

World Changing, November 13, 2006

Worldchanging Interview: Thomas Homer-Dixon

HASSAN MASUM

Why is Thomas Homer-Dixon so worth listening to? There are many writers out there taking on energy issues: Vaclav Smil's works, Out of Gas, Paul Gipe on practical wind power. Society's robustness to breakdowns? Jared Diamond and Joseph Tainter. Climate change? Sure, a few, though that niche is just heating up. Biodiversity and environmental damage? Yes: some of my favorites are Red Sky at Morning, Something New Under the Sun, and the Millennium Ecosystem Assessment. The inappropriateness of focusing on GDP as the default measure of progress? That's an interesting one, with an intermittent thread of scholarship through the last 40 years: Scitovsky's The Joyless Economy, Hazel Henderson, Herman Daly, some recent Ecological Economics. But there are few if any authors who are writing books which cover this range of topics in a sensitive, contextualized, way.

In his new book, The Upside of Down (book review forthcoming), Thomas Homer-Dixon does just that. Many of us here at Worldchanging liked his previous book The Ingenuity Gap. This book takes an even longer view at how we can navigate successfully through societal breakdowns, leaving societies stronger and more resilient.

We wanted to know more about the man behind the book, so he and I sat down for a conversation (distilled below).

With regard to the potential of online tools, what do you see as the next simple step beyond transmitting and sharing information?

One thing we need to achieve is winnowing - we need to increase the signal to noise ratio. But it has to be a democratic process - you can't have people on the outside saying "I like this idea but I don't like that idea, this idea is going forward and that one isn't." Instead, it needs to be internally legitimate, in the sense that the community as a whole decides what ideas are going to be winnowed out, and what ideas are going to go to the next stage.

One of the remarkable things about the Wikipedia environment is that there seems to be a general cumulation of quality - entries tend to improve over time. I had occasion when writing this book to go and look at the entries on thermodynamics, and they were terrific, but I'm sure they're not the result of a single person's contribution. Many people have been contributing, and the quality over time has improved.

I don't think anybody except the diehard advocates would have predicted, 5 or 10 years ago, that you would have been able to have an information source of such high quality that was produced entirely by volunteers, collaboratively. So there is a winnowing and cumulation of quality process there that's very effective. But, and here's where Wikipedia seems to run into trouble, there's the hijacking problem. Especially when you have morally fraught issues, or issues that have strong value conflicts or connotations for people - capital punishment, abortion, the nature of capitalism, some celebrities doing things that annoy people a lot. You get so many divergent interventions that you won't come to a consensus in terms of the entry, and what they've had to do is implement a series of protocols for cooling off discussion or limiting the range of people who can intervene.

Hijacking tends to happen when issues are value-fraught, and a lot of the problems that I think we need to address within an open-source democratic framework will be value-fraught, and so they're going to be vulnerable to hijacking by small groups of highly motivated and not terribly tolerant people who are fixated on one idea, one solution, or one enemy.

When it's possible to replicate your voice easily with the push of a button, hijacking becomes much more of a problem than it does in a personal conversation or a room. It's like somebody in a town hall meeting getting hold of the microphone, and nobody can take it away. So in terms of the institutional design, there needs to be a capacity to legitimately reduce the risk of hijacking, and sideline people who aren't prepared to engage in a cumulative winnowed conversation over time about a particular problem.

I think this is a very important institutional requirement for an open-source democratic decision-making system for dealing with complex social problems. Another is the relationship between lay people and experts. Some of the most difficult problems we're facing - climate change, energy - are technical problems that are enormously complex, and it's very easy for experts to just take over the discussion.

In fact, when I was having a conversation with Paul Martin (the former Canadian Prime Minister) about this issue at one point - this was before he became Prime Minister, and before he was even leader of the Liberal Party - I had a conversation about open-source problem solving. I said, you know, we have this difficult health care problem in Canada - wouldn't it be remarkable to have a hundred thousand people involved in thinking about how to solve that problem? And his first reaction was, well, my thinking would be to get the twenty best experts in the world around the table for a conversation.

