, as much as anything else, is to frame the discussion of this matter, to push the boundaries of that discussion a bit, and to help stimulate a terrific conversation at this conference tomorrow.

Right at the beginning I would like to remind everybody of the seriousness of the situation we're facing right now. At the time the 2007 IPCC reports were issued, scientists generally thought that the IPCC's A1B emission scenario was the best representation of what might be called "modified business as usual." Two years later, actual human emissions of carbon into the atmosphere are far above the A1B level. In fact, they are well above A1F1 emission scenario, which is the worst-case outcome postulated by the IPCC. If we continue along our current trajectory, total emissions of carbon will rise about three percent annually, and the atmosphere will contain well over 800 parts per million of carbon dioxide by the year 2100 – that's more than triple the preindustrial concentration.

We are, in other words, going in diametrically the opposite direction of where we should be going if we want to get the climate problem under control.

At its core, the climate problem is fundamentally an energy problem. But it's an energy problem in ways a lot of people don't recognize. Most people think it's an energy problem because the carbon released from the combustion of fossil fuels significantly drives climate change. Of course that's true, but it's also an energy problem because we are probably close to the peak global output of conventional oil.

Even conservative experts now acknowledge that that a peak in conventional oil output has arrived for non-OPEC countries. The world increasingly depends for its oil-production "surge capacity" on OPEC countries. But that extra capacity is not infinite. Sometime within the next ten years we will see the peak in total output of conventional oil – we may be seeing it now – and a steady decline after that. This change will sharply boost our use of coal. As oil drillers have to go farther into more hostile natural environments to drill deeper for smaller pools of lower quality oil, they will have to work harder for every extra barrel. Oil's "energy return on investment" – that is, the amount of oil energy obtained for every unit of energy invested to get that oil – will steadily drop, which will make coal progressively more attractive.

Indeed, in the last six years every single region in the world has moved from a steady decarbonization of its energy supply – a trend that extended over a period of almost two centuries as humankind moved from wood, to coal, to oil, and finally to natural gas as a principal energy source – to a re-carbonization of its energy supply. This turnaround has been significantly driven by the increasing energetic cost of finding conventional oil.

Yet there is another and even more fundamental reason why our climate problem is an energy problem. If we study the demise of the great civilizations through history – and one that I have studied quite closely is the collapse of Western Roman Empire – we find that in many cases they ran into severe economic, social, and political problems when they started to run out of cheap energy – that is, when these civilizations' energy return on investment (again, the amount of energy they got back for every unit of energy they used to produce that energy) started to drop fast. And that's the type of transition that we're about to enter now. We've used up an enormous pool of extremely cheap high-quality energy, and from here forward

energy is going to be harder to get and much more expensive.

Human beings deal with their problems, mainly, by applying energy to them. They create more complex technologies and institutions that require more energy to invent, implement, and maintain. Just at the time when we are going to need enormous amounts of energy to cope with increasingly difficult problems like climate change, we're entering a transition from a world of abundant, cheap energy to a world of scarce and much more expensive energy. In particular, we will need a lot of energy to address climate change: we are going to have to rebuild coastal infrastructures, move people away from coastlines, drill deeper for water, move water from newly wet areas to newly dry areas, desalinize water along coast lines, and pump billions of tons of carbon dioxide underground. All of that activity will demand huge amounts of energy right at the time when energy is becoming much more scarce and much more costly.

For the most part, our policymakers do not grasp the significance of this fundamental contradiction. It means, I believe, that we are entering a time of enormous turbulence, in which instabilities of various kinds will develop in global systems. There will be crises, including systemic breakdowns of economies and agricultural systems. Sometimes whole societies will succumb as converging climate, resource, and economic stresses produce internal instability and violence.

It may be a time of crisis, but it doesn't have to be a time of catastrophe. It's in times of crisis that human beings are often most creative and ingenious and that they pull together most effectively to solve their problems. I am convinced that we won't really address the climate change problem until it produces some major shocks or instabilities that mobilize broad populations.

Much will depend ultimately on whether we're prepared to seize the opportunities these shocks create – opportunities that arise at what I call "moments of contingency" when people are scared, angry, and prepared to consider alternatives to the status quo. In moments of contingency, much will depend on whether people who have thought in advance about these challenges are able to suggest avenues for change that are appealing and constructive and that lead to a more humane global society. Only at these historical junctures will we have the potential to achieve sharp, non-linear shifts in fundamental aspects of our societ-

ies – in their cultures, patterns of social behavior, and institutions.

