



A Physicist Who Models ISIS and the Alt-Right

The rise of new extremist groups has served as both an impetus and test-case for Neil Johnson's models of terrorism and insurgency.

By Natalie Wolchover



[Juan Cristóbal Cobo](#) for Quanta Magazine

Neil Johnson (center) in Bogota, Colombia, in July.

[Neil Johnson](#) used to study electrons as a buttoned-up professor of physics at the University of Oxford. Then, a decade ago, he decamped to the University of Miami — a young institution that he sees as unconstrained by rigid traditions or barriers between disciplines — and branched out. In recent years, the 55-year-old physicist has published research on financial markets, crowds, superconductivity, earthquake forecasting, light-matter interactions, bacterial photosynthesis, quantum information and computation, neuron firing patterns, heart attacks, tumor growth,

contagion and urban disasters, not to mention his extensive body of work on terrorism and other forms of insurgent conflict.

Johnson models the extreme events and behaviors that can arise in complex systems. The author of two books on complexity, he has found that the same principles often apply, regardless of whether a system consists of interacting electrons, humans or anything else. After the terrorist attacks of Sept. 11, 2001, he began modeling extremism in human society. He had also spent time in Colombia during the war against the FARC guerrilla army, and grew up near London during the era of IRA bombings. "I started wondering what the patterns of attacks in the respective places might be telling us about how humans do terrorism," he said. "Terrorism suddenly became, for me, an urgent problem that I might be able to help society understand, and perhaps even one day predict."

The rise of ISIS has served as both an impetus and test case for Johnson's models. Even more recently, he has begun using his models to study the growth of white nationalist groups in the United States.

Quanta Magazine caught up with Johnson to discuss his findings by phone in June, before he left to spend the summer working with collaborators in Bogota, Colombia. A condensed and edited version of that conversation, and a subsequent email exchange after the events in Charlottesville, follows.

What's a physicist doing studying terrorist networks, financial markets and all these other systems?

In all these complex systems, the pieces of the system interact with each other and they evolve over time. And there's something that a collection of objects like that can do which a handful of coins cannot do. I can throw up a set of coins and it would always come down pretty much 50-50 heads and tails, and there will be a little bit of variance around that — it obeys something called a bell curve. We base so much science on the bell curve. Bell-curve distributions arise when you deal with coins, or any collection where the pieces aren't connected, like heights of people in a room. However, in most of the systems we're interested in — the hard problems, be they of science or society — those distributions look very different than bell curves; they're so-called fat-tail distributions.

Thinking about heights, instead of everybody being 5 feet 10 inches, on average, and maybe down to 4 feet and up to 7 feet, but certainly not 70 feet, with the distributions you get in these complex systems, you can get the 70-foot person. In fact, you can get the 700-foot person. There's something about the way the pieces interact with each other that makes these extreme events happen: the 700-foot person, the stock market crash, the 9/11. So the interesting question is, is there a general science that can govern and tell us about these extreme behaviors? And if we can understand that for one system, can we transfer that understanding over to another one and therefore do something about it?

When you began working on terrorism in the early 2000s, where did you start?

We looked at the shapes of the distributions of terrorist attacks. Given 9/11 and an attack half that size — how frequent are the two relative to each other? That gives you the statistical distribution, like a distribution of heights. In doing that, you find common features across all these different conflicts and across terrorist events, regardless of their specific details. Now, you talk to a social scientist and they think that's absolutely awful to hear that. Because they're focused on those details. You can have someone who's an expert on the Second World War, or the Vietnam War or

Iraq, and it's kind of strange for them to get a message from another discipline that somehow those details don't matter.

What unites the different conflicts?

When you plot the frequency of events versus the size of events, you get this straight line and it has a slope of 2.5 for conflicts and terrorism. And that was a surprise to us. So here starts the physics. We started to build mathematical models of what might be going on in a system of insurgents or terrorist groups; what would create a 2.5 power law?

Were you able to figure out where the 2.5 comes from?

We were looking at Iraq and Colombia at the time — 2003. And we said, OK, one's in the desert, one is in the jungle; let's think about what people do. And we thought, well, the only thing we can really say is both of those conflicts are sort of irregular. There might be a state army, but they're fighting against an insurgency — terrorists or guerrillas, where you've got loose groups that come together and then if they sense danger from the opposing military unit, they might fight and then break up and scatter in all directions. Almost like fish under the sea. They build into schools of fish, and when a predator comes along they scatter; and then they kind of re-form again because there's an advantage to re-forming. In guerrilla warfare there's an advantage because you get to attack in a group. And then they scatter.

We took those two features and we built a mathematical model, writing down differential equations for collections of objects like coins, but now, instead of just sitting on their own, they try to come together over time — this is called coalescence — and then when they detect some kind of imminent danger, they break apart. And then they form up again, then break apart. We had to change the usual physics and chemistry equations and put in this feature of gradually coming together and then completely fragmenting. And when we sat down and solved the equations, it comes out to be a power-law distribution for the sizes of the clusters with a slope of 2.5 exactly. And then, curiously, when you go around and try and vary some of these rules, such as how quickly they fragment, it doesn't seem to matter too much how you vary them. So suddenly we had this model. If we assume that's the mechanism for forming insurgent groups, and if they have some probability of acting, the size of the events I observe should look the same as the distribution of the size of groups. If you make that leap you suddenly have explained the 2.5 power law. We [put that out in 2005](#).



[Juan Cristóbal Cobo](#) for Quanta Magazine

Then a few years later, along comes ISIS and the pro-ISIS support online.

We had this idea in 2014 when ISIS first appeared on the horizon to start tracking what we saw. Implementing it was hard, because Facebook was really good at shutting down this type of activity. But we found there's one social media entity that's really popular in central Europe and has 350 million users worldwide, called VKontakte. It's exactly like Facebook, and it has this interesting feature Facebook has, which is groups. Pro-ISIS supporters would form themselves into these groups and they would exchange information about weaponry, financing, recruitment and events.

I had Ph.D. students and post-docs who were Russian speakers, Arabic speakers, and from political science — and we tracked pro-ISIS groups over time. We found that exactly like the fish under the sea, ISIS supporters slowly build up into groups and then the groups get shut down by the moderators, in which case they scatter. And people don't disappear; they just go off and form other groups. So not only did we find exactly the mechanism that we proposed in our model, but when we looked at the size distribution of the online pro-ISIS groups, we found it was a power law and its exponent came out to be 2.5. That was [a Science paper](#) a year ago.

So once you have a model that explains where the power law comes from, how does that help? What does it actually tell you about combating terrorism?

Most of the approaches to dismantling the online support — recruiting and financing, et cetera —

are at the individual level. They always seem to want to find the bad guy, the needle in the haystack, the ringleader. What our work shows is that is not the way to go. It's like the fish: Imagine I want to stop schools of fish forming. You try to catch one fish; will it stop the grouping? No, of course it won't. Fish No. 3 becomes No. 2, No. 2 becomes No. 1, and in fact there may not even be any hierarchy, there's just a collection of objects. So you need this systems-level approach or you'll never understand this behavior.

It's a classic conundrum: Nobody knows how to look "[left of boom](#)." Security agents are very good at finding who's actually buying explosives, who's just about to do something. But what about when the people themselves don't necessarily know where they're heading? If you can understand how people move through these groups, then you're going to get a sense of who is developing momentum toward at least having the intent and the capability. It certainly seems that this dynamic systems view is better than just watchlists based on immigration status.

In [one recent paper](#), you analyzed individuals and groups on VKontakte that were banned for promoting violence; what did the research suggest?

It turns out that most people that get banned, it's because they've been members of certain types of groups. But it's not true that the more groups I join, the more likely I am to become banned. We were able to find out that people who are most likely to get banned are those who join one pro-ISIS group. Join two, and your probability of becoming banned is less. So might it be that by joining two, I sort of confuse my message to myself? Then the chance of being banned after joining three groups is less than for two groups, et cetera.

We also find that the people on the way to becoming banned tend to go for the small groups, the ones that are focused on, not the news, but something more to do with the spiritual or ideological side. It doesn't seem to be the case that people go along, and then there's a piece of news that bothers them and then they go out and do something; it really is this ratcheting up in ideology. And the people who develop more quickly do that in a more predictable way. The ones who take longer to cook, as it were, they tend to fluctuate around more. Which is interesting, because that means there's probably more opportunity to persuade them away, for instance by trying to get into one of the groups that seem to be where they are heading and soften the message and deflect the person away. Now, that's not my business; I do the science. But there are interesting possibilities that we hope might be looked into.

Can you check that your model actually identifies the terrorists, rather than just ISIS sympathizers?

There's a whole bunch of people who are members of these online groups who don't end up doing anything. But there are many whom we identified who are known from media reports to have eventually been killed in combat. It's an awful thing to be talking about, but I think it's an important thing to be doing. Because all of this is open source information. We could sit down in a Starbucks, open up group pages on VKontakte, we'd see everything, because these groups keep themselves open to try and attract recruits and new people.

Do the intelligence agencies take note of your findings?

We've given a lot of talks and I'm very impressed by how much interest U.S. agencies showed in this work. The unfortunate thing is, it's basic science that we're still trying to work out at the same time that we're addressing the problem. So we don't have daily interaction with those agencies. They may be doing something in private; I've seen our work mentioned in a lot of reports that are in the public

domain.

Does your research also apply to the rise of white supremacist groups in the U.S.?

Yes, what we are doing is very relevant since the alt-right groups live, recruit and coordinate (and hence evolve) online. And from what we can already see, they do so pretty much exactly like the pro-ISIS groups evolve and coordinate, but Facebook has so far been less quick to shut them down. So the question is: What was the activity of the online groups before [Charlottesville](#)? And if we look at their evolution (as we did for pro-ISIS groups) from now on, can we foresee the growth to an outburst like a future Charlottesville, but elsewhere in the U.S.?

I am writing a paper showing results that we already have, for how two opposing sets of online groups (for example, alt-right and their opponents, including Antifa) interact online and under what conditions one online group will be able to overcome the other.

In general, this tendency for people to flock together and then break up is, I think, common to any human activity which is under stress. In the future it won't be ISIS, it will be something post-ISIS. I would claim that none of those details actually matter too much. It's the same idea.

This article was reprinted on [TheAtlantic.com](#).