

## Video message by Dr Paul Bernal

To the 3<sup>rd</sup> Myanmar Digital Rights Forum, 18 January 2019, Yangon

Hello, I'm Doctor **Dr Paul Bernal**. I'm a Senior Lecturer in IT, Intellectual Property (IP) and Media Law at the University of East Anglia Law School. I'm a specialist in Digital Rights, in particular, in Human Rights as they apply on the Internet. For that reason, as well as being someone who has a long-standing interest in Myanmar, I'm particularly happy to be invited to give a talk for the Myanmar Digital Rights Forum.

My particular area used to be privacy. Privacy is where I started from, and privacy still remains an important part of what I do. But as my research has developed, I have found that it has gone in various different directions, all of which are relevant to what we are talking about here.

My new book, "**The Internet, Warts and All**", has a subtitle, "Free Speech, Privacy and Truth". I think it's really important to understand the main contention in this book, which is that those three issues, free speech, privacy and truth, are intrinsically and inextricably linked. One of the biggest problems that we have at the moment is that we try to deal with them one by one, rather than understanding how they work together and that works in lots of different directions.

Most importantly, when you do something to do with **privacy**, it has an impact on **freedom of speech**. Most directly, even though people imagine that privacy is somehow in conflict with freedom of speech, privacy in practice underpins freedom of speech. Unless you have privacy rights, you really can't protect freedom of speech. And, similarly, the measures that we take to deal with things like freedom of speech often involve invasions of privacy.

So we want to deal with extremist websites, what do we do? We have to find out who is visiting those extremist websites. We do that through invasions of **privacy**. All of this makes it harder for us to find out the truth, in so far such a thing as '**truth**' exists. This is played out in all kinds of different ways, in many of which are relevant to the things that you'll be discussing today. One of these is **fake news**, and another of these is **trolling**. I'll be talking both about fake news and trolling a little later on.

I think, though, the most important thing to understand - and this is really the central point to my book - is that a lot of the things that we do to try to **regulate the internet**, to try to put laws in place for the internet, and to try to deal with the big companies that so much dominate the internet, are based on **myths**. To a great extent that's what I deal with in my book, to address these myths. Some of these myths are kind of universal myths, like the myth that everything, once it's put on the internet, is there forever — I call it the **myth of permanence**.

The **myth of perfection**: this is that somehow things that are on the internet are all right, and that we should be very wary of interfering with the stuff that's already on there. Very often, for example, the **right to be forgotten** - which is actually a really important right - is treated as '**rewriting history**'. But actually 'history' sometimes needs to be rewritten, where 'history' has been done falsely, or badly, in ways that are politically biased and so forth. Actually what we may need to do is go back and look at 'history' and rewrite it. So we need to understand that.

So we have the myth of permanence, we have the myth of perfection, we also have - and this I think perhaps the most important in relation to Myanmar at the moment - **the myth of neutrality**. The idea of neutrality – as a kind of fundamental concept - is something that has pervaded the internet from the very beginning. It pervades it even now. We somehow imagine that the internet search engines work in a neutral way. We imagine that the internet social media services - Facebook in particular - are in some way neutral. We imagine that the information that's on the internet is presented in a neutral way.

Neutrality - the neutral point of view - is one of the **five pillars of Wikipedia**, for example. We imagine that somehow, there are humans who are able to be neutral. We imagine perhaps even more misleadingly, that **algorithms** are neutral. That somehow, something that is organically generated through an algorithm, when we

use algorithms to curate news, when we use algorithms to decide who is recommended as a Friend for you, or something to Follow, that those are somehow neutral because they're automated.

None of these things is true. **Algorithms** are not neutral. They are not neutral for a number of reasons. They are not neutral because the people who created them aren't neutral. They are not neutral because the data used to generate them, things like machine learning and artificial intelligence, basically use large quantities of data in order to generate further quantities of data, and to work out correlations, and from that make decisions. If the original data is skewed, the original data-order is skewed, then the algorithm based on it will also be skewed.

We also imagine - and this is the last and perhaps the most important part of the neutrality myth - that the corporations involved in them are neutral: that Facebook, as a corporation, is neutral. It likes to present itself as a **platform**. Even the word 'platform' implies neutrality. But it is not neutral. Facebook is a business; it serves itself, as a business. It's there to make money. It will do what it needs to make money, both short term and long term, regardless of whether that's neutral.