Experts certainly have a role, but they can hijack the agenda and deprive the whole process of legitimacy just because they have so much knowledge. So one of the problems with democracy that we have in the world right now is that people just don't think it achieves anything for them - that's why you get participation declining so dramatically in many Western democracies. I think this kind of open-source institutional environment could give people a sense of participation that would be very valuable, but the relationship between the experts and the lay people is critically important. The experts have to provide the information that allows lay people to make informed decisions, without taking over the process.

So I see the relationship of experts to decision-making, and the problems of cumulation, winnowing, and hijacking at the centre of figuring out the institutional design for open-source democratic decision-making.

Interesting. One issue is that it's easy to have a process where one feels as if one's participating, without actually having input into the final solution. So I'm trying to picture any kind of large institution where we've had 10 000 people, or even 1000 people, giving ideas and had them filtered and used in a way that is actually democratic. Any examples?

No, not really... But Wikipedia's interesting - there are some very smart people who spend a huge amount of time creating entries, monitoring entries, making sure the system works OK. They're not well known, they don't get their name put up in lights, but they serve a very important social function within this apparatus, as a kind of glue that holds the system together. It's a voluntaristic culture - not particularly egocentric or narcissistic, like much else on the Web. So that's the kind of culture we want to create.

Now people still need to feel that they're being listened to and that they can make a difference, but they need to understand that it's a meritocratic system, that there's a legitimate mechanism for improving the quality of ideas over time, and that maybe their idea won't get forward or maybe only a little portion of their idea will morph its way through to the end. I think most people are remarkably responsive if they feel they're actually being listened to - that they're not just saying something and it disappears into a void, which is the way so many of us feel with

our contemporary democracy. You write a letter to your Member of Parliament, and you get a form letter back, and what difference does it make? Better than not getting any answer at all, but you don't really think you're making any difference.

One of the things about Wikipedia is you can see what's going on. You can see the conversations, you can see who the people are - in many cases they put up their names - and that leads to a certain transparency. If you want to see the genealogy of certain ideas, you can archive the whole discussion, see how it's been discussed, see the whole process...

Trace it through time.

Trace it through time, exactly. And if somebody says, I made no difference, then you can say, well let's go back and look at the history - here's a point where someone raised an argument which was decisive in the face of your idea, and your idea just dropped out of the process. Or you might say, well look, your idea contributed to this thread of the discussion, and there it is right there, there's a little bit of it still remaining...that's how it influenced it. In either case, you can't possibly say you had no influence - even in the first case when your idea met a counterargument and dropped out, it still was an important component for the progress of the discussion beyond that point.

And I think ultimately, that's all people want. It's like the person working on the line - a lot of manufacturers have found (and the Swedes in their Volvo plants realized this early on) that it's important to provide some power on the line, so that people who are working in the interstices of the system, the fine-grained detail of the system (building cars in this case) can say, this set of procedures isn't working. This is a problem that's costing us money, it's dangerous, it's reducing the quality of the end product.

And they can bring that into a larger discussion, and then there can be a conversation about how to solve that. Sometimes it might involve fairly large changes in the overall structure of the system, but it's the people on the line who frequently have the best knowledge about why things are going awry. And what I suggested to Paul Martin is that you need to provide those people with the opportunity to make their suggestions. And as long as they think they're being listened to, even if their suggestions don't go anywhere because somebody comes up with a better idea, I think they'll feel much more a part of the system, and they'll be eager to participate.

It would be interesting to have a way of routing those suggestions to the place where they'd do the most good - some sort of "reverse Google", in a way.

Right, that's an interesting idea. But (just to make a jump) the underlying ontological assumption here is that there are emergent properties of these systems - that you can get a lot of people together, and if the institution's designed properly, the intelligence exhibited by the whole is larger than any one individual of the whole.

Unfortunately, I think what's happening with many of our decision-making institutions now is that we're not seeing positive emergence but "negative emergence": the intelligence of the whole is less than any of the individuals. Our societies behave like beasts, frequently - with no thought for the future, often extremely violently, with very little moral or ethical guidance or conscience, and what we want to do is reverse that.

