WHITE PAPER

The X-Events Index

John Casti
The X-Center, Vienna

Overview

The X-Events Index is a pilot project involving the development of quantitative tools to characterize and measure when we are in the vicinity of the occurrence of an extreme event ("X-event"). The tools for exploring this question were introduced in the books *Mood Matters* (Copernicus, NY, 2010) and *X-Events* (HarperCollins, NY 2012). The first book emphasized the role of mass psychology in the biasing of collective human social events. The second book focused on how the relative complexity of interacting subsystems of the social domain gives rise to stresses leading to X-events. The goal of this project is to make those tools operational.

The anticipation of an X-event is a bit like weather-forecasting. We take the drivers of atmospheric change, quantities like temperature, air pressure, humidity, wind velocities and the like, and use mathematical equations governing the change of these quantities to come up with a forecast like, "twenty percent chance of rain tomorrow." Here we might think of a kind of landscape of weather events, consisting of sharp peaks, narrow valleys, flat plateaus and the like. At a given time, the current weather is located somewhere on this landscape. But the landscape is dynamically changing, so you may be on a flat plateau right now but on a mountain peak tomorrow. This means that tomorrow a small random trigger could send you tumbling into a new weather domain, while that same trigger today while you're on the plateau would cause no major movement, at all.

The goal of this pilot project is to investigate this landscape picture not for weather events, but for human-caused extreme events, such as the breakup of the European Union, a terrorist movement like ISIS, a meltdown of the global financial system and other such major life-changing events.

To carry out this study, we have to first agree upon what types of X-events we will focus attention upon. Then we have to determine the drivers of the landscape change and see where we are today on that landscape and see how it is changing. Finally, we have to identify when we are nearing a peak on the landscape, where even a small random triggering event can push us off the peak into one or another of the valleys.

The pilot study shows how to attack these questions, and then employs the tools to a case study of the collapse of the European Union.

Now let's talk in more detail about all the foregoing concepts and their application to anticipation of X-events.

The Landscape of Events

Every human event, ordinary or extreme, is a combination of two factors: a context within which the event takes place, and a random trigger that selects from the set of possible next events, the one that is actually realized. So

Event = context + random trigger

To illustrate, suppose you're driving your car down the street and come to an intersection. The space of possible next events contains the possibility of driving straight on, turning right or left, stopping for a pedestrian and a lot of other things. What actually happens is determined by an unpredictable (random) trigger. Possible triggers are (1) nothing happens, in which case you continue on through the intersection, (2) a pedestrian runs out in the street in front of you, whereupon you hit the brakes or turn the wheel, (3) you suddenly remember you have to buy some milk at the food store before going home, so you make a right turn to stop by the store, and so on. So what event actually *occurs* is determined by this random trigger, with the context being you reaching the intersection. Once the trigger has acted, all the other possibilities disappear except for the action you take as a result of the trigger.

Generally speaking, the trigger leads to an event that not at all remarkable, rare or surprising. Namely, you continue on your way without incident. But if the trigger is a pedestrian unexpectedly appearing in front of you, and you happen to actually run into that pedestrian, then the event becomes an X-event, in the sense that it is rare, it is surprising, and it may well have dramatic consequences for you (killing//injuring the pedestrian, financial damages in a lawsuit against you, loss of your driver's license and so forth).

This homely example of driving is a kind of template for every possible event. The context sets the space of possible events, the random trigger picks out a specific event from that space and that is the event that's actually realized. (Note: Readers familiar with the Schrödinger equation in quantum theory will recognize this situation. What we call "context" is the Schrödinger wave function, while our "random trigger" is the act of making an observation of the system, where here I'm using the standard Copenhagen interpretation of the wave function).

