Extreme Events in Human Society:
The Xevents Observatory and Simulator

John L. Casti
International Institute for Applied Systems Analysis
and
The Kenos Circle
Vienna, Austria

In this plenary talk, I discuss extreme events (Xevents) created by humans, not nature. These include things like terrorist attacks, pandemics, political revolutions and financial system meltdowns. The talk explores the types of methodological tools needed to develop early-warning signals for such events—and what to do with such signals once they are obtained. We also present the outlines for a new research venture at IIASA, involving an Xevents "observatory" for development of methodology and an Xevents "simulator" to serve as a laboratory for both testing of tools, as well as identification of Xevents that have never before occurred.

Keywords: Extreme events; agent-based simulation; extreme events methodology

1. Xevents in Nature and in Life

Consider the following events:

- A virulent strain of the avian virus jumps to humans in Hong Kong, sweeps across Asia, and ends up killing more than fifty million people.
- A magnitude 8 earthquake centered on the Ginza in Tokyo kills two million people outright, with property damage mounting into the trillions.
- Bees around the world begin dying off in massive numbers, interfering with pollination of plants worldwide, precipitating a global food shortage.
- An asteroid ten kilometers wide crashes into the Atlantic Ocean, setting off a tsunami that destroys all life on Earth.
- Iranian terrorists set off a nuclear weapon in Times Square during rush hour, leveling much of Manhattan, killing a half a million people, and permanently reducing New York City to rubble.
- A tanker car carrying chlorine derails in Rio de Janeiro, spilling its content and killing more than five million Cariocas.

This list could be carried on almost indefinitely. The point is that surprising events capable of killing millions, if not hundreds of millions, of humans happen. Moreover, even without huge loss of lives, capital stock is decimated, setting back development worldwide for decades. Not a single one of the items on the foregoing list is impossible. And, in fact, some of them like an asteroid impact or the spill of a deadly chemical have already happened—many times!

Each of these events is what has come to be called recently an "extreme" event, or *Xevent* for short. These are events generally seen as deadly surprises whose likelihood is difficult to estimate, and that can turn society upside-down in a few years, perhaps even minutes. Xevents come in both the "natural" and "human-caused" variety. The asteroid strike illustrates the former, while a terrorist-inspired nuclear blast serves nicely for the latter.

But what really qualifies as an Xevent, anyway? I don't think a consensus has emerged on the question. But two factors are certainly part of the answer: (1) an Xevent is rare, something outside everyday experience, and (2) the event is capable of causing massive destruction of human life and property. In regard to the first property, if sufficient data on the type of event exists to enable us to apply statistical tools to estimate the likelihood of an Xevent taking place, the distribution describing that likelihood will be one with what are termed "fat tails". The default gaussian distribution so beloved by probabilists and statisticians everywhere simply does not apply to measure how often such events are likely to occur. Xevents are a lot more frequent than the normal distribution would lead us to believe.

Evolutionary biologists and geophysicists have noted that many long-term processes are ultimately driven by the Xevents. For instance, the theory of "punctuated" equilibrium in evolutionary biology argues that most of the time evolution isn't doing much of anything. Then along comes an Xevent, such as the Cambrian explosion about 530 million years ago in which most major groups of complex animals appeared over about 70-80 million years, and the whole system is set off onto an entirely new trajectory. There is evidence suggesting that the same phenomenon applies in the human social domain, as well.

Generally speaking, when we use the term "extreme event" we think of something that's negative or damaging to society, as with the events mentioned above. But this bias stems mostly from the linguistic fact that in everyday speech the word "event" refers to an occurrence that's localized in space and time. In other words, it's something that takes a relatively short period of time and is confined to a limited spatial region. Hurricanes and tornadoes are good examples. They unfold over a few hours and generally affect a limited geographical area. With this interpretation of what's meant by an "event", it's not surprising that almost all Xevents are negative, in that they result in many deaths and/or do great property damage and they do it quickly. But if we abandon such a limited interpretation of what constitutes an event, things begin to look rather different.

The first point to note is that events have a characteristic *unfolding time* (UT), a period that measures the time between the start of the event and when it ends. For some events, like earthquakes, the UT might be just a few minutes or even seconds. For other types of events, such as a war, the UT is more likely to be measured in years. Moreover, events are not necessarily localized in space, either. Something like a financial crisis or a pandemic may well encompass the entire world.