The thinking that you are undertaken at this meeting can help prepare the ground for that nonlinear transformation – the great transformation – that will come in the future, likely starting in a time of crisis.

This evening I will focus on what I think are four essential components of the coming great transformation. Again, I could generate a very long list of components, but these four will give you a sense for the range of possibilities. The transformation must incorporate:

- a cognitive transition,
- an economic transition,
- a political transition, and
- a normative transition.

I will spend a few minutes talking about each.

The cognitive transition

The cognitive transition will help people better understand the implications of complexity and in turn help societies better cope with complex problems like climate change. By complexity I mean something technical, and I want to spend a bit of time explaining the term.

Most fundamentally we need to move from believing our world is composed mainly of machines to understanding that it's composed largely of complex systems. When I explain this idea to my students, I take into class an old mechanical windup clock that sits on the mantle in my office. I note that I can disassemble the clock into its component parts – its bushings, springs, cogwheels, and screws, and I can understand how all the bits and pieces work. I can understand exactly how they fit together. And when I put the parts back together, if the clock isn't working properly – and, by the way, it means something to say that the clock isn't working properly – I can point to one or two parts where the problem resides. Something is bent, or something is not on the right place.

In other words, I can have a very accurate and complete understanding of this machine.

Many of the challenges we face in the world today arise in significant part because we think the world is composed mainly of machines like this clock, instead of recognizing that it's composed largely of complex systems. Machines like my clock can be taken apart, analyzed, and fully understood. Ultimately, they are

no more than the sum of their parts. They exhibit “normal” or “equilibrium” patterns of behavior, according to how we define “normal” for that particular machine. Most fundamentally, they show proportionality of cause and effect. Small causes cause small effects. Big causes cause big effects. And because of these characteristics we can often predict and manage the behavior of machines very precisely.

I believe that this kind of understanding has been at the root of disasters such as the collapse of the huge cod fishery off the east coast of North America in the 1980s and early 1990s. That fishery was the most productive ecological system in the world. It produced more biomass annually than any other ecological system. Today it's gone. Many very smart Canadian scientists were in charge of managing a significant portion of that fishery, and they thought they understood how it worked really well. Based on this understanding, they estimated the fishery's annual “sustainable” yield – the total mass of cod, in hundreds of thousands of tons, that could be harvested each year without damaging the productivity of the underlying resource. But it turned out that many other factors affected the fishery's behavior and productivity, like changes in ocean temperature and salinity, variations in seal predation, and increasingly heavy harvesting beyond Canada's exclusive economic zone by foreign factory trawlers. In the context of these other factors, Canadian extraction was far too large, and the system flipped from one equilibrium to another.

Complex systems like that cod fishery are ultimately more than the sum of their parts. They have what specialists call “emergent” properties: when their parts are put together, things happen that one wouldn't expect even with the complete knowledge of each individual part. They can flip from one pattern of behavior to another – that is, they have multiple equilibriums. They show disproportionality of cause and effect: small causes can cause enormous effects, but sometimes really big changes in the system don't seem to produce any effect at all. This is what specialists mean by non-linear behavior.

As we found with the cod fishery, systems exhibiting these characteristics can't be easily managed, because their behavior can't be easily predicted. Going back to a distinction made originally by Frank Knight – a leader of the department of economics at the University of Chicago in the 1920s through 1940s and one of the most prominent American economists of the time – we need to distinguish between risk and uncertainty. We also need to recognize that as our world in many respects becomes more complex – or

as we increasingly perturb complex systems Earth's climate – we're moving from a world of risk to a world of uncertainty.

In a world of risk we have data at hand on which we can base probabilistic judgments about what pathways a system we're interested in – such as an economy, an ecology, or Earth's climate – might move along in its future behavior and about what the costs and benefits of moving along those pathways might be. In a world of true uncertainty, we don't have data to make such judgments. In the famous words of Donald Rumsfeld, the former Secretary of Defense of the United States, we face “unknown unknowns.”

Now, I'm not a big fan of Don Rumsfeld – as I imagine is true for a lot of people here – but he did say one very smart thing during his career as Secretary of Defense. Ironically, he was probably ridiculed more for this remark than for any other. You probably recall the incident. At one point when things were not going well in Iraq, he commented: “We have known knowns, we have known unknowns, and we have unknown unknowns.” The press had a field day with this remark, of course: “Don has certifiably left the planet,” was the general consensus. But, it turns out, Rumsfeld was absolutely right. And it also turns out that this concept of unknown unknowns is so well understood in the American military that it's abbreviated to “unk unks.” People who have fought on the battlefield know that surprises come out of the blue. Things come out of the blue. Here in Germany, you might remember that Carl von Clausewitz, in his famous 19th-century treatise *On War*, talked about “friction” on the battlefield and the “fog of war.” Clausewitz understood that surprises happen – that, in the vernacular, stuff happens that you can't possibly anticipate. In a world of unknown unknowns, we don't even know what questions to ask; we're ignorant of your ignorance.