Earlier, we visualized the context as a kind of rugged landscape, consisting of hills, valleys, sharp peaks and plateaus, a bit like the landscape depicted below in Figure 1. Ordinary events are like those a point A in the Figure, which shows a small peak. If the peak is flattened into a plateau, then every random trigger just pushes the system into a nearby point and the event that occurs is unremarkable, i.e., ordinary. But since the landscape is dynamically changing, we might find ourselves at point B instead of A, in which case even a small random trigger can send the system spiraling down to one of the lower valleys, giving rise to a much more dramatic type of event, an X-event. The "X-ness" of the event in that case being measured by the height of the drop from point B to the valley.

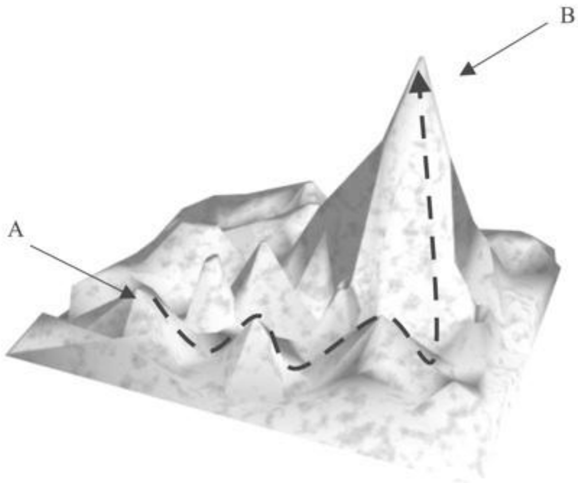

Figure 1. A rugged landscape.

Just to hammer home this point, suppose we're considering the X-event of a major financial crash. In that case, the landscape represents all the possible states for the financial system at a particular moment in time. Point A, then, could represent a state of heightened risk due to political and/or social uncertainties, but is not a point with especially high X-ness. But as the drivers of change work their way, maybe we would find the system at B instead of A, in which case the early-warning signals of a major crash would be sounding their alarm since even a small trigger like a fruit-seller in Tunisia burning himself up in the street or the collapse of a bank like Lehman Brothers could send the system into the abyss.

The take-home message at this point is that it is flat out impossible *to predict exactly* what event will happen next for the simple reason that the trigger setting off the event is *random*, i.e., it has no pattern. So the best we can hope for is to understand the dynamical change of the context and look for ways to see if we're near a point like B and not sitting comfortably at A instead. Our goal is to understand the drivers of change in context, in order to clarify where we are and how close that is to a critical point like B.

Drivers of the Change of the Landscape of Events

How does the landscape of events change during the course of time? In particular, what are the drivers of that change? These are the questions that underlie the feasibility of developing a workable procedure for creation of the X-Events Index. Earlier, we alluded to the drivers of change in the landscape being two elements: The *social mood*, which is a measure of mass psychology in the population, and the *complexity gaps*, which are the ultimate sources of stress in the social system. Let's look a bit more deeply into each of them.

Social Mood

Put simply, the social mood of a population at any moment is the overall beliefs the population holds about its future. If people are optimistic and looking forward to the future, then the mood is positive; if they are pessimistic and fearing the future, then the mood is negative. As we'll see in a moment, the social mood has a very strong biasing force on the types of events that we can expect to see on a given timescale. A positive social mood tends to lead to events that we would describe with words like "joining", "global", "happy", and "welcoming", while if the mood is negative we would expect to see the opposite polarities being dominant, "separating", "local", "sad", and "rejecting".

It's very important here to note the matter of timescale. Each event has its own typical unfolding time. For example, events involving popular culture, such as the type of movies people, the type of books that are popular, the clothing fashions that are in vogue and the like typically unfold on a short timescale, generally a year or less. On the other hand, events such as changes in political ideologies or the outbreak of wars unfold on a longer timescale of years to a decade or two. Finally, there are very long timescale events such as the rise and fall of a global power that take place over many decades or even centuries. So when we speak about the anticipation of a particular X-event, we have to keep the timescale of the event in mind.

Finally, the big question: How can we actually *measure* the social mood? A first thought might be to just use a questionnaire or a survey of the population and infer the social mood from the survey results. There are lots of reasons why this is both an impractical and quite useless approach, not the least of which is that people might say just about anything. What's much more useful is to use a measuring stick that reflects what people actually *do*, not what they say. In short, actions speak louder than words.

