



HOUSE OF COMMONS

Digital, Culture, Media and Sport Committee

Oral evidence: Fake News, HC 363

Tuesday 22 May 2018

Ordered by the House of Commons to be published on 22 May 2018.

Watch the meeting

Members present: Damian Collins (Chair); Julie Elliott; Paul Farrelly; Julian Knight; Ian C. Lucas; Christian Matheson; Brendan O'Hara; Rebecca Pow; Jo Stevens; Giles Watling.

Questions 3146-3190

Witness

[I](#): Tristan Harris, Co-founder and Executive Director, Center for Humane Technology (via video link).

Examination of witness

Witness: Tristan Harris.

Q3146 **Chair:** Good morning. Welcome to this further session of the Digital, Culture, Media and Sport Committee's inquiry into disinformation and fake news. Tristan Harris, thank you for joining us from Paris this morning to give evidence. I appreciate that you have a pretty busy schedule today, and I believe you are meeting President Macron later this afternoon. Is that correct?

Tristan Harris: This morning. Thank you for grabbing me—I am really excited to get all this information out there.

Q3147 **Chair:** That's great. Many people, in speaking about Facebook in particular, have said that they believe that the early success of Facebook was based on the extraordinary amount of time people spent on the site, which made it quite an outlier compared with other platforms. Perhaps you could explain, for the record and the benefit of the Committee, how the addictive quality of Facebook has been such an important part of its success.

Tristan Harris: Sure. I guess, if it is helpful, I could give you a taste of my background in answering this question. It is an important lens to have. We often see people who use Facebook as making a conscious choice—think of the 150 times a day people check their phones, which we often think of as 150 conscious choices. When I was a kid I was a magician, and that teaches you to see this word “choice” in a very different way. The whole premise is that, by controlling people's attention, you can control their choices.

The thing about magic—this is important, because it really speaks to the framing for all the issues we are going to talk about today—is that it works on all human minds, no matter what language they speak, no matter how old or young they are, and no matter whether they are educated or not. In fact, if they have a PhD, they are usually easier to manipulate. In other words, there are facts about the hardware of the human mind, and that gives you a reference point for asking, “Are there facts about how to manipulate and influence any human mind that your audience won't know?” The answer is, of course, yes.

There is a set of techniques that are used in the tech industry under the guise of creating engagement that mask other problems like addiction. They are basically about hijacking the deeper underlying instincts of the human mind. A simple example, which you have probably seen before, is the slot machine, which has a variable schedule of rewards. A slot machine handles addictive qualities by playing to a specific kind of pattern in a human mind. It offers a reward when a person pulls a lever. There is a



HOUSE OF COMMONS

delay, which is a variable—it might be quick or long. The reward might be big or small. It is the randomness that creates the addiction.

Whenever you open up, say, Facebook, you do not know what you are going to get. You are literally opening it up and playing a slot machine to see, “What is it going to show me from my friends’ lives today?” When you click on notifications—when you tap that red notification box—you are playing a slot machine to see, “What am I going to get?” Every time you scroll, you might as well be playing a slot machine, because you are not sure what is going to come up next on the page. A slot machine is a very simple, powerful technique that causes people to want to check in all the time. Facebook and Twitter, by being social products—by using your social network—have an infinite supply of new information that they could show you. There are literally thousands of things that they could populate that news feed with, which turns it into that random-reward slot machine. That is one technique.

Another one is something called stopping cues, which are like what a magician does. How do I know when to stop drinking this glass of water? Well, it has a bottom—there is a stopping cue—meaning that my mind has to wake up and ask, “Do I really want more water?” If I wanted people not to think about that and to manipulate their choices, I might remove that stopping cue. This is related to a study in which six people were sat down at a table with different bowls of soup. Two had a tube at the bottom that refilled the bowl with more soup as they were drinking it, and the others did not. The question was, “If you remove the stopping cue and the bowl of soup doesn’t finish, will people just continue eating?” The answer was yes. I think they consumed 76% more calories.

The reason I bring this up is that the way social media products are designed is that they want to take out the bottom—they want to take out the stopping cue—of the usage. The reason that there are these infinitely scrolling news feeds is that they act, basically, like a bottomless bowl of content, and you do not want a mind to wake up and think for itself, “When do I want to stop doing this?”