Today we're in the midst of an economic crisis of extraordinary proportion. I believe one significant yet wholly overlooked cause of this crisis was the assumption among members of the finance community – at a deep cultural level – that they were operating in a world of risk. They didn't realize that they were increasingly operating in a world of uncertainty. Many of the most technically sophisticated members of this community came out of the mathematics and physics departments of institutions like the one where I received my PhD (the Massachusetts Institute of Technology). They applied the powerful analytical tools they had learned as graduate students to estimate precisely the risk associated with financial inno-

vations like collateralized debt obligations, structured investment vehicles, and credit default swaps. Over time, assets produced by these innovations propagated into personal and corporate balance sheet far and wide and created an enormous edifice of credit that spanned the whole planet.

Nearly everybody in the financial community thought these developments posed little danger, because they thought they (or at least someone somewhere) understood the risks. They didn't recognize that as they together piled layer upon layer of these innovations on top of each other they were creating a world of unknown unknowns. But when people outside the financial community started to recognize what had happened, they lost all confidence in the value of all those assets on all those balance sheets, and the world economy quickly tipped into a breathtaking abyss. People wouldn't trade with other people, they wouldn't engage in an economic activity with counterparties, if they couldn't be certain whether their counterparties could fulfill their obligations. This was the deep source of our economic crisis.


Notice that I'm suggesting that many of the causes of the crisis lay at the level of people's perception of the world's most basic features. That's why I call this first transition we must make a cognitive shift: It's a shift in our deepest world-view – that is, in what philosophers of science call the “ontology” we use to understand the building blocks of the world around us.

Climate change is an example of what specialists sometimes call a “wicked problem.” Such problems almost always emerge from systems that are complex in the way I have defined the term here. These systems tend to have three defining characteristics: high uncertainty and nonlinear behavior, which I've already discussed, and time lags of indeterminate length in response to perturbations. These time lags might be short or they might be long, and one usually doesn't know how they're operating in the system. I'll illustrate these three characteristics with specific reference to climate change.

Let's look at uncertainty first. Although our climate models do give us a fair amount of certainty about what might happen in the short term, as we go further out in this century uncertainty increases rapidly. Figure 1 shows an illustration from the most recent IPCC report. On the left we see a set of superimposed bell curves representing probability distributions. Each curve is a probability distribution generated by a general circulation model using the IPCC's A2 emission scenario. Researchers entered the A2 emis-

sion figures into a general circulation model and then ran the model repeatedly to generate a probability distribution of likely temperatures. On the right of the figure, we see a set of superimposed probability distributions for 2090 to 2099.

Figure 1.



For the 2020 to 2029 time frame, we have a pretty good consensus about what the average warming is likely to be – somewhere just under 1 degree Celsius. The probability distributions are tall and narrow. But as we move to the end of the century they flatten out. The distributions develop what statisticians call “fat tails,” and they seem to have particularly substantial fat tails on the high side.

There’s a vigorous debate in the economics community right now – the principal protagonists being Martin Weitzman of Harvard University and Bill Nordhaus of Yale – over the implications of this uncertainty. Weitzman argues, and I agree with him, that if we take these models’ findings seriously – that is, if we accept that there’s an indeterminate but nonetheless significant possibility of a warming of four, five, six or even seven degrees Celsius this century – then we should be doing everything we can to avoid that outcome. Put simply, avoiding a fat tail outcome should drive public policy.

Yet almost all our discussions about the economic costs of climate change this century assume that we will get something like the mean predicted warming of about three degrees. They resolutely ignore the possibility that we might end up in a fat tail

Discussion of fat tails brings us to the Arctic. I mentioned earlier that climate change seem to be happening faster than expected in the Arctic, where developments put us 20 to 50 years ahead of model projections – a reflection, again, of uncertainty in our understanding of the climate system.

You probably recall the dramatic acceleration of sea-ice loss in the summer of 2007, which produced a record minimum in sea-ice extent in the Arctic basin that September (2008 showed a slight recovery in the September minimum). This event was possibly an indication of the non-linearity of the climate system. Some scientists think that ocean and air circulation patterns over the Arctic have fundamentally re-organized or re-configured themselves, in the same way that the cod fishery off the east coast of North America re-configured itself and hasn’t returned to its previous equilibrium. In the course of relatively few years, we might be seeing a fundamental nonlinear shift in a major feature of the earth’s climate, the cryosphere, which consists of the system of ice, mainly at the North and South ends of the planet.