For the purposes of this study, I used a financial market index as a way of characterizing the mood of a given population. The reasons are manifold for employing an index like the Dow Jones Industrial Average or the S&P500 in this connection. First of all, any such index represents what investors are thinking about the future, in the sense that if they are optimistic about the future, investors tend to buy, which if they are pessimistic they lean toward selling. Thus, a rising market is a sign of increasingly positive mood, and conversely. Moreover, the financial market data is easy to obtain on all timescales and is reasonably free of measurement error. Finally, such data exists over a very long period of time. As a result, it's easy to do historical studies of social mood employing financial data as the measure of mood.

With all these matters in mind, let's look at a concrete example in order to fix these notions of polarity, timescale and anticipation.

The Middle East Crisis (Long Timescale)

Tension in the Middle East is a bellwether rift that has an almost perfect record of erupting into open hostility right at the onset of major downturns in social mood. Figure 2 shows that virtually every major change in social mood was presaged by either the eruption of active hostilities in the Middle East or by a period of easing of tension and hopes for a better, or at least more peaceful, future. (Note: On the chart, the vertical scale is the DJIA adjusted for inflation via the consumer price index.)

Figure 2. Ups and downs in the Middle East, 1925–2000.

A couple of points on this diagram are of special interest. We see that the longest period of goodwill and cooperation between the warring factions was a Middle East *Era of Good Feelings*, which ran from September 1993 through January 11, 2000—***3 days*** before a major high in the DJIA. By the summer of 2000 tensions were rising, as the Palestinians threatened to declare statehood. As the social mood declined through the fall of 2000, riots were commonplace and Yassir Arafat's popularity underwent a major shift upward. By May 2001 newspaper headlines were asking, "Is It War Yet?" and Israel employed US warplanes against the Palestinians for the first time since the 1967 war. As stock prices were in a downturn in September 2001 that would only reverse one year later in October 2002, it would certainly be consistent with the idea of a major war breaking out after a trend change in social mood. So where do we stand today on the prospects for the Middle East during the next few years?

Recall the general principle operating here: good things tend to happen in the Middle East when the mood is positive, bad things when it's negative. But now let's get a bit closer to the action and look not at the US Dow Jones Industrial Index but at the index of the Israeli stock market in Tel Aviv. Figure 3 shows this index over the entire history of the state of Israel, together with several of the major political and military events in that period.

In general, the pattern of positive mood/nice events was borne out during this period with a few nontrivial exceptions marked in grey. These exceptions underscore the fact noted earlier that social mood is not a 100% sure-fire indicator of what will or will not happen. It only sets the tone for what is *more likely* to occur than not. This graphic illustrates that maxim in spades. It also strongly suggests that if the Israeli social mood is indeed flattening out as it appears in this chart, we could be in for some "interesting" times in the Middle East over the next several years.