I am enumerating just a few techniques that are taught at conferences and in books like “Hooked” that basically teach you, “These are different ways of influencing people’s choices.” Together, they create what we are calling addiction. I think the word “addiction” is not a very helpful word, because it makes it seem as if there are lots of substances that are addictive, such as tobacco and alcohol, and phones are just another thing that is addictive. It is more nuanced than that, because behind the screen are 100 engineers who know a lot more about how your mind works than you do. They play all these different tricks every single day and update those tricks to keep people hooked. There are many more techniques that I could go into, but that is probably a good entrée.

I will give one more example, which is more about how Facebook would do it. Let us say that there is a user who goes dormant for a while. There is a user who was active, and then they stop using the product and use it not at all for a week. If I am Facebook and I want to reactivate that user,



HOUSE OF COMMONS

using social psychology, I can see that Damian has been browsing Facebook and just looked at a photo that includes that user. Now, I am going to pop up a little notification that says, "Hey Damian, do you want to tag your friend"—this dormant user—"in this photo?" It is just a quick yes/no question. You are sitting there, and you think, "Yeah, sure, why not?" So you hit yes. This dormant user then gets an email saying, "Your friend Damian tagged you in this photo." It makes it seem as though Damian made his own independent choice, when in fact Facebook is kind of a puppet master. It gets one person to tag someone in a photo so that this dormant user has to come back to Facebook. It is very powerful. It taps into their social approval, their vanity and their social validation. They have to come back and see the photo they were tagged in. It is very powerful.

I am not saying that Facebook does that exact thing specifically, but there are variations on those kinds of techniques that are used across Facebook, Twitter, YouTube and LinkedIn. LinkedIn does it a lot with getting you to invite other people to LinkedIn that you have emailed with recently.

This is the entire environment that we have immersed human minds in. There are more than 2 billion people who use Facebook, which is about the number of conventional followers of Christianity. There are about 1.8 billion users of YouTube, which is about the number of conventional followers of Islam. People check their phones about 150 times a day in the developed world—the millennials. These things have a totalising surface area over people's lives. They spend huge amounts of their time on them. The amount of time they spend is not an accident of, "Oh no, we made it addictive. We could never have predicted that." Actually, there is a deliberate set of techniques that can be used to create that addiction and that time-spent usage.

Q3148 Chair: When Mark Zuckerberg appeared in front of the Senate, he was famously asked how he made money out of a service that was free, and he said in a very clear way that Facebook runs ads. Do you think, regardless of the intentions of Mark Zuckerberg when Facebook started, that the company has become addicted to capturing as much of the attention of its users as possible so that it can run more ads against that audience?

Tristan Harris: That is a great question. My experience and knowledge of the industry—I should say that I am from Silicon Valley, I am the same age as Zuckerberg and I was at Stanford University when Facebook was getting started, so a lot of my friends worked there in the early days. I have a lot of experience of people who are in the industry, the things that motivate them and how they thought about these kinds of issues.

To answer your question, I do not think that in the beginning there was this evil, aggressive idea of, "We are going to steal as much of people's attention as possible," or, "We have to make so much money, because of advertising." In fact, advertising in this model did not show up until much later. What really drove a lot of this was the desire to win and grow. In Silicon Valley, there are teams called growth hackers. Growth hackers are



basically the magician teams or the psychological manipulation teams. Chamath Palihapitiya was the famous one from Facebook who came out saying basically that the “dopamine-driven feedback loops we have created are ruining the fabric of society”. He stepped back from Facebook.

The teams are incentivised—literally how their role is measured inside the companies is based on how much growth they can drive. Their job is to figure out what the new species of notifications are. They have to think, “What are the new kinds of things we can do to email people to seduce them to come back to the product? What are the new techniques we can add to the design of the product? Do we need a red coloured button or a blue button?” Say they want consent for the GDPR and users have to click a consent button on the product. They say, “Let’s change the choice architecture. We will wait to show you the consent box on a day that you are urgently looking for the address of that place you have to go to. We will put the big blue consent box on top of that place so you basically have to hit consent.”

Q3149 Chair: There is a case in China, I believe, where video games companies have—whether required or not, they have done it—changed the algorithms of some of their games to make them less addictive, so that players aren’t rewarded for extended periods of play. Do you think something similar should be looked at for social media platforms such as Facebook, so that people are not rewarded as they are currently for engaging with the platform as much as they do?