To me, this is about institutional design - it's fundamentally a collective action problem. The greatest things that humankind has ever accomplished have been accomplished by an institutional design that gets people working in the same direction, in ways that are very creative, so that resources and ingenuity are effectively mobilized.

How important do you think it is to have ways of seeing patterns that are not obvious? For instance, you talked about society acting “like a beast” - that might be apparent to you having thought about it..

Well, it's a really important question, and there are a couple of things here. In some respects that question is about values, and in some respects that question is about facts. My interpretation of a society behaving like a beast is first of all a values judgement. I think Guantanamo is beastly behavior on the part of the United States - it's morally bankrupt, and it's also not at all supportive or helpful to the enlightened self-interest of the United States - it's counterproductive, just in a purely narrow political sense.

There are two things happening there. First of all I'm making a value judgement, based on a certain moral code, and that's something that people may well not share - they might come up with a different set of values where the behavior in this case is entirely legitimate, entirely reasonable, and morally appropriate. Now that's an important discussion. We may not be able to resolve our value differences clearly, but we certainly need to be able to understand them better, and see if there's a possibility for some kind of overlap or consensus from which we can build to arrive at a solution.

But the second part of my statement, when I say this is beastly behavior, is in a sense a factual judgement about the consequences of this behavior for American society. It turns people against the United States, it's making American foreign policy a lot harder, it's making Americans more vulnerable to terrorist attacks because it makes so many people angrier in the world and hate the United States. Now that's not a value-based judgement - it's an assertion about the facts on the ground and their consequences. That's something we can have a factual discussion about, and at this point we can bring some experts in.

On the value judgements, the experts can participate a little bit, the moral philosophers can participate, but much of that discussion you can have without the involvement of experts. Yet it's important on the non-value side, on the factual side, that we can have foreign-policy experts from other areas saying, “This is what this policy has done in the Muslim world. This is how they interpret it, this is how they see it.” And that input can have a very important role in us understanding factual, functional consequences.

Our ultimate decision about this foreign policy has to involve both components: the “ought” and the “is”. And it seems to me that an open-source environment could provide the framework within which that's done, if you get the institutional design right. I'm not saying you're going to reach an agreement on everything, but you're certainly going to understand where the points of disagreement are much better, and then you might be able to find “kludges” (to use that old computer science word) - ways of living with those disagreements that allow you to get on and do something everybody agrees is worthwhile.

A sort of state of maximal consensus. And in fact one might hope to find a way of mapping out these factual consequences in a way which was adaptive and predictive, so you could actually see them visually.

Yes, although I'm persuaded enough by complexity theory and so forth that, as I say in my book, I think our capacity for prediction is very limited. But you can certainly define a rough boundary between plausible and implausible.

And scenario development is really important in this - part of the factual exploration would be thinking about possible scenarios for the future. What is Guantanamo and like foreign policy going to do for American well-being in the world, and the well-being of humankind as a whole? And you could chart out a range of scenarios from positive to negative, and have a very vigorous debate about whether those scenarios make sense or not.

Again, if you've got the winnowing and cumulation institutional design, you might be able to come out with five or six scenarios which distill the essence of the debate, and that could have very useful policy implications.

And then you could see your values and the value discussion in the context of those scenarios, and it provides a much more powerful framework for thinking about what decisions we're going to make, and coming to some consensus on those decisions.

What do you see as being some low-hanging fruit for individual action on these kinds of issues?

I've been thinking about this...I would like to see some beta-testing of these institutional designs pretty quickly. I think you need to start with a couple of tractable problems.

One potentially tractable problem that we've thought about here is to design better indicators of social well-being, i.e. alternatives to GDP. It's a technical problem, so experts have to be involved. It also involves complex value issues, it involves complex ontological problems about how you aggregate data and things like that. And we thought of using that task of beta-testing an open source environment to explore the development of alternate GDP indicators - we have [a paper about a methodology for comparing alternative social well-being indicators](#), looking at a large number of them.