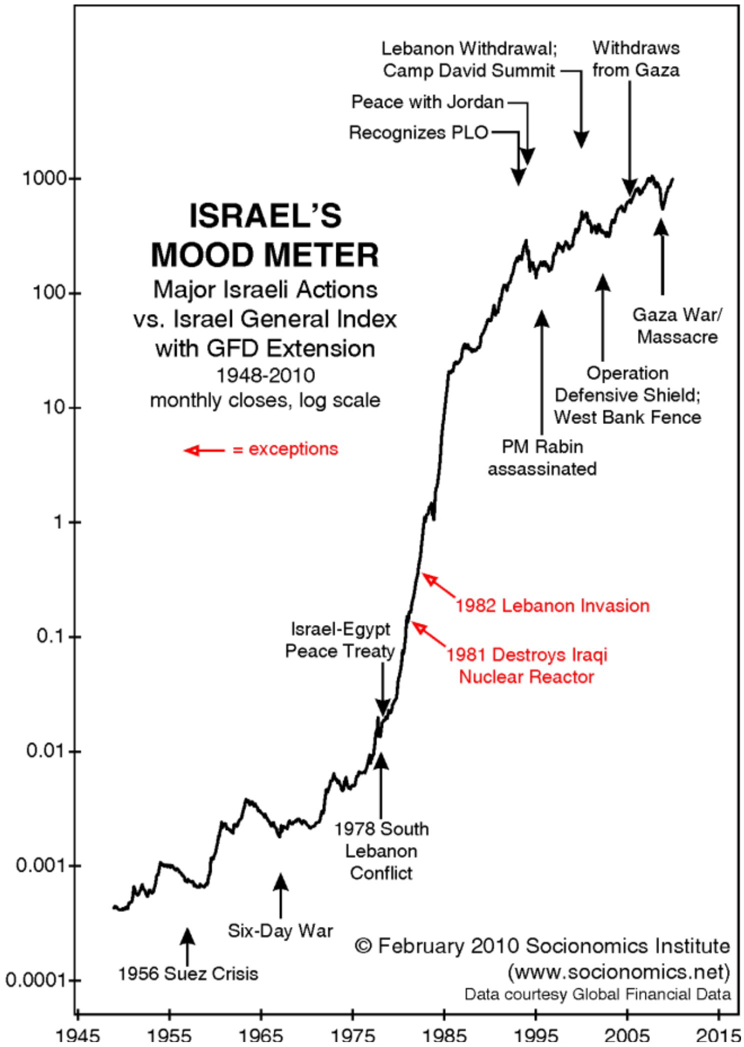

Figure 3. Israeli social mood and events, 1945-2010.

In terms of our landscape metaphor, the spectrum of possible events suggested by the social mood constitute the valleys in our diagram. The bigger the bias in favor of a particular possible outcome, the larger the valley. As a consequence, the random trigger pushing the system into one valley or another leads to a much greater chance of an X-event corresponding to a big valley than to a small one.

Now let’s turn attention to the other driver of landscape change, the complexity gaps.

Complexity Gaps

SYSTEMS IN COLLISION

In recent years, we've seen long-standing regimes in Tunisia, Libya and Egypt sent packing almost overnight, with Iraq, Yemen and Syria all now being torched by the very same revolutionary flames as rebels battle entrenched governments in an attempt to overturn decades of oppression. On the surface, these types of civil disturbances give the appearance of arising from public discontent with governments over high unemployment, rising food prices, lack of housing, and other such necessities of everyday life. But such explanations are facile, failing to address the root cause of the societal collapse. A civil disturbance is not a game-changer by itself but simply a precursor, or early-warning signal, for the coming X-event of regime change.

The real source of the X-event that moves regimes sits much deeper in the social system. It is a widening "complexity gap" between a government and its citizens, with revolution breaking out when that gap can no longer be bridged. Think of a rigid authoritarian government confronted with a populace that has awakened to new freedoms through their contact with the outside world and that is coordinated by diverse social-networking platforms. The gap between the complexity of the control system (here the government) and the increased complexity of the controlled system (the population) has to be bridged. One path is for the government to repress the population—imprison leaders, send soldiers to disperse crowds of protesters, and other measures to tamp down the situation thereby reducing its complexity. Alternatively, the government's complexity must be increased so as to speed up the holding of free elections, quickly remove restrictions on an open media, and open possibilities for upward mobility of the population.

This notion that an X-event is human nature's way of bridging a chasm in complexity between two (or more) systems is the leitmotiv running throughout the human-caused X-events we examine in this study. An X-event is the vehicle by which a disparate level of complexity between two (or more) systems in competition or even cooperation is narrowed. In particular, it is the default path when humans themselves fail to voluntarily narrow a widening gap. Let me just give a flavor of how this principle plays out by quickly recalling a couple of recent X-events where this complexity mismatch is very apparent.