When it comes to the effect of an event on society, we have what might be termed an *impact time* (IT) over which the event's effects can be felt. So, for instance, the IT of an asteroid strike may well be millennia, while the IT of a hurricane like Katrina is probably just a few years.

Putting these two notions together, consider the quantity

$$X = 1- UT/(UT + IT)$$

This quantity is always between 0 and 1, which is convenient Here are a few examples to illustrate the basic idea of this measure:

- *A Force 5 Hurricane Striking Miami Beach*: Here the *X* is near 1, with a short UT and a much longer IT. So this is an Xevent.
- *A Force 5 Hurricane over the Caribbean Sea*: In this case, $X = 0$, t since the IT is zero (there is no societal impact, at all). Thus, it is not an Xevent, at all.
- *The Post-WWII German "Economic Miracle"*: This event *might* be an Xevent or not, depending on what you mean by the event versus its impact. Is the event the implementation of the Marshall Plan? If so, the UT is around 5 years, while the IT is probably about 25 years. So the $X = 5/6$, which would characterize this as a "moderate" Xevent.
- *The Development of Agriculture*: Best accounts would measure the UT as about 8,000 years ending around 2000bc. In this case, the IT is still taking place, 4,000 years later. Then *X* would then be around 1/3—and growing. I'd characterize this as a "mild" Xevent, one that's developing into a major one.

The last two examples are Xevents that are by no means negative or damaging, but just the opposite. Here we see why the conventional interpretation that an Xevent always refers to something negative is very misleading. If the UT is short, the overwhelming likelihood is that the impact on society will be destructive for the simple reason that it's a lot easier–and quicker–to tear something down than to build it up (2nd Law of Thermodynamics!). So if you insist that the term "event" refers to an action or occurrence very localized in time, you're virtually forced into regarding it as an Xevent that leads to death and destruction, not growth and better life quality.

Of course, there will be problems with almost any measure like *X* in the sense that examples can be cited when everyday common sense says one thing while the measure suggests something else. So I think *X* as defined above should be taken as a kind of 밝ule-of-thumb? a starting point for deeper consideration of any potential Xevent.

2. Fingerprints of Xevents

A careful examination of numerous natural and man-made Xevents turns up a few distinguishing fingerprints of such events. These include

- *Statistical*: If there exists sufficient data to employ extreme event statistics, the distributions for the occurrence and magnitude of an Xevent is generally from a family of probability distributions, like that *(stable*) *Levy or (stable) Paretian distribution*. These families include the familiar gaussian distribution as a very special case, as the gaussian is the *only* member of the family that has a finite variance. All other members display the fat tail property, which technically means they have an infinite variance. Many Xevents, though, are so rare that we have little, if any, data upon which to bring statistical procedures of any type to bear upon them.

We also note the ubiquity of *power-law* distributions in the world of Xevents. These distributions measure the relationship between the frequency of an event and its magnitude. For instance, an earthquake twice as large is four times as rare. If this pattern holds for earthquakes of all sizes, then the distribution is said to "scale". This means that there is no *typical* size of earthquake in the conventional sense of "size".

It's now known that the distributions of a wide variety of quantities seem to follow the power-law form, at least in their upper tail (Xevents). Scientific interest in power-law relations stems partly from the ease with which certain general classes of mechanisms generate them. The demonstration of a power-law relation in a collection of data can point to specific kinds of mechanisms that might underlie the phenomenon in question, and often suggests a deep connection with other, seemingly unrelated, systems

- *Dynamics:* Systems displaying Xevents live far from equilibrium. This means that system variability and collective effects from the components of the system in interaction are the most important determinants of the behavior of the system.
- *Evolutionary Processes:* As noted earlier, the Xevents shape the direction that the system's evolution takes. Small, incremental changes operate for long periods of time. Then, *wham*, an Xevent takes place and the whole system is shot off onto an entirely different course.
- *More Effect than Cause:* The commonalities of behavior of Xevent systems is focused more on the effect the events have on human life than on the specific cause that gives rise to the event. Deaths from the Xevent are large, financial losses are large, and environmental destruction is large.
- *The Mere Possibility of Disaster or Boom:* In the developing world, the very possibility of disasters, combined with the lack of social safety nets, limits risk taking in the society (i.e., lack of credit due to disaster risk) and appears to be a major contributor to poverty traps. On the other hand, in the developed countries, research and exploration of "possibility space" is likely to lead to new technological innovations that can be expected to lead to an Xevent-generated boom leading to "the next big thing".
- *Policy Response: Society typically* underestimates the chance of an Xevent, even those that are not so unlikely. For example, expenditures on disaster assistance involve investing far more in reacting to disasters than in taking steps to prevent them. Moreover, probabilistic tools are rarely used to support policy decisions involving Xevents.