There might be only a few dozen people in the world involved in this exercise, but it would allow us to figure out how to make them work together. Have some of our students involved who aren't experts, and have some experts involved - then you have to work out the interfaces, to make sure the experts are providing enough information but are not dominating the process, along with all the challenges I discussed before.

To me, development of alternative social well-being indicators is a very important stage in this overall process, because if we shift from GDP to something else it lengthens the "shadow of the future" - it gives us a tighter, more obvious connection to future generations and to other biota on this planet. That can change the discourse really dramatically - change the whole calculus of values and factual assumptions within which we see human behavior.

It's the kind of thing that's very complex, hard to wrap your head around, and maybe we can create one of these open-source environments where the whole is more than the sum of its parts. So that any expert coming in ends up going away with knowledge that could never possibly have been generated just by that expert, or even with a few other experts together; so that the whole is producing something that is much more valuable than any sub-cluster of people could produce.

That's an excellent idea! And I like too the fact that you're actually, in the process of doing this, looking at how you're doing it, and therefore improving the process of tackling similar problems in the future.

Right. If the process works, you learn something about architectures for open-source problem-solving, but you also get some real progress on designing indicators for social well-being.

For me, one of the most resonant focal points of your book was the dual theme of resilience and catagenesis. I wonder if you agree that a particularly practical avenue is to adopt “low-regrets” technologies and systems?

More generally, what creative new kinds of institutions, customs, or ways of thought would you like to see arise, that could help spur catagenesis on an ongoing basis?

We need to build buffering capacity in our societies and systems that's fungible, that can be moved back and forth between different eventualities.

In the first part of the book, I talk about people's desire to hold on, to keep things the same. But we can't always keep things the same, since we don't have as much control over reality as we think we do. This is very different from being fatalistic. The whole idea of the prospective mind is to develop a new set of customs - proactive, anticipatory, comfortable with change, and not surprised by surprise.

Institutionally, we could build in tax incentives and subsidies for people to make households more resilient. For example, if we have an energy grid that's unreliable, maybe we shouldn't build condo apartments that are totally dependent on electricity for elevators, water, and air conditioning. In some business towers, the windows don't even open without power. This kind of housing is fundamentally reliant on large-scale centralized power production.

But what if our economy provided tax incentives for residents and commercial centers to have autonomous power production? If these kinds of incentives were incorporated into everyday policy - whether transportation, electricity, food or water - our systems would evolve to be more capable of withstanding shocks.

I'm sure if we got smart people around the table to think about this, we would generate thousands of specific ideas. Right now resilience isn't treated as important, so people don't pay a premium for it - and there will be a cost associated with the necessary capital investments, a cost that draws resources away from other things. But you're buying resilience - a positive externality in the system, that benefits everybody to the extent it's there.

It's partly the role of government to provide encouragement to do these kinds of things. Distributed open source problem solving would also be an essential feature of a society which recognized that catagenesis is an important part of adaptation - that you're going to have growth, increasing complexity, breakdown, recombination, regeneration, regrowth, and so forth in cycles again and again. A system able to incorporate those cycles in a natural, “standard operating procedure” kind of way is going to require non-hierarchical distributed problem solving.

Let me take an analogy. Within our market economy, the Schumpeterian notion of creative destruction is manifested every day - the growth of new industrial sectors, their decline and obsolescence, their replacement by new technology. Individual companies will start and grow, but may eventually go bankrupt. This creative destruction is part of the everyday world in market economies - it's part of life, and a reason why market economies are so adaptive.

Somehow we need to take that normative comfort level that markets have to adaptation, and introduce the same kind of culture into our social and political worlds. Remarkably, in economies nobody assumes that things stay the same. But in political systems, everybody assumes change is an anomaly, and to the extent change is allowed, it's incremental and managed.

The last point about this market analogy is that markets are highly decentralized networked systems. Adam Smith's hidden hand is all about how the independent pursuit of one's self-interest can create an outcome that benefits everybody in the long run. So we have a distributed problem-solving system going through controlled breakdowns all the time, that for certain ends produces remarkable adaptivity.

One of the reasons economies and political systems differ in their comfort with change is that we've not always had pleasant experience with revolutions in the past...