Egypt had a state-controlled economy that was wildly mismanaged for decades. Even the noticeable improvement in the past few years is a case of too little, too late. Moreover, the country was (and still is) monumentally corrupt, as crony capitalism runs rampant throughout the entire social structure. Such a system of corruption relies upon bribes to officials to get contracts, obtain jobs or to find adequate housing. Amusingly (and tellingly) the impotence drug Viagra was reportedly kept off the market because its manufacturer, Pfizer, failed to pay a large enough bribe to the Egyptian Minister of Health for its approval.

This type of parasitic mismanagement and corruption worked to freeze in place an already low-complexity government, one that had very few degrees of freedom in either its structure or means of dealing with social problems as they arose. But as long as the Egyptian population had even more limited ways to express their dissatisfaction with a lack of proper housing, rising food prices, minimal health care and the like, the government had no motivation to create the framework(s) necessary to provide these services. Of course, there was a ministry charged with health care, for example. But it served mainly as a sinecure for career bureaucrats and cronies of those in power, and provided health care only as a kind of spare-time "optional extra." Who would expect this to ever change as long as the spectrum of actions available to the citizenry was kept at a low level (low complexity) much lower than that of the government itself? But times change. And when modern technologies like instant global communication, widespread higher education, and rapid transportation started making their way into the Arab world, citizens quickly became empowered. At this point, the handwriting was on the wall (more to the point, Facebook walls) for entrenched regimes throughout the region.

Modern communication and social networking services like Google, Twitter, and Facebook serve to dramatically increase social complexity—but now it's the complexity of the population, at-large that's enhanced, not that of the government. This fact is why such services are routinely restricted or even shut down when governments are under attack, as when the Egyptian government totally closed the Internet for a few days, since these services allow more voices to be heard and more highly-connected social networks to be formed. At some point the gap between the stagnant level of government complexity and the growing level of complexity of the general public becomes too great to be sustained. Result? Regime change in Tunisia, Libya and Egypt, along with the likely change of the Assad dynasty in Syria and/or the monarchy in Bahrain.

A complex system theorist would recognize immediately that the principle at work here is what's called the *Law of Requisite Complexity*. This "law" states that in order to fully regulate/control a system, the complexity of the controller has to be at least as great as the complexity of the system that's being controlled. To put it in even simpler terms, only complexity can destroy complexity. An obvious corollary is that if the gap is too large, you're going to have trouble. And in the world of politics, "trouble" is often spelled "r-e-v-o-l-u-t-i-o-n."

Examples of such mismatches abound. Take the Roman Empire, in which the ruling classes used political and military power to control the lower classes and to conquer neighbors in order to extract tax revenues. Ultimately, the entire resources of the society were being consumed simply to maintain an ever-growing, far-flung empire that had grown too complex to be sustained. The ancient Mayan civilization is another good case in point, as is the former Soviet Union. Some scholars, including historian Paul Kennedy, have argued that the American empire, which spends over $23 billion a year on foreign aid and consumes far more than it exports, is in the process of coming undone for much the same reason.

This type of mismatch is not confined just to complexity gaps in political and governmental domains either, as evidenced by the disruption of everyday life in Japan arising out of the radiation spewing forth from the Fukushima Daiichi reactors damaged by the March 2011 earthquake. The ultimate cause of this social discontent is a "design basis accident," in which the tsunami created by the earthquake overflowed the retaining walls designed to keep seawater out of the reactor. The overflow damaged backup electrical generators intended to supply emergency power for pumping water to cool the reactors nuclear fuel rods. This is a two-fold problem: First, the designers planned the height of the walls for a magnitude 8.3 quake, the largest that Japan had previously experienced, not considering that a quake might someday exceed that level. What's even worse, the generators were located on low ground where any overflow would short them out. And not only this. Some reports claimed that the quake itself actually lowered ground level by two feet, further exacerbating the problem. So everything ultimately hinged on the retaining walls doing their job—which they didn't! This is a case of too little complexity in the control system (the combination of the height of the wall and the generator location) being overwhelmed by too much complexity in the system to be controlled (the magnitude of the tsunami).