Tristan Harris: Yes, Facebook has said themselves—since they adopted Time Well Spent—that they were able to reduce the amount of time people spent on Facebook. I think it was something like 50 million hours that they were reducing usage by—simply achieved by the design.

For example, Facebook can make the news feed as addictive as possible—there’s a spectrum, and they control how far along the spectrum they want to go. They could make it so that it is maximally addictive or least addictive. For example, they currently sort the news feed based on, “What are the things that I can tell you that most cause you to want to keep scrolling?” They have a machine casting billions of variations of this across these billions of human animals, and they get to figure out that these are the things that kind of work, and these are the things that don’t work.

On that spectrum, they could set the mark way over here at maximum addiction, or they could set the mark so that if they decide to get the captive audience less interested, they would show you the more boring cat photos of your friend’s life, which you are not interested in. They set the dial; they don’t want to admit that they set the dial, and instead they keep claiming, “We’re a neutral platform,” or, “We’re a neutral tool,” but in fact every choice they make is a choice architecture. They are designing how compelling the thing that shows up next on the news feed is, and their admission that they can already change the news feeds so that people spend less time shows that they do have control of that.



HOUSE OF COMMONS

The core change and shift in responsibility needs to be to, "It's not up to us, the user, to use this product differently", from, "If you are having a problem with it, it's your problem. You have to not use it as much", "It's free choice—if you don't like this product, there is a different one", or, "It's a free choice—if you're addicted, just have more self-control, use it differently." Instead, they need to see it as their responsibility.

It's not just about the time spent—that's not the core part. The key thing about the addiction is that it is almost like it sets up the matrix. Now you have 2 billion people with this addictive element, and social psychology and all this puppet master stuff knocking about. That kind of sets up the 2 billion-person matrix where now everyone is jacked in, because from the moment they wake up in the morning and turn their alarm off, they start seeing things that come from other companies, to the moment they go to bed, when they set their alarms but keep themselves jacked into a 24-hour cycle—that sets up the first layer, although my concern is less about the addiction than about what this then does when people's thoughts and beliefs are propagated through society by the fact that you so effectively hijacked 2 billion human beings into this system.

Q3150 Chair: A lot of our work in this inquiry has looked at the impact of Facebook on political campaigns, and the spreading of information about issues and events. Is this something that you have looked at? Do you feel that maybe this is something that has taken Facebook by surprise—that they were not aware of how open their platform was to being exploited by people who wish to spread disinformation or highly partisan content?

Tristan Harris: I am very interested in this issue and have studied it. Facebook has been aware of the spread of certain kinds of information for a long time. They have had efforts, stretching back years, to work on clickbait, dramatically reducing the kinds of things that were just sensational—looking at the kind of language patterns used in headlines, the way that people write those headlines to get the most clicks—to try and disincentivise those things from showing up at the top of news feeds. So they have been aware of that for a long time.

What has happened now is that there is divisive and incendiary speech, combined with clickbaity-style headlines that have gone through. It has been shown through different research that there is an intrinsic bias in favour of a word that is the most sensational, the most divisive and the most incendiary.

Someone who works with us at the Center for Humane Technology, Guillaume Chaslot, was a YouTube engineer who showed that YouTube systematically rewards the more divisive, conspiracy-theory-engaging speech. For example, a pattern like, "Hillary destroys Trump in debate"—that language pattern, "A destroys B in debate"—will be recommended at the top of the recommendations in YouTube so that, over time, if you airdropped a person on to a random video of Hillary and Trump, two videos later you would be in one of the more extreme ones. He showed that if you airdropped on to a regular video about 9/11, two videos later on autoplay it will put you into a 9/11 conspiracy theory video. If you



HOUSE OF COMMONS

airdrop a human being into YouTube and they land on a regular moon landing video from the 1960s, two videos later, they are in flat Earth conspiracy theories. So there is an intrinsic bias of these platforms to reward whatever is most engaging, because engagement is in both the business model and the growth strategy.

If I am Facebook, I will probably say, "Look, these guys at the Center for Humane Technology may think we are all about the advertising business model, but we do not even optimise for revenue." But they do optimise for engagement. They just said it in the last week. The guy who ran the news feed said that they do optimise for daily active users. If they are going to keep coming back every single day, which is how they measure success, and they have to grow that effort, then they have to show you engaging things. If it is not engaging, they cannot keep a business going. That is the core problem behind all of this: we have put this kind of Pavlovian Skinner box thing in people's pockets where you play slot machines all the time, and it is sorting what it shows you by "most engaging", so it is this race to the bottom of the brain stem, to the bottom of the lizard brain, to activate our core instincts.