And in some countries, that means that Facebook will cooperate with the authorities in particular ways, because that's the only way it can get the relevant permissions to function in the country, and that's the only way that it will be supported in the country. In other places, it resists the government because there are big corporate interests that are more interesting to it, or because the people involved are more interesting to Facebook than the government is. Either way, it's not neutral. So we have a corporation that's not neutral, using algorithms that are not neutral, and using individuals that are not neutral where those algorithms are not functioning.

So if we understand that these things aren't true; that the neutrality myth is important; the perfection myth is important; the idea of things being permanent is also important, then that should change our perspective of how we consider the governance of the internet, and how we consider how we behave on the internet.

In relation to Myanmar, there's a particular thing - and this is one of the reasons I've been talking about **Facebook** in particular— which is that Myanmar is a place where Facebook has deep importance. It plays a really, really important part. I know that in a session during your forum, you'll be discussing **content management** on Facebook. Content management is really important, it's really important in lots of different directions. It's important because we have problems with extremist content, we have problems with 'fake news', we have problems in the way that content gives you the way effectively to bully and to harass and so forth, and to create narratives that are then, potentially, destructive.

But the biggest problem we have with Facebook is its universality. Facebook has such a strong and important position that it is in a position to control a great deal of the agendas of people and a great deal of what happens to people. The important thing to understand here, is not just that Facebook is not neutral - which is important - but that Facebook has significant power to influence and to manipulate.

We see this all over the world. Myanmar is just one of the many places where Facebook has significant influence. But I think it's important to understand that this is not just anecdotal stuff. This has all been empirically evidenced.

Facebook itself has done a number of experiments which have demonstrated how effective it is. They have demonstrated, they have proven in the US, they are able to get people to register to vote to a statistically significant effect. They've been able to demonstrate, further to that, that they've been able to get people to actually go out and vote.

They've also been able to demonstrate that, through the curation of their feeds, they're able to manipulate people's emotions. And, this is perhaps one of the areas that people don't really quite understand. They did a massive experiment in 2014, using approximately 700,000 people that they broke up into three separate groups.

With one group, they used their curation of their newsfeeds to suppress sad stories and to enhance happy stories. They did this by using automatic analysis - what they call **sentiment analysis** - on the content of the news story. The second group, they did the opposite. They raised up the sad stories and suppressed the happy stories. The third group, they suppressed both the happy and the sad stories, and enhanced the emotionally neutral stories.

Now, this got a lot of headlines at the time, but again, I think this has been a little bit forgotten. The first group, the happy group, got happier. So they were able to make people happier by making the happy stories appear higher and sad stories appear lower. They measured this happiness by doing the same kind of sentiment analysis on the posts done by those people. The second group, the one where they enhanced the sad stories and suppressed the happy stories, they got sadder. So they were able to make one group sadder and one group happier. The third group, which is the one that wasn't reported enough - but is in some ways, perhaps the most important - where they suppressed both happy and sad stories, and enhanced the emotionally neutral stories, they spent less time on Facebook. They engaged less with Facebook.

So what does this mean? This means that Facebook has an incentive to support, and to promote emotionally charged content, and to suppress neutral content. It's not to say that they necessarily do this. But their algorithms will be designed to create the greatest level of engagement with Facebook. That will, in natural consequence, mean potentially, the suppression of emotionally neutral stories.

When we start to look at fake news and extremist content, you'll see that this provides an incentive for Facebook, in practice, to promote fake news and to promote extremist content. That is in tension with their stated suggestion that they are going to do the opposite: they are going to deal with extremism and deal with fake news. In practice, they have a conflict here, that is very difficult to know about, particularly because what Facebook actually does is distinctly far from transparent. That's one of the things that I think needs to be changed. This is not a problem just for Myanmar. This is a problem for the whole world. I think the whole world is starting to get to grips with some of the problems with Facebook.

A bigger problem with Facebook is the **fake news** problem. I think we need to understand lots of things about fake news. I've done quite a lot of research on this: as well as my book, I have some separate publications on specifically Facebook role in fake news.

The first thing to understand with fake news is that it is not new. There is evidence of things very similar to fake news throughout history. In my articles, I have talked about, for example, 15<sup>th</sup> century Romania, where Vlad the Impaler<sup>1</sup> was a victim of fake news. Effectively, his reputation was tarnished forever by the use of fake news. There are similar versions of fake news all the way throughout recorded history. Some people may have heard a recent example of a German reporter for a very distinguished German newspaper being called out for faking news in America. What's very interesting is this precisely echoes a 19<sup>th</sup> century German phenomenon, where fake foreign correspondents were effectively used to create stories from abroad because the newspapers couldn't afford, or didn't want to send, real reporters overseas. So they just made up stories based on reports that they had already got.