Right about now, a conventional risk analyst at an insurance company or a bank might be asking, What's new here? If we want to assess the risk of particular event Y taking place, we calculate the probability that Y happens, evaluate the damage done if Y does indeed occur, and multiply these two numbers together. That calculation tells us the expected damage if Y happens. No muss, no fuss. So what are you suggesting here that differs in any important way from this? For those readers who skipped over the last several pages of this introduction, let me summarize why this question from the conventional risk analyst process is the right question to ask—for "normal" events. But it's far less than the right question, a dangerous one even, to ask and assume the answer is Nothing when we start talking about extreme events. Here's why.

First of all, the very rarity of an X-event means that we do not have a database of past actions and behaviors sufficiently rich to actually calculate a *meaningful* probability for the event Y to actually take place. While probability theorists and statisticians have developed an ingenious array of tools ranging from subjective probability theory to Bayesian analysis to extreme-event statistics to try to circumvent this obstacle, the fact remains that nailing down a probability you can believe in for a rare event to occur is just not possible. If it were, we would not have experienced things like the Great Recession of 2007-08, the 2003 East Coast power failure, and the damage to New Orleans from Hurricane Katrina. And people would not be wondering when the next game-changing shock was going to come out of the closet and kick us in the you-know-where. So when it comes to X-events, we need to invent/discover ways to measure risk that captures what we mean when we say that this shock is more likely to take place here and now than it was previously. Here, our take on this problem is to advance the idea that the level of complexity mismatch between human infrastructure systems serves as just such a measure.

The second ingredient involved in a conventional risk analysis of normal happenings is the damage the event inflicts on society, should it take place. To assume the game changer actually happens at a particular time and place and then calculate the damage done doesn't differ from our thinking here, at all. The only problem is that if it's a shock that's never happened before to a system that underpins our human way of life, then it's problematic as to how to assess what the actual damage will be. To make such an evaluation, we generally need to compare the shock with comparable shocks from the past. We then look at how the assumed situation differs from what's in the historical record, and build a fudge factor into our damage assessment to account for this difference. But how could this process work if there is no historical record to draw upon? When the real world doesn't supply the necessary data we have to build a surrogate world in our computers to obtain that data. This, again, is an approach very different from what's employed in studying “normal” events.

In summary, there are two very different regimes here. There's the *normal regime,* consisting of events that have taken place many times in the past and for which there is a database available upon which to apply our tools for probabilities and possible damages. Then there is the *X-events regime* for which those tools simply do not apply. This study offers a perspective for creating a framework to systematically investigage the X-events regime, a framework that complements what's used to calculate expected risk in the normal regime. We present this line of argument both by precept and by example in the case of the European Union collapse in the Supplement to this paper, leaving the technical details for a research program that needs to be carried out in the months and years to come.

The X-Events Index: Where Do We Go From Here?

In order to create a full-scale X-Events Index, we must scale-up the work shown on the attached illustrative example of the European Union collapse in a variety of ways. These include

- Choosing a representative sample of potential X-events to focus the Index upon. Examples besides the European Union collapse might include political uncertainties in the USA, much in the spirit of what we've seen in the media about the potential X-events that might arise from the election of Donald Trump or Hillary Clinton in the upcoming November presidential election. Other interesting candidate X-events are a China-USA showdown in the South China Sea, a huge global financial meltdown, a worldwide epidemic, escalating tensions on the Korean peninsula, and/or Russian meddling in central Europe. Notice that none of these candidates are X-events from nature, such as a hurricane, earthquake, asteroid impact and the like. Our concern with the X-Events Index is with human-caused X-events only.
- Investigation of the subsystems in interaction acting to create the target X-event, as well as the "right" definition of complexity to use to measure the complexity gaps.
- Study of the timescale of the target event and how to measure the social mood in the relevant part of the world that acts to bias the occurrence of the event.
- Identification of suitable data sources for accessing the information needed for the these complexity and social mood indiators.