Finally, there is the matter of time lags – something that few people outside the scientific community adequately grasp. Our climate and energy systems are replete with lags. There is, for instance, the lag between changes in our emissions of carbon and the climate’s response, which might be as long as centuries. We also have lags between policy decisions to change our energy infrastructure and the actual completion of this change. The shift to a zero-carbon world economy is going to take at least a century. Every time we build a coal-fired electrical utility plant, we’re building into our energy infrastructure at least fifty years of coal consumption. As of last year, the Chinese were installing about 1500 megawatts of coal-fired electrical generating capacity every week, or somewhere around 70,000 megawatts a year, which is roughly the equivalent of the total electrical production of France. That commits the Chinese to huge long-term carbon emissions, and it’s going to take a very long time to reverse that commitment.

In a situation of high uncertainty and long time lags, we have great scope for procrastination. But procrastination isn’t a solution, because if we don’t deal with underlying trends and rising stresses, at some point the system we’re dealing with will likely go non-linear – it will probably flip to another state. At that point, it’s very likely that we won’t be able to get back to the system’s previous state.

Just to conclude this point about the cognitive shift required: Changing our perceptual framework – shifting our cognition – is not enough. We must also change our behaviors and our political and economic institutions to cope better in an increasingly non-linear world. I don’t know exactly what this means in practical terms. But I’m convinced that the first step is to help people realize that we’re not living in a

mechanistic world that can be precisely managed but in a complex world characterized by high uncertainty, non-linearity and time lags.

The economic transition

In my treatment today of the economic transition, I will focus on an issue rarely considered in conventional discussions about the economics of climate change. These discussions quite rightly emphasize the importance of internalizing external costs. They stress that we have to, for instance, start to pay a price for using the atmosphere as a garbage dump for our carbon dioxide.


We pay a price for using landfills as garbage dumps. In my town in Canada, as I'm certain is also true in Europe, I have to pay an amount for every bag of garbage that municipal workers take away on a weekly basis. But at the moment the atmosphere is essentially a common garbage dump that we all use for free. In many, many ways, the external costs of using the atmosphere this way are ultimately going to be borne by our children and our grandchildren.

Internalizing such externalities will be an essential part of our coming economic transition. But another part of this transition is discussed far less – and that's a shift to a steady-state world economy. Yet this shift, which will include efforts to enhance systemic resilience, must be a fundamental part of the cultural transformation that you're considering at this conference. I'll explain what I mean.

I am going to be controversial: There is a common argument, which often comes under the label of the "environmental Kuznets curve," that as poor societies become wealthier (conventionally, as their aggregated GDP increases) the amount of damage they cause to the environment at first increases and then at some point starts to decrease (see Figure 2).

Figure 2.

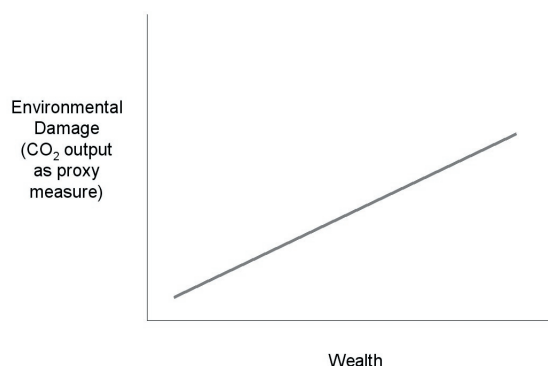
What conventional economists hoped would happen
Growth reduces environmental damage



When we look at the historical data for wealthy countries in Europe, North America, and Asia, this "inverted U" relationship indeed holds for certain pollutants such as sulfur dioxide and lead. But if we examine data for what I think is the best proxy measure of total human load on the environment – carbon dioxide output – the relationship looks more like that shown in Figure 3. For any modern economy in aggregate, as wealth has increased total carbon dioxide output has increased steadily too – maybe not as fast as wealth has increased, but still significantly nonetheless. When we look at the data for carbon dioxide, we don't see anything approximating an environmental Kuznets curve.

Figure 3.

What really happened
Growth increases environmental damage



So, what is going on here? Why is it that we commonly hear one story from conservative economists about the environmental Kuznets curve and yet observe another story when we examine CO₂ output?