Even if one company does not want to grow the amount of attention that it has—say you are Facebook and you just have 30 minutes of people's day, and you don't want to grow 30 minutes, you just want to hold on to that 30 minutes—well, guess what? Years go on and months go on and there are more and more competitive threats because there is only so much attention out there: Facebook is competing with YouTube is competing with Discord is competing with TBH is competing with Reddit. As more and more things show up, they can't get attention the way they got it yesterday—they have to get more and more aggressive, so they start adding slot machines, the puppet controls, the bottomless bowl and things like that. It gets more and more persuasive over time, which is why we are feeling, as a species, across the world, that we are losing agency and we are feeling more addicted. It is because every year each of these companies has to reach deeper down in the brain stem to hijack us at a deeper sociological level.

It goes very deep because it is not just our behavioural addictions. It is not just behaviour modification. You can also go deeper down to people's identity and self-worth. For example, Snapchat has taken over the currency of whether or not kids believe that they are friends with each other. Keep in mind that Snapchat is one of the No. 1 ways that teenagers in the United States, and abroad, communicate. It is very, very popular. Snapchat and Instagram are the two most popular products. The way that kids feel like they are friends with each other is whether or not they have been able to keep up their streak, which shows the number of days in a row that they have sent messages back and forth with each of their friends.

In other words, if you have two friends, a third party—a puppet master—has made these two human beings feel like they are only friends if they keep that 156-days-in-a-row streak going, so now it is 157 days in a row



HOUSE OF COMMONS

that they have sent a message. It is like tying two kids on two treadmills, tying legs together with string and hitting start on both treadmills, so now both those kids are running and they keep sending the football back and forth every single day, and if they don't, the other one falls off—except the thing that has fallen down is not the other kid, it is their friendship. Kids actually believe on a self-worth level, or a social approval level, or a social validation level, when they are very vulnerable on that stage, that they can only keep up their friendship if they continue doing that. That is all orchestrated by the companies.

I think that we have to see that there are many different issues emerging out of the attention economy. The externalities range from public health, addiction, culture, children's wellbeing, mental wellbeing, loneliness, sovereignty of identity and things like that to election, democracy, truth, discernment of truth and a shared reality, anti-trust and power. There are multiple issues. There are even more, because when you control the minds of people, you control society. How people make sense of the world and how they make choices are what ID is, and that can affect every aspect of society.

Q3151 Chair: A brief final question from me before I bring in other members of the Committee. From what you are saying, on the political side the danger is that no matter how people start to engage with political content, the platform will direct them to more extreme versions of whatever it is they have started to look at, because that is the way you maintain people's engagement.

Tristan Harris: Yes. I don't have specific evidence about what Facebook does to do that, but in general, if you just look at it as a game theoretic situation, each actor is locked into needing to be more and more persuasive and engaging. They have to be persuasive over time. Netflix's recommendations have to get better and more perfect; YouTube wants to make you forget why you came to YouTube and what you wanted to do originally, and instead make you go to the next video on YouTube, and that is going to get stronger and stronger over time. Tinder wants to show you one more perfect reason why you shouldn't be with the person you're with, and Facebook wants to show you one more engaging article. This is only going to continue, which is why we need to change course.

The structure of the larger conversation is about how you make that work, and instead build on human nature, which is fixed. We are living inside an increasingly adversarial environment in which super-computers are pointed at our brains playing chess against our minds, and trying to figure out what move they could make inside the chessboard of our mind, and what is the most engaging thing, before we even know what's happened. That is the reason why we go to YouTube and we think, "I'll watch just one video and then I'll go back to my work". Suddenly we wake up two hours later and we are like, "What the hell happened?" It is because it was playing chess against your mind, and it knew there was a video it could show you before you knew that that video would keep you for the next two hours. Again, the background trend is engagement.

Chair: Thank you. I will bring in other members of the Committee. It is possible that some people might be off screen for you, but hopefully you will bear with us.

Q3152 **Ian C. Lucas:** Thank you. Tristan, what do you think that the platforms fear most from legislators?

Tristan Harris: That's a great question. They fear any kind of regulation because they are obviously thinking about its value. They fear anything that would slow them down because they need to grow fast.