It's partly because we overextend the "growth phase", sticking with what has worked in the past for too long - like forest managers who constantly put out small fires, yet by doing so create ideal conditions for a really big fire eventually. There must be ways of doing that differently - of most people being able to live in the vicinity while the small fires are happening.

Revolutions occur because there's an accumulation of rigidity, like the buildup of potential energy along a fault line. Because we haven't had a lot of small breakdowns and adaptations, we get a big one. People say, "we don't want a revolution, they're terrible" - but if we got used to change as in our economies, we might be more prepared to accept moderate, manageable breakdowns.

It's also interesting that although economies are distributed problem-solving mechanisms, they also rely fundamentally on centralized institutions to set and police the "rules of the game".

You've got it - the government provides rules on which the market functions. We're not going to do any of this without the right frameworks within which distributed problem solving can take place - the devil is in the details of architectures. And in the research I want to do in distributed open-source problem-solving, we're going to have to explore a variety of architectures, because they affect outcomes.

If anything gets set up for a very large-scale effort, it will probably need to be set up with the involvement of governments, just as happens in economics - these are public goods. It's possible it may arise independently, but at some point we may need to set up institutions - though it might arise in an iterative, internally validated way by creating rules, rules for overturning the rules, metarules, and so forth, as with constitutions.

I don't think we should jettison the role of government in this sphere. We want to see the emergence of these systems in a way that demonstrates their validity, and sees them used more and more for demonstrating solutions to democratic problem-solving. I'd hope that would be an evolutionary process, but in case of extreme breakdown, we as a society need to be prepared.

You mentioned extreme breakdowns a couple of times in the book, and the distinction between breakdowns from which we can recover and irreversible collapse. Can you talk a little more about that?

I'm arguing that we need to find a balance between stasis and shocks. There's a window of opportunity along this spectrum of system disturbances, between those so gentle they don't motivate, and those so extreme they cripple reaction. One of the reasons for building up resilience is to prevent synchronous failures of multiple systems, which is when the really catastrophic shocks can occur.

I'm at the limits of my imagination here. I'm not sure what the best solutions here would look like, but I'm calling for a community to think about what they would look like.

This includes the open scenario development I talk about. And we want to try to develop plans that are fungible for different outcomes. Not just “if a nuclear bomb goes off in New York, do this” - instead, we want a variety of buffering capabilities in the system.

If you have a stock of canned food in the basement, you don't know what it might be used for. Maybe you didn't have a chance to go to the store, and guests show up. Maybe the power goes off. Maybe there's an awful terrorist attack and the grid goes down for an extended period of time. Maybe avian flu breaks out and the food supply is limited because truckers can't move across the continent...

This sort of action that we can take ourselves is different from the centralized incentive structures that governments could set up to promote resilience - sounds like we need resilience at a whole range of different scales.

Yes, and we want to get away from what my colleague Janice Stein calls the “cult of efficiency”.

This can be done individually, but can also be encouraged through government incentives. For example, geothermal energy can provide the majority of heating for houses once set up - if we were to establish tax structures encouraging homeowners to install those, our dependence on centralized systems like natural gas supplies could be decreased significantly. There's an asset that you get once it's installed - so let those assets be easily paid for and transferable, through tax structures and incentives, and 20-year amortization terms as with mortgages.

So there would be more autonomy. Not 100% off-grid, but there are a lot of intermediate positions where people can be encouraged to provide a bit of their own needs. And over time, that presumption of resilience - that you need to think about resilience in your day to day and corporate actions - will become part of the background assumptions of our lives, and assume a sort of autonomous reality independent of governments and particular incentives.

I'd like to see this ideas of resilience so ingrained in us as individuals that it doesn't seem unusual or remarkable in any way. As a friend of mine said, “This is the way Scottish Presbyterians think!” You live in a harsh environment, there are vicissitudes you can't anticipate entirely. What do you do? You take in stores of food, build strong-walled houses, save for the future, always think about how things could go wrong - recognizing that you don't really understand the world around you, and sometimes it can produce surprises that hurt you a lot.