I believe that the missing part of the story is economic growth. As our economies grew fast in rich countries in the 1960s, 70s and 80s, we invested in end-of-pipe environmental solutions: We cleaned up our power plants, our sewage outfalls, and to a large extent the emissions from our car tail pipes. But this was just the visible stuff – the stuff that was accumulating in our cities' air, that we could see in our streams, rivers and lakes, and that was palpable to people. As citizens became wealthier in our democracies, and as they started to be concerned about their natural environment, they turned to their politicians and said: "Fix this mess. It could be making us sick. We don't like the look of it. It smells. Clean it up."

In almost all cases, the end-of-pipe solutions that our societies adopted involved higher energy consumption, and because most of our energy comes from fossil fuels, higher energy consumption meant higher

carbon dioxide output. For instance, the catalytic converters we plugged onto the ends of our car tail pipes lowered our cars' overall gas mileage. We had cleaner urban skies but emitted more CO₂.

Still, rich countries managed to improve their efficiency quite remarkably over this period of time – efficiency defined in terms of the amount of energy or material used per dollar of GDP. Rates of improvement of two percent a year were common. Over a period of thirty years, these improvements added up to a very substantial change. Between 1970 and 2000, rich countries saw gains in many cases of 40 percent or more in material and energy efficiency.

But economic growth swamped all these improvements, which meant that these societies' total environmental impact, if we use a measure like total emissions of carbon dioxide, steadily increased. Efficiency improvements of 2 percent a year were overwhelmed by real economic growth of 3 percent a year or more. And because we got rid of all the visible problems – the visible pollution in our rivers and our streams and our lakes and the pollution in the air over our cities – we couldn't see the extra environmental load we were putting on the environment and on the planet. We had transferred our load beyond the horizon in time and space. The people who will pay a price for our increased carbon dioxide output are elsewhere on the planet or elsewhere in time – like our children and children's children in the future.

So we need to ask: Why are we so deeply committed to economic growth? Here we come to some profound cultural issues. We have internalized four equivalencies in our cultures and our societies, in part because they are backed by a fair amount of evidence.

1. *Growth equals solvency.* After the Second World War, all the deeply indebted countries that had fought that war grew out of their debt. In the same way, we often use growth to ease the burden of household debt: if our household incomes and wealth rise over time, our debt becomes a smaller proportion of that wealth, and payments on the debt a smaller fraction of income. We maintain the putative solvency of our pension funds by assuming that economic growth will allow us to meet these funds' obligations in the future. In all kinds of ways, both explicitly and implicitly, we associate growth with maintaining solvency over time.

2. *Growth equals freedom.* This equivalency probably dates back as far as the Renaissance with the development of modern notions of social, intel-

lectual, and scientific progress. It became even more deeply embedded in Western cultures during the Enlightenment. Today, some conservative economists, such as Benjamin Friedman at Harvard University, cogently argue that the only way we can be truly free is to live within societies that experience continual economic growth.

3. *Growth equals happiness.* There's an association in our minds between wealth and happiness: if we get wealthier, we're happier. Researchers argue about whether this relationship actually holds. The evidence is interesting and complex, but there's enough of it to support the claim that there is – especially at lower incomes – a strong correlation between wealth and happiness. Beyond a level of around 20,000 dollars per year per capita the relationship likely starts to weaken, as every extra dollar of income produces diminishing returns (that is, as people's basic needs are satisfied).

And then there's what I believe is the most persuasive and culturally powerful equivalency.

4. *Growth equals peace.* This is a lesson we learned in the 1930s: an economic collapse leads to the rise of political extremism and horrible outcomes like the Second World War. John Maynard Keynes understood this relationship and gave us tools to maintain perpetual economic growth. How ironic that the man regarded by many people on the left as the icon of liberal economists gave us the tools to destroy the planet's environment. And today in this time of economic crisis, the world's central bankers and policy makers are using every single Keynesian tool in their tool kit to sustain consumption. It seems that the only way we can reduce human-kinds' load on the natural environment is to have an economic collapse.

We need an alternative to economic growth that addresses these four real equivalencies. It's a cultural challenge because these equivalencies are deeply embedded cultural assumptions backed by lots of empirical evidence and historical experience. This cultural change won't come easily. When we move this century to a steady-state economy – as we will – we will need to have a clear sense for where we're going. In other words, we will need to know that we can maintain our solvency, happiness, freedom, and peace even though we don't grow.

Keep in mind that about fifty percent of the human population lives on 2 dollars a day or less. For them economic growth is absolutely central and essential.