Q3153 **Ian C. Lucas:** What do you think the legislators should be doing? You have outlined some quite harrowing evidence concerning young people. What should legislators be doing in that area?

Tristan Harris: I think there are several things. Fundamentally, this is a whole new class of political, social and economic actors. We know what a bank is, and we know how to regulate a bank. Banks do one thing. We know what health products do because we live inside the health domain. When you have a new tranche of products, they are not just products—they are basically societal environments. They are the public sphere; they are democratic norms; they are election laws.

The way to think about this—I am zooming out first before I answer your specific question—concerns Marc Andreessen's insight that technology can do every industry more efficiently than anything that can be done without technology. We have given the tech sector a free pass on regulation. The phrase "software is eating the world" means that deregulation is eating the world. Just take media and advertising. In the United States laws prohibit certain kinds of political advertising—fair pricing for political advertising on TV and so on. When software eats that, it eats it up and you have Facebook running that infrastructure. You just took away all the protections. There is no more regulation for any of the things that we used to have regulation for.

In every single sector, whether that is media, advertising, children's health, what we allow children to see, data privacy—all that stuff—we have to reintroduce that in the digital landscape. One thing you can think about is, in all areas where technology is eating up parts of society in different industries, what were the protections that we had in place that we would need to re-establish in that part of the digital world? For example, I don't know how it works in the UK, but in the US we have fair pricing for election ads. It should cost the same for each candidate to run an ad at the same time in the same TV slot. That is regulated by the FCC. On Facebook that is not true. In fact, at the last election there is some private information showing that there would be a 17 to one difference in the cost of ads for one party versus another. That is one example. Take all your laws that protect children, elections, media and ask, "What are the ways in which you are stripping us of those protections and what would it take to reintroduce that?" That is going to put it on some companies. But more than that, that is kind of table pick—it is like, what do we lose when we introduce all these things?



HOUSE OF COMMONS

Then there is the new class of threats for which we didn't have regulations on previous domains. We didn't have that new class of threat. A good example is micro-targeted advertising. So long as I can target you, based on something about you that you don't know that I can target, I can find conspiracy theorists in the language of a country that I don't speak but I can know that that country has conspiracy theory beliefs around—I don't know, I can make it up. Then I can go to that country and start amplifying those conspiracy theorists and I can do it all the way from here in the United States, from Paris, from Malaysia, from anywhere. That is a new kind of threat that we haven't needed regulations for in the past, because we didn't have that class of threat.

I personally believe that we need a ban on micro-targeted advertising. I think there is no way in which Facebook can control the exponentially complex set of things that people let advertisers enter into the platforms to target users' baseline. For example, basically engineers only speak so many languages and they are stirring the pot of 2 billion people who have languages that they don't even speak, including in very emergent sensitive markets where they have just got the internet like in Myanmar and Sri Lanka. It is possible for people to create culture wars in those languages because there is much less oversight by Facebook and different parties for how that is happening. I think that is one of the areas that we are going to need regulation on—the micro-targeting of advertising.

We are going to need to reclass what these services are—specifically reclassing the legal status of what these services are. If I asked you, "What is Facebook?" Is it a technology, is it a social network product, therefore it is in the same category as other social networking products, basically an advertising company? Or is it an AI company? These are different ways to see what Facebook is. I would argue, right now, Facebook is being governed by the wrong kind of law, which is contract law, meaning peer-to-peer. You are a user, you hit okay, and you sign on the dotted line, therefore if something happens, like something goes wrong, or your data doesn't work or you don't want it to go, or you get advertising for something that you didn't know, or pressure or something like that. Too bad, because you have already signed on the dotted line. You have accepted these evolving terms and conditions. If you don't remember accepting, then we will get you to reaccept and we will do it again on a day when you have to accept it. You probably won't think about it. They're really good at doing it that way. There are hundreds of engineers making sure we hit that big blue okay button.

That is a peer-to-peer contract relationship. Instead, what we need to have is a fiduciary relationship. A fiduciary relationship is where one party has the ability to exploit another party because they have asymmetric powers—think of an attorney-client relationship. The attorney knows more about the law than you do and you give them all this personal information that can completely exploit your interests. If the attorney does exploit the interests of the other party, they can't just say, "You signed on the dotted line, so it's your fault, because you agreed to it". Of course, we say that an attorney has a different status because they have asymmetric power

over their client. It's the same way that a psychotherapist can exploit you in the ways that are amenable to needing protection. That is also a fiduciary relationship.