3. Methodological Issues

Let's now consider a few conceptual issues that serve as a framework for a coherent research program on Xevents.

- ✓ *Anticipation:* Without a doubt, the single most important tool we could develop for dealing with Xevents would be a systematic procedure for early warning of a possible event. As the saying goes, "forewarned is forearmed", and to have a reliable, consistent procedure for anticipating major discontinuities like a financial meltdown or a crash of the Internet would be a huge step toward effectively addressing such crises.
- ✓ *Forecasting:* Some may argue that anticipation and forecasting are the same thing, for all practical purposes. But this is not the case. Anticipation deals with early warning of the *possibility* of an event occurring; forecasting has to do with claims that an event of a certain type and of a certain magnitude will take place at a certain place and time with a given *likelihood*. As noted earlier, the notion of likelihood is a slippery one, especially in the context of Xevents where we have little, or perhaps even no, actual data upon which to base any kind of standard probabilistic forecast. What we need here is more like a

"theory of surprise" than a theory of likelihood.

- ✓ *Trends:* The majority of people operating as "futurists" generally make the following type of forecast: "Tomorrow will be just like today except a little better or a little worse". In short, they are trend followers and simply extrapolate whatever the current trend is into the future. For such people, surprises never occur and trends never change. Strangely, this kind of trend following is almost always *right.* But it's also almost always useless as well, and you certainly shouldn't pay any money to a so-called futurist for this type of forecast. What you should be ready to pay for, though, is information about the *turning points*, those moments in time when the current trend is rolling over and beginning to change. That type of information is golden (and not fool's gold, either). Mathematically, such turning points are called "critical" points and there is a very well developed theory about them in the dynamical systems literature. Oddly, though, that theory has been very little employed for the kind of practical questions about Xevents that concern us here.
- ✓ *Modeling:* Traditional mathematical modeling a la physics will have to give way to what Stephen Wolfram has termed "a new kind of science". This is a science in which computer programs replace mathematical formalisms. The creation of an agent-based modeling (ABM) laboratory for testing hypotheses about Xevents is necessarily an important component of whatever form an Xevents initiative might take, since we require some type of "laboratory" to do controlled, repeatable experiments with systems that we cannot experiment with "in the wild".

To summarize, here's a telegraphic list of the types of methodological foci a possible Xevents research activity might comprise:

- Global catastrophes (including the human-caused variety)
- Early-warning/horizon-scanning systems
- Extreme risk analysis and management procedures
- Surprise, resilience and tipping points
- Methods of forecasting collective social events and behaviors
- Computer simulation and scenario construction as laboratories for studying Xevent
- Development of counterfactual thought experiments in social processes
- Tools for the analysis of the fragility of critical infrastructures

REFERENCES

1. May, R. S. Levin and G. Sugihara. "Ecology for Bankers." *Nature*, Vol. 451, 21 February 2008, 893-895. (Short, provocative report on a meeting organized by the US National Academy of Sciences and the New York Federal Reserve Bank to investigate systemic risk in the global financial system. The article suggests many analogies between ecosystems and financial systems that have yet to be fully explored.)
2. Newman, M., A.-L. Barabasi, and D. Watts, eds. *The Structure and Dynamics of Networks.* Princeton University Press, Princeton, NJ, 2006. (A collection of pivotal articles in network analysis, together with very interesting commentary by the editors.)
3. Albeverio, S, V. Jentsch, and H. Kantz, eds. *Extreme Events in Nature and Society.* Springer, Berlin, 2006.
4. Carlson, J. and J. Doyle. "Complexity and Robustness." *Proc. National Acad. Sci. USA*, Vol. 99, 19 February 2002, 2538-2545. (Overview of the theory of "Highly Optimized Tolerance (HOT)", a theoretical framework showing how networks can be highly structured internally, yet robust and fragile in their external behavior.)
5. Barabasi, A.-L. *Linked.* Perseus Books, Cambridge, MA, 2002. (Popular account of the multiyear odyssey leading to the discovery of the ubiquity of scale-free networks. Along the way, the author provides many insights into why networks are the way they are and not some other way.)