Good engineers also think that way, with fudge factors and so forth.

Right - they overbuild bridges, they design for worst-case scenarios. They're natural problem-solvers, and want to make sure things work under a lot of unexpected circumstances.

But you see, this comes out of the bottom line, and that makes it tough to retain in a competitive market. There's a tension, because engineers are trained to be resilience conscious - yet the corporate ethos in highly competitive industries is to have a high discount rate. Paying for resilience is somewhat analogous to paying for the externality of ecosystem goods and services - it needs incentive structures to help it along.

Maybe we can encourage a shift in culture through collaborative communities? Value discussion will be central. Our challenge is not just solving technical and engineering problems. It's also our response to other people in the world, with respect to inequality and suffering - principles for distribution of wealth and power guide what we aim for in our society. And the discount rate we use for the future strongly impacts our time horizon for planning and investment.

These are really central issues which we don't have a coherent conversation about. You want to get everybody talking about these issues in parallel, and perhaps evolve toward pointing more or less the same way.

Going back to engineers and resilience, perhaps it would be worth exploring which engineering companies have successfully made it part of their culture, and figure how and why? And you'd think smart investors would be willing to pay for resilience in what they buy, at least where they know they'll be owning the assets for many years.

But everybody turns over assets so rapidly today. Notice how values come in - if everybody assumes that a worthwhile life includes lots of throughput, acquisition, and movement up the social ladder, then you're going to be consuming a lot of materials and turning over assets a lot - moving from a bicycle to small car to midsize car to SUV, with each step carrying social value.

If instead a good life is one where you have time to spend with your family and friends, where you have lots of time for really good conversations, maybe the material aspects aren't so important. Maybe then there's value in living in the same house for 20 years - you establish a network of friends who are really interesting, who drop over for dinner.

So our value judgements have very practical implications for things like discount rates, and acquisitiveness, and the lifespans of material things in our lives.

In a world where hardly anyone lives in one place for 20 years, there's going to be a culture of short time horizons and transience, not decades and generations of experience of a place or family. That's one of the reasons why we moved into a stone house - there's something very permanent about stone. There's something very important about stopping, about thinking. But there's a tension between that and having a prospective mind.

It's sad - I have a house that's 150, 160 years old, a stone house with a lot of old trees on the property. But I know that lots of them may die within my lifetime or my children's lifetime, from climate change. So there's a tension, because I know that this landscape, this space for my family, is going to change from outside pressures. We want to extend people's time horizons, but it's a hard thing to reconcile.

How are you feeling about this book, now that it's complete and about to be released?

TH: It's scary because it's an encapsulation of 30 years of thought, and I know I'm going to take some bruises on this one. One of the reasons I spend much of the book talking about the problems we face in detail is that I'm sick of people dismissing the fact that we're in a serious situation. By the time you finish chapter 8, if you've been listening and thinking, it's going to be hard to deny at least the possibility that there will be serious breakdowns in the future.

The last part of the book is about denial, about what breakdowns might look like, and what we might do about them. [With the potential of open-source democratic problem-solving and resilience], it's the first time I've started to see a glimmer of something that offered a way out. At the end of the last book, I didn't have a sense of what the way out might be.

It was by reading [Buzz Holling's work](#) that I realized what normally seems bad can be an enormous opportunity. It's a radical idea that some people are not going to like, and they may ridicule and pigeonhole it.

And then there's your focus on the urgency of moving beyond GDP as a measure of progress...

Yes, there's a contradiction between biophysics and our economic system. You can only create an armillary sphere with more and more bands for so long before it breaks down.

If we're going to pull ourselves out of this mess, we're going to have to bring together a huge number of people. And there's lots of that's happening, but I don't think the architectures so far are very good for cumulative problem-solving.

There's lots of different jobs to be done, and I'm not a believer in the "great man theory" of history. I collected a lot of information about leadership in the course of my research, but you'll notice in the book that I didn't talk about leaders.

I decided that this is about individuals taking on responsibility themselves, and doing things themselves, in a flat yet structured collaboration architecture that moves decisions away from the apex. We're all leaders in this model.