If you wind up how much power an attorney has over their client, how much power a psychotherapist has over a client, next to how much power Facebook has—how much does it know about how to exploit the user compared with those other two examples? Based on that alone, it is orders of magnitude. It is like a priest in a confession booth who is listening to 2 billion people's confessions, the same priest, except they know the confession that you don't even know you are making, because they pointed AIs at your brain and they know what colours, buttons and words light up your brain. And they know every conversation you have ever had with anybody else on Facebook. Their entire business model is to sell access to that confession booth to a third party.

That is an enormous problem, in that we should never allow priests in confession booths to have a business model of selling that to a third party. To me, that instantly indicates that the advertising business model is a no-go, for something with that kind of asymmetric power to exploit people. It is one thing for them to claim, "We care about our users," but its business model is explicitly to enable anyone, including Russia and foreign state actors, to manipulate the populations of anyone in that confession booth. The worst part is that people do not even realise they are in that confession booth.

People also do not realise that it is a priest with a supercomputer by their side, which is making calculations and predictions about confessions you are going to make that you don't know you are going to make. There was an article in *The Intercept* about Facebook's long-term business model, which is about the ability to tell advertisers about when you are going to churn or not be loyal to something before you know you are not going to be loyal to something.

This is an enormous and new species of power that we have never seen before, and we have to have sensible ways of reining in that power. We have to do it in a geopolitical context, knowing that other companies in China and elsewhere will have similar kinds of power, so we have to think very carefully about what it means. But at the very least, in the same way that banks have a fiduciary responsibility to the people they serve because they have asymmetric knowledge and the ability to exploit them, we have to think about that for these social media companies—especially Facebook.

Q3154 Rebecca Pow: Thank you very much for joining us. I wanted to ask you something quickly. You touched on the health aspects of some of the things you think Facebook might be causing. In a word or two, do you actually think Facebook is dangerous to children's health?

Tristan Harris: Is it dangerous to public health? It is less about Facebook and more about the model. There are probably some specific ways in which Facebook endangers parts of public health, but before I go there let me just think about it. I think about it this way: we are evolved animals



HOUSE OF COMMONS

and we have a limited brain, a limited mind and certain vulnerabilities that work on us. If you took a human animal and put it in a cage, and when it woke up you made it look at photos in sequence of its friends having fun without it—photo after photo after photo of its friends having fun without it—that would have an impact on any human being. If I could show you photo after photo of your ex-romantic partner smiling with someone new, that would have an impact on your mental wellbeing.

The question is, how have we designed this environment that surrounds 2 billion human animals? Because this business model of engagement and these products such as Instagram—and photos are much more engaging in many cases than Facebook and just lots of posts, especially for younger generations—we are basically surrounding human beings with evidence that is actually false. We are constructing a magic trick, a false social reality of the highlight reels of our friends' lives.

Q3155 Rebecca Pow: Do you think there would be any way of analysing that, so that you could prove that this model was causing a danger to society? I only ask because that might help us to target this if we need to.

Tristan Harris: I'm sorry, I didn't hear the very beginning of that.

Q3156 Rebecca Pow: Earlier, you said that 2 billion people's minds were being hijacked through this system. We have to come up with some recommendations if we think this is a danger; if this is one of the aspects that is a danger to society, how are we going to measure that to prove it? Is there any way of doing that?

Tristan Harris: It is a challenge. It will be very much like sugar and tobacco. Let's take sugar as an example for social psychology and health. Sugar will always taste really good, even after we know it is bad for us. We cannot beat the fact that it tastes good for us. We can increase the price of it, but it will always be in our food supply. Social validation and social approval will never go away. They are like sugar, especially for younger people. It is more a product of a public and cultural awakening to the fact that this has such a tremendous impact on us. For as long as we are surrounded by evidence of our friends having fun without us, it is very hard to regulate that and say that people should not be surrounded by that. I just think people need to know that, actually, it is going to have that negative impact, and there is going to be a cost on our lives; and the costs do show up in analysis—specifically attention and cognition. More frequently switching between tasks—back and forth, multi-tasking—lowers our IQ. It has been shown that when we get interruptions in a work context it takes us 23 minutes to resume focus after an interruption. We cycle between two or three other projects, usually, before we get back to the thing that we are doing; so every time you get distracted by a Youtube video, or something like that, you don't just go right back to working and you are right back in the flow. You check your email and you do something else, and something else, and then you get back to maybe starting to get back in the flow. The costs are about so many different layers of society.