Fake news has been used for a very long time. It's been used either to make money, or to politically manipulate, and or to do both, which is perhaps more common than either. The thing about fake news is that although it existed pretty much forever, in the new era, the social media era, Facebook era, it's particularly bad. There are lots of reasons for this.

The first thing to understand about fake news is that it is more likely to be believed than real news or at least, it can be. It can be more likely to be believed than the real news because it is generally created to fit in with the prejudices of the people that they are trying to get to read it. So people are naturally predisposed to believe these stories. If, for example, they have anger with a particular group of immigrants, then, a piece of fake news making those immigrants out to be bad will be more likely to be believed.

Fake news can also be created to be more emotionally charged, and hence, more likely to be shared and believed, and to have fewer plot-holes. The real world is a bit messy. Real stories are little less convincing. When you create a story, you can fill in all the plot-holes. You can make it all sit together very nicely. That makes it even more likely to be believed. This is a phenomenon spreading all over the world right now. One of the reasons for this is the way that social media works. This is being played out particularly in things like the Cambridge Analytica saga as just one example, which uses the way that Facebook works specifically to create an effect.

---

<sup>1</sup> Otherwise known as Dracula

So what happens? What happens is you use big data analysis on the user data on Facebook in order to identify what the things are that people are likely to believe, what are their prejudices, and so forth. That allows you to then create fake news that fits perfectly with those prejudices. That makes it more likely to be believed. You can then use Facebook's ability to target. So having created the fake news, you can make sure that it gets to the right people. What's more, you can create it in a way that is more likely to be shared. Again, there's empirical evidence about what content is likely to be shared: more threatening content is more likely to be shared, more emotionally charged content is likely to be shared and so forth. And you can use Facebook's social maps, the links between Friends, Friends of Friends, and so on. By Friends, I mean Facebook Friends rather than real friends. Therefore, you can work out where to place your news so that people will automatically share it for you. You just need to create it, stick it somewhere, and it will then, automatically spread.

Now most of the attempts we have to deal with fake news, deal only with the actual news itself, rather than the fundamental environment underneath it. And because fake news is very, very easy to create, and you can put up a Facebook Page that looks identical to a proper news page in a matter of minutes, dealing with individual examples is fundamentally frustrating. It's a game of Whack-a-Mole. It just does a few things and then fails. You get more coming out all the time. You can get rid of one, the next one pops up, and so on.

The only real way to deal with it in the future is going to be to work on two levels. These are both challenges, particularly for people in places like Myanmar. The first is to make sure that the real reliable news exists. So the bedrock for dealing with fake news is to have real news that can actually counter it. For that, proper journalists have to deal with this, and to try to avoid falling into the trap of being creators of fake news, following government lines too easily, and so on.

The second - and this again is something that comes back to the basic theory in my book - is to work better for privacy. If we protected privacy better, if we did not allow the likes of Facebook to use the kind of big data analysis that they do, if we stopped that data from being gathered in the first place, it will be much harder to create fake news. If we didn't allow them to **micro-target** in the way that they do, it would be much easier to stop the spread of fake news.

Can we do that? Can we address Facebook's entire business model? Very hard. But we need to though. Because unless we do, it's inevitable that these problems will reoccur. And what we also need to do -and I think that this is where the biggest challenge is for places like Myanmar - is to try to reduce our overall dependency on Facebook. Other social media as well, but Facebook is by far the dominant one. Part of our problem is that everybody is on Facebook, everybody uses Facebook, all the news services put their news out on Facebook, and so on. While we do that, Facebook has the power. While it has the power, it's going to be in a position where we will be manipulated.

This is not just a problem for Myanmar. But Myanmar is perhaps one of the most vulnerable in this way. It's a problem for the whole world. We need to try to address it. We need to find a way to be more 'savvy' about the media and about social media, and be more able to address all these problems. We need to be less on Facebook, we need to be on Facebook for less time. We need to find alternatives, particularly for distributing things like news. If we can do that, we have a chance to deal with it.

This is not a challenge just for Myanmar. This is a challenge for the world, although Myanmar is in some ways more vulnerable than others. And all I can say at this point is that 'we're all in it together'. We're all trying to help, but we need to work together and we need to understand that we need to make some significant changes.

I'm sorry for not being able to provide more answers, rather more questions.

I wish you luck with the rest of the Forum. Thank you very much.

[paul.bernal@uea.ac.uk](mailto:paul.bernal@uea.ac.uk) Twitter @paulbernaluk (...and no longer on Facebook)