It can mean the difference between life and death. A steady-state global economy doesn't have to be one in which no one is growing. But if some parts of that economy are growing, other won't be – they may even have to shrink. Fundamentally, a redistribution of wealth, income, and opportunity for growth is essential within the world economy.

So the climate problem – because it's linked to growth – is ultimately a problem of planetary equity, an issue that's rarely discussed. It's ultimately a problem of our world's enormous gaps between rich and poor. John Kenneth Galbraith, the late American economist, noted that growth is the best lubricant between rich and poor. If we don't put in place alternatives and our global economic growth comes to an involuntary halt, perhaps because of the increasing damage caused by climate change and energy scarcity later this century, rich and poor will be at each other's throats, both within our societies and around the world.

What's the alternative? Whatever the alternative, I expect it will involve the notion of resilience. My thinking here is grounded in the ideas of Buzz Holling, a renowned Canadian ecologist. Resilient people, institutions, and societies can withstand shocks without catastrophic failure. In a complex, tightly coupled world exhibiting increasingly frequent and severe system shock, the balance of economic and social investment, I argue, should go towards increasing resilience rather than towards increasing efficiency, productivity, and growth.

In the future, we will see many instabilities and crises. These moments can be opportunities for deep change in behaviors, cultures and technologies, but only if we make sure that things don't fall apart when they occur. Resilience is an essential feature of a world that is able to adjust effectively to crises and exploit the opportunity for real change they present.

I haven't the time to go into great detail here, but briefly resilience means putting aside and protecting reserves; loosening coupling in food, energy, economic, and other systems vital to our well-being; increasing redundancy of critical components in these systems; increasing the diversity of entities and procedures across these systems; and decentralizing decision making. Decentralizing is key because it allows for what Holling calls "safe-fail experimentation" – that is, experimentation that takes place at the local level and doesn't produce cascading failures outwards if it doesn't work. Cascading failure is like a row of dominoes falling over, and it's exactly the phe-

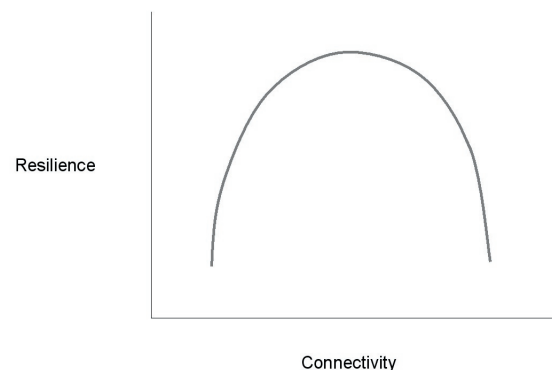
nomenon we have lately seen in the world economy. Tightly coupled systems are prone of such failures.

All of these changes would maximize flexibility in response to an uncertain and rapidly changing world. When we don't really know what is going to happen in a world of unknown unknowns, we need to be able to respond in a wide range of ways to new contingencies.

I also argue, controversially, that there is an inverted U-relationship between connectivity and resilience (see Figure 4). I disagreed with the inverted U we saw in the case of the environmental Kuznets curve, but here I think the representation is quite accurate. In a complex system, resilience (the capacity to withstand catastrophic shock) can increase up to a point as connectivity in the system rises. In a food system for instance, we don't want all the food producing regions to be disconnected from each other. The food system in aggregate needs some internal connectivity so that a producing region affected severely by disease or a drought can reach out to draw the food it needs from the rest of the network.

Figure 4.

Connectivity and Resilience



But beyond a certain point, represented by the top of the inverted U, the risk of cascading failures starts rising. Greater connectivity in the network actually reduces resilience. This is a profound critique, I would argue, of dominant notions of globalization and of the economic benefits of rising global inter-connectivity. In the last two or three decades, rich countries have shared the deep culturally embedded presumption that the more connectivity we have in the world, the better off we all are. I suggest, in response, that we're too far up the connectivity continuum now and that we are seeing the negative results in the form of phenomena like our recent economic crisis.

I will spend less time on both the political and the normative transitions, because I don't have quite as much to say. Nonetheless I believe these two are in many respects as important as the cultural and economic transitions.

The political transition

Human beings have a striking characteristic: As we aggregate ourselves into larger and larger groups, we seem to behave more and more stupidly. And in fact, at the level of the whole world – that is, taking the human population of the whole planet – we seem to have the intelligence of something like a population of protozoa in a petri dish. We're gobbling up all our resources, poisoning up our environment, and poisoning ourselves.