Q3157 Rebecca Pow: So who would this be a role for academics or regulators;



HOUSE OF COMMONS

or, indeed, should we be interfering at all? I only give this example because in the UK we have just brought down the limit on gambling machines—a £2 limit—because of the dangers of gambling. You could say is this a similar thing, and how on earth do we pinpoint it to say this is a danger?

Tristan Harris: On the gambling thing, I think that what we need are certain standards at the very least, and this is something that we are advocating for with our organisation—our humane design standards. So, for example, how should a phone or an object that sits up against the human skin of 2 billion people not operate like a slot machine. If we knew the threshold by which at a certain frequency of dosing out these little rewards of dopamine, that that creates addiction, can we set a threshold or a standard by which mobile phone manufacturers would want to stay underneath the threshold for trying to be this dopamine-and-slot-machine-style incentive? That is kind of like what we do in gambling; I don't know all the limitations but I know that the Nevada Gaming Commission and other organisations are trying to set a standard for how slot machines work, so that it has to be genuinely random—the house can't always win. There are certain kinds of mathematical standards.

The challenge here—it is more complex and messy, because it is not as if Apple and Google, who make the phone themselves, are trying to set up what that dopamine slot machine thing is. It is that they created an environment in which, with all these apps competing, the apps act like slot machines. So it is going to take redesigning it from an OS level, from Apple and Google, to be more considerate; and there is a shareholder letter to Apple right now about the toxic public health problems that emerge from addiction, for children especially—and I think that conversation needs to be amplified. I will say that Google recently announced the Google wellbeing initiative, which is just baby steps, honestly, in starting to answer these kinds of questions. What we really want to do is turn this from a race to the bottom of the brain stem to capture attention to a race to the top of the marketplace where people protect and care about users—care about people.

Q3158 **Rebecca Pow:** Apparently, Tristan, it has been shown that addictive gamblers don't want to win; they just want more time on the machine, so would not that actually be playing into Facebook's hands—this kind of addiction?

Tristan Harris: What was the bit about more time on the machine?

Q3159 **Rebecca Pow:** Sorry, I probably wasn't facing the microphone. It has been shown that addictive gamblers don't want to win. They want greater time on the device. So wouldn't that actually be what is happening on Facebook as well? That actually then plays into the hands of Facebook, because then you are going to be targeting them with more adverts and more of the stuff you want to indoctrinate them with.

Tristan Harris: I am sorry; the audio quality didn't come in super-clearly, so I didn't hear all of that: something about how Facebook could benefit

from some kind of regulation around gambling in the same way that we can creatively—is that what it was about?

Q3160 Rebecca Pow: Shall I just repeat it one more time; I do not know if this microphone is working properly? It has been shown that addictive gamblers actually just want more time on the device. That is what they want. They are addicted to using the device; so isn't that actually what is happening on Facebook? The more they use it, the more they are addicted, and the harder it is going to be for anyone to come up with a plan or a policy to get them out of that addiction.

Tristan Harris: Yes, exactly; and in the case of gambling you have to at least drive a few miles in a car, as a gambler, to do that. We can say there is some amount of responsibility that a person has in that gap of choice making, where they really have to make a conscious choice. In the case of smartphones, you don't get to choose. There are ways in which the slot machine is kind of built right next to your alarm clock, so right when you turn off the alarm on your phone—they are right there in the app switcher. They are already open oftentimes, or there are notifications there. This challenge is about regulating choice architectures and how we enable people to have the agency to at least have the freedom to put slot machines far away from them, and also to regulate how they work and how addicted they are, versus being in a world where they cannot control how close or far away from them those slot machines are.

A good and slightly different example of this is that blue light screws up our circadian rhythms as human animals. When you look at screens and see the light that comes out—that white, shiny, bright light—it has blue light in it. If you look at that late at night, it actually has an impact on your circadian rhythms, and you don't get to choose that. In this space, Apple introduced something called night shift and Google introduced some night setting modes to basically strip the colour out of the screen, or to turn it into greyscale.

These are the kinds of things that, if we really cared about people, we would design them to align with our human biology. A slot machine redesign would do the same thing. We would say that, given that human beings are vulnerable to being dosed with dopamine rewards at a certain rate, we would have to design