Host: Robert Frederick
It's called computational propaganda. It's primarily spread by Twitter bots and fake Facebook accounts. It's how fake news — designed to sway public opinion, sway your vote — it's how all that gets around.

Speaker: Philip Howard
At the moment, the democracies that seem to be suffering the most is Brazil, possibly Germany, and the United States.

Host: Robert Frederick
On this episode of The Conjectural, manufacturing consensus, online. I'm Robert Frederick.

Host: Robert Frederick
If you're on social media very much, you've probably seen evidence yourself, read output from, or even interacted with a highly automated social media account—a bot. Philip Howard is the principal investigator of The Computational Propaganda Project at Oxford University.

Speaker: Philip Howard
Our goal is to produce large amounts of evidence gathered systematically so that we can make some safe, if not conservative, generalizations about where public life is going.

Host: Robert Frederick
Because public life is going more and more online. Howard was speaking in a lecture hall at the European Conference of Science Journalists held in Copenhagen in late June, which I attended.

Speaker: Philip Howard
At the moment, the democracies that seem to be suffering the most is Brazil, possibly Germany, and the United States.

Host: Robert Frederick
And in the United States, we have a history of using propaganda — a term that now has chiefly negative connotations, but was studied quite extensively with the idea for pubic good: that sometimes the public needed to be guided when important decisions needed to be made quickly. One of the early pioneers who applied science to this public guidance was Edward Bernays, nephew to Sigmund Freud, the famed psychoanalyst. In 1947, Bernays published an essay well-known to communications scholars, titled "The Engineering of Consent." He wrote:

"With pressing crises and decisions to be faced, a leader frequently cannot wait for the people to arrive at even general understanding. In certain cases, democratic leaders must play their part in leading the public through the engineering of consent to socially constructive goals and values."

Of course, the same tools for the engineering of consent "to socially constructive goals and values" can also be used for the engineering of consent to socially destructive goals and values, or, say, marketing. And Bernays — who may be quite correctly referred to as the father of modern public relations — is behind why we think disposable Dixie cups are more sanitary than reusing a washed glass, why American society was so quick to accept the addition of fluoride to the public water supply, and why bacon and eggs is often considered the all-American breakfast: marketing, marketing, marketing.

Fast forward to 1988, when Noam Chomsky and Edward Herman published "Manufacturing Consent: The Political Economy of the Mass Media," which begins with a description that sounds like an indictment that—quote—

"The mass media serve as a system for communicating messages and symbols to the general populace. It is their function to amuse, entertain, and inform, and to inculcate individuals with the values, beliefs, and codes of behavior that will integrate them into the institutional structures of the larger society. In a world of concentrated wealth and major conflicts of class interest, to fulfill this role requires systematic propaganda."

Bernays, Chomsky, and Herman, though, of course, were constructing or examining messages to convince real people. Today with automated bots, Philip Howard says it's truly manufactured consensus — not even the people are real.

Speaker: Philip Howard
So, in the first Presidential debate both Hillary Clinton and Donald Trump had automated accounts tweeting — oddly — about about how successful they were. But over time, between the first debate, the second and the third, Trump's bots started announcing that he had won the debate earlier and earlier in the day, such that by the third debate, Trump's bots over Twitter were announcing that he had won the debate before the debate was broadcast.

Host: Robert Frederick
Would that be convincing to people, to voters? One wonders. But Howard says that convincing people of something isn't necessarily the goal of such campaigns to sway public opinion.

Speaker: Philip Howard
The earliest of campaigns originated in Russia. The best of them that we've caught involved the spin around the Malaysian flight that was shot down over Ukraine in the summer of 2014…. The goal for this kind of communications campaign and the automation behind it, isn't to put out one counter narrative, or put out an alternative perspective, but to seed multiple conflicting stories that different proportions of the public will all believe in equal, small measure. The effect is that now there's at least four of these stories explaining why the Malaysian Airlines flight was shot down. And having multiple conspiracies in play helps, at least from the Russian perspective, helps to prevent any political reaction.

Host: Robert Frederick
And as you might imagine, in the days before the 2016 U.S. election, the bots were at it, spreading more and more junk.

Speaker: Philip Howard
And we found that in Michigan, there was about a one-to-one ratio: for every piece of professionally produced news content there was one piece of junk.

Host: Robert Frederick
Some of the data for the Computational Propaganda Project came from Twitter itself and from analyzing what Howard and his colleagues could from public groups on Facebook.

Speaker: Philip Howard
The big data analysis involves big scoops of mostly Twitter data. Facebook doesn't collaborate well with researchers — it doesn't have an API that we can actually make much use of — so there's a lot of phenomena that we think that we can identify on Twitter that we hope speaks to what goes on on Facebook. But without being able to check these things, without collaborations with Facebook, it's difficult to know.