This relationship does not always hold, of course. In certain circumstances, the larger the collectivity of human beings, the smarter the overall behavior. In these cases, there's an emergent phenomenon in which the whole is clearly more than the sum of its parts. Sometimes this emergence is visible in market behavior; other times it is visible in behavior on the World Wide Web and the Internet. Wikipedia, for example, is a remarkable product of the interaction of huge numbers of people, none of whom knows even a tiny fraction of everything that is represented in this online encyclopedia.

One of the deep institutional and political challenges humankind faces in coming decades is to provide for better democratic problem solving – to raise our collective intelligence, so to speak. I don't necessarily mean the formal procedural democracy that we're familiar with in western societies, but real democracy at the level of community decision making. We are not going to make genuine progress solving our large-scale environmental, climate, economic and social problems unless we can mobilize people, and coordinate their problem solving capacity, through new democratic processes.

These new processes might apply some of the lessons that we've learned in the open source movement – the movement that has produced Wikipedia, for instance. There are many potential obstacles to such an endeavor. Somehow, though, we need to create communities that are smarter rather than stupider than the sum of their parts.

The real political challenge here is the following: we need democratic mobilization to increase the power of our publics relative to special interests that are blocking change. This challenge is the same across all

today's economic, political, and social systems, and it's the main reason why we aren't making progress on climate change. It's true in Canada, for instance, with Alberta's tar sands. Canada has an abysmal national policy on climate change because tar sands companies reach right into Canada's federal cabinet to block any kind of meaningful climate policy. Politicians in Ottawa don't have the political capacity to push back.

This kind of political gridlock is visible all over the world. We can address the problem only through a reinvigoration of democracy, in the process getting beyond purely procedural democracy and, perhaps, engaging in some kind of open-architecture process of democratic problem solving.

Here is something that gives me a little bit of hope: At the very time that our species faces some of the biggest challenges it has ever faced in its history – perhaps THE biggest challenges that it has ever faced – we happen to have developed a rudimentary technological infrastructure for species-wide democracy. It's the Internet. But we don't use this technology effectively. Outside of phenomena like Wikipedia, the Internet is mainly a venue for a cacophony of narcissism. We blog at each other, bully each other, and flame away. This behavior doesn't create anything larger than the sum of our parts.

This pathetic outcome isn't inevitable, but changing it is going to require some deep rethinking of what forms of political engagement we as citizens can undertake in our societies.

Any form of open-architecture democracy will face four specific challenges: winnowing, cumulation, preventing hijacking, and managing experts. First, any ideas put forward or solutions to problems that are suggested in the democratic discussion – and there are going to be a multitude of them within any real democratic space – somehow have to be winnowed those down to a few that merit further conversation and focused investment of resources. Then, second, we need to take these remaining suggestions and improve them over time, which is a process of cumulation. Wikipedia, interestingly enough, cumulates very well: in general, the encyclopedia's entries steadily improve over time as people contribute to them. Third, the process must not be hijacked by special interests that are extremely vocal and can focus their resources to block solutions. Fourth, and perhaps most fundamentally, we need to manage experts effectively.

Now, I guess I qualify as something of an expert, and there are many experts in this room. But I'm increasingly convinced that unless the general public is an integral part of the process of solving our climate change and energy problems – that is, I'm convinced that if these problems remains fundamentally technocratic challenges addressed by experts – they probably won't get solved. Our communities of experts cannot generate a sufficient stream of creative ideas fast enough. More importantly, without the involvement of large democratic communities in their development, any proposed solutions, especially those that require significant sacrifice, will not be legitimate enough to be widely accepted. Experts need to contribute ideas and knowledge to the democratic discussion, listen very carefully and respond constructively, but not stand in the way or occlude debate in any way. Unfortunately experts do stand in the way far too often.

Before I leave aside the subject of the political transition, let me tell you a very quick story: During the height of the Cold War in the early 1960s, the nuclear arms race looked like a fundamentally intractable problem – perhaps a lot like the climate change problem right now. But there was a mother in Connecticut who read a newspaper story about some recent scientific research. That research indicated that Strontium 90 from nuclear testing in the atmosphere by the United States and the Soviet Union was accumulating in cow's and mother's milk, which in turn was increasing the likelihood of leukemia in children.

This mother was nursing a baby, and she recognized that she was potentially killing her child. She was outraged. She started a little organization in her community that quickly expanded beyond her community – and remember, this was long before the age of the Internet. Within a few months, hundreds of thousands of mothers were involved around the world, from the United States and Europe to the Soviet Union. Within a year, in large part in response to this pressure, the United States and the Soviet Union had signed the Partial Test Ban Treaty, which put all testing underground. It wasn't a perfect solution by any means, but it was still the first major step in addressing the nuclear arms race.

So don't think that individuals, even those who would seem to be as marginal as isolated housewives in the early 1960s, can't mobilize themselves in large numbers to create major change quickly. In fact, such processes should be vastly easier now, given the communication technologies available to us.

The normative transition

I come finally to the normative transition. This will be in some respects the deepest cultural shift. We need to shift our common understanding of value and to expand the scope of what we mean by "we."

I have come to the conclusion that there are three fundamental categories of value that human beings use. There are, first, what I would call "simple preferences" or "utilities." These values are basic likes or dislikes that might be expressed in statements of the form "I like vanilla ice cream more than chocolate ice cream." Modern economists use this conception of value. Then there are "moral values" or "oughts" that concern issues of fairness mainly in relations between people across space and time – for example, issues of distributive justice. Finally, there are values that get discussed much less: what one might call "existential" or "spiritual" values. These values concern what we believe brings meaning into our lives, what our notion of the good life is, and what we imagine is the relationship between the cosmos and ourselves.

Existential values are connected to the big "why" questions that we all ask when we're young but that, sometime about when we're 10 years old (and I'm sure this is true in every society), we realize from the expressions on adults' faces that we just shouldn't bother, because the questions are too disruptive and annoying. The adults say: "Go to your local religious institution, and you'll find some answers there." But when we walk in the door of our local church or mosque or synagogue, we're not given a place to think about these issues, instead we're given a creed. We're told what to think.

As a result, even in our liberal societies, we don't actually spend a lot of time talking about moral values and even less talking about "existential" or "spiritual" values.

Yet we won't be able to address climate change and energy problems unless we come up with a rough consensus about what the good life is and about what we want as a good life for our children and grandchildren. In the absence of that general conversation in our societies about moral, spiritual, and existential values, utilitarian values fill up all the space. We deal with our existential angst by going to the mall and shopping. Really!

This is now a deep challenge. Ultimately our climate and energy problems come down to what we think would be a good life, spiritually and morally, for our children and our grandchildren and whether we can

achieve some kind of rough agreement on these matters across the societies around the world.

This brings me to my final point, which is about expanding the scope of the “we” –that is, our notion of collective identity and community. Anatol Rapoport was one of the 20th century’s great mathematical psychologists. He was at the University of Toronto when I was a young faculty member there, and he was something of a mentor to me. In my last lunch with him, he said this: “The moral development of a civilization is measured by the breadth of its sense of community.” That statement has stayed with me for many years now, because I think it’s profoundly true.

I’d like to emphasize this point by saying something a bit more personal. My wife Sarah and I have two young children, a four-year old little boy and a one-year old little girl. These children are a source of both enormous worry and enormous hope. They are source of worry and concern because they are the most intimate connection to the future imaginable. Every time I look at Ben and Kate, I feel that I’m looking through a window 80 years into the future. Because of the uncertainty inherent in the complex systems in our world, I don’t really know what that future looks like. But I know enough to be very worried. I’m deeply concerned that theirs may be the first generation in the last two centuries in wealthy countries that will live substantially worse off than their parents’ generation.

But Ben and Kate are also a source of great hope. This hope comes partly from my experience travelling around the world. I have visited 60 countries or so and traveled extensively in several dozen. In these countries, as one does when one travels, I’ve talked to people about their concerns. And I have learned one vitally important thing. Across all our world’s divisions of race, ethnicity, religion, language, cast, and class, and even across the world’s divisions among civilizations that some people insist tear us apart, all people care about their kids. And people actually want more or less the same thing for their kids. Parents say: “Well, we’d like a world, a future, in which our children are safe and secure and in which they can flourish as human beings.” People use slightly different language depending on their culture, but the sentiment is basically the same everywhere.

Now, if we can agree on this fundamental point across all of these societies, we have the basis for a common conception of community that will allow us to address our monumental common problems. We’re not going to start making the difficult choices about

who is going to give up what – or about how wealth and resources might be transferred from one place to another – unless we think of ourselves as part of one “we.”

I think it’s possible. But we have to remember our children. And if we do, it’s conceivable that in 100 or 120 years we’ll live on a substantially warmer planet, a planet that has lost 30 percent or more of its biodiversity, with coastlines that are receding, but one on which, nonetheless, we have achieved a *modus vivendi* among ourselves – the silver lining, in a sense, of climate change – and created a more prosperous world, by some definitions, and a more humane world certainly, a world in which our children and grandchildren can flourish as we have in our lives.

Thanks very much.