Host: Robert Frederick
Other data for the Computational Propaganda Project was from interviews with a few of the people who program bots. How do you get such an interview?

Speaker: Philip Howard
As with any kind of ethnography of marginal or deviant communities, the only secret sauce to the method is time, and finding one respondent who gives a good interview and then agrees to be a fixer

Host: Robert Frederick
Setting up other interviews. And it was through these interviews Howard and his team learned such things as who these people were, where they lived, and why they were programming bots — these highly automated social media accounts — to sway public opinion.

Speaker: Philip Howard
So, some of are subjects are from Seattle, and from San Francisco and Brooklyn, and they're from Montana — they're libertarian programmers who are expressing their citizenship. This is political speech for them. They're expressing themselves. And that's another group that does this work, and their clients are not the U.S. government. They're political parties and lobbyists.

Host: Robert Frederick
But, to be clear, Howard says, the large campaigns of bots, the ones working to affect our 2016 election that he and his team were able to discover, they all originated in Russia.

Speaker: Philip Howard
We did not find them from China. Everything we found so far seems to originate in Russia. Over the last year, there are also these predictable crises in democracies. In fact, this is to some degree what makes democracies soft targets. Elections are sensitive moments for political structures and in our democracies. And so we've done a series of memos over the last year about automation and junk news. We started with the Brexit referendum. We did the three U.S. presidential debates and the election itself. We did the German presidential mostly as an exercise to do some benchmarking for the next presidential voting in October. We did the two stages of the French election and the two stages of the UK election.

Host: Robert Frederick
In response, Howard says that civil society groups in every country are struggling in the face of these automated political attacks via social media. And they're struggling because they don't know how to mount an effective response without also resorting to bots. And so sometimes these civil society groups, and/or individuals who represent a particular view or stance, sometimes they just end up leaving social media as a result of these attacks.

Speaker: Philip Howard
Many of the campaigns that we've studied are particularly good at driving women off of social media. So prominent feminists, prominent feminist intellectuals, female reporters, female politicians are also soft targets. And there are multiple examples of prominent female intellectuals being driven off social media with campaigns against their public life.

Host: Robert Frederick
After Howard's presentation and the question-and-answer period was over, I went up to talk with him further.

Interviewee: Philip Howard
Hi.

Interviewer: Robert Frederick
Hi, you mentioned there was various techniques for discouraging women from remaining on their accounts. I wonder if you could talk a little bit about what those are.

Interviewee: Philip Howard
Well it's mostly sexual harassment over Twitter, and it includes dirty words and nasty commentary. And there have been now multiple — in the U.S. — pop stars, sort of, local political figures who have just been driven off Twitter because of the sexual harassment online.

Interviewer: Robert Frederick
And so this is a highly automated account doing this kind of work, or…

Interviewee: Philip Howard
Sometimes it appears to be highly automated accounts from the alt-right in the U.S.. So, do you remember the GamerGate story that involved… so there was some automation that was involved in those attacks.

Interviewer: Robert Frederick
OK. Is there something that people can do, is there something — if someone starts to receive or is a subject of these attacks — that they can say, 'Oh, well, I've just been targeted.'

Interviewee: Philip Howard
The platforms are getting better and better at responding to user reports of hate speech or sexual harassment. So, I mean, that's the default thing: report the other user. If Twitter doesn't act fast enough or the network of bot attacks seems to be really expansive, then people pull off social media. And sometimes there's a goodbye tweet where people say 'I'm leaving 'cause I'm done — this is too much' and they pull off, and that tends to have an impact, too.

Interviewer: Robert Frederick
Does that stop the bot from continuing the work, or…

Interviewee: Philip Howard
Many of the bots that we've… many of the automated accounts that we've studied shift from issue to issue or from focus to focus. Some of them go quiet after an election, but there was a handful of bots from the Brexit conversation that switched to being interested in U.S. politics, right, and then

switched to being interested in the Italian referendum, and then they switched to being interested in the German election, and then the two French… so there's a little bit of migration, and, I mean, that suggests that there's direction behind it all. But most of the code that goes into creating a bot gets used and repurposed in different ways.

Interviewer: Robert Frederick
Do you have a current list of bots that you, or those highly automated accounts that you think are out there, but haven't…

Interviewee: Philip Howard
If you, yes, if you have a particular country that you're interested in, then we make lists of the top-100 highly automated accounts. There's occasionally false positives in them, but yes, in principle — I mean, I don't have one right now on my laptop but I would have some old ones from the U.S. election. But if you send me a note I'm happy to — and it's a dataset that we already have —I'd be happy to carve out some material for you.

Interviewer: Robert Frederick
OK. But you don't publish that as part of your bot prevention. It's more about studying it as it happens.

Interviewee: Philip Howard
This summer…. No, we're going to be putting up our replication data this summer, and so for each of the reports, we'll…. Twitter has a policy where you can't share the data. If you share data, Twitter won't let you study it anymore. So we will be providing the Tweet IDs so you can see which accounts are there, but we can't provide any of the content. So I guess I'm saying, they have a policy that kind of discourages us from collaborating too much.

Interviewer: Robert Frederick
Does that seem like a good idea to you?

Interviewee: Philip Howard
Nope, I agree. Yeah, I agree (unintelligible). You know, we can identify hundreds of accounts, and then we can publicize the Tweet IDs. Somebody else could, sort of, reverse engineer and look up those accounts fairly easily. But your questions is great: Twitter has never asked us to; we've offered to collaborate on this, but they don't, they've never asked for our….

Unidentified Speaker
That's the news.

Interviewer: Robert Frederick
Yeah…thank you.

Interviewee: Philip Howard
Of course. Thank you.

Host: Robert Frederick
Indeed, as you heard my colleague in the press said, that's the news—the lack of apparent interest by these social-media companies in collaborating with researchers to shut down these highly automated accounts that attack people. Howard and his team are expected to release their data

and next set of reports this summer, which you can find it on their website, The Computational Propaganda Project. You can find a link to it on our website, TheConjectural.com.

Finally today, I have an announcement. We're taking a break from the show because I'm running for political office,with an election this fall, 2017. We may or may not return because if I win, I'll be mayor of my small town on top of my full-time job, and the original purpose of this podcast — to experiment to find better ways to talk about science news — well, the world's changed significantly in the two years since I started this monthly show, and, yes of course, that includes my world, too. For example, in July, 2015, when I started this show, I was a freelance science journalist. Now I'm digital managing editor of *American Scientist* magazine, producing their monthly podcast. In July, 2015, there was science denial-ism, but it was limited mostly to the fringes, including a brand new, extremely unlikely — or so we thought at the time — candidate to be President of these United States. Now there's science denial-ism in the mainstream, including in the White House. The experiment to find better ways to talk about science news, well, I do that in my job. There are different problems with higher priority now, at least for me, in what I do when I'm not working my job.

But I've learned a lot from this experiment in science news, and although the results of this experiment would be difficult to summarize in only a few words, of course — for you — I'll give it a try.

There is a power in a story told well, and the most popular podcast episodes — the ones downloaded the most, received the most emails about — they were ones in which I used traditional story forms with characters, I explored tensions, stories that had a beginning, middle, and end. This result is not terribly surprising, but it is fundamentally different than science itself, which doesn't have an end and doesn't depend on characters but often the scientific enterprise succeeeds despite characters. That's because the power of scientific thinking is not in reaching a conclusion — thinking deductively like Sherlock Holmes does, mostly — but in disproving guesses. And we do not live our lives by disproving guesses. I don't decide that walking down a dark alley is safe just because I've experimented by walking down dark alleys.

And speaking of dark alleys, or things in the past that I've avoided, for the most part, before, I'm entering politics because I enjoy identifying problems in communities and working creatively to solve them. It's been a characteristic of every career I've had, from management consulting to textbook editing, teaching to science journalism. This show's experiment had identified a problem — a problem with how science news is delivered as something new, something wrong, something new, something wrong, so that the public learns through repetition that what is new in science today may be wrong in science tomorrow. So why even bother to learn what's new? And through this show I worked, with your input, to explore new forms of talking about science.

So the takeaway: in doing a good story about science news, it's not enough to describe a problem — a problem carefully stated is only half solved, as the saying goes — it's not enough to describe the characters who are working to solve it, not enough to describe the tension surrounding the problem, not enough to describe the reason the problem matters, not enough to give the story a beginning, middle, and end — a conclusion. That's because, as I've learned from doing this show through the second-most popular podcast episodes, which all shared the characteristic of this: my taking extreme care to figure out a way to give the story an ending but not the science a conclusion. I struggled to find ways to allow the problem — the scientific problem — to continue,

to allow that more data would be gathered, to allow for more experiments that would be done, to allow that there might be some scientists who disagree, to allow that all knowledge that we create is based on incomplete information. And that is okay, because all knowledge, even scientific knowledge, is in some sense, Conjectural.

You've been listening to The Conjectural, a show that's been running an experiment. The data for this experiment? Your feedback to TheConjectural.com, where you can still give feedback and download transcripts, at the very least until this November's election. Support for what may be this final episode comes from the European Science Journalism Conference and from *American Scientist* magazine. I'm Robert Frederick. Thanks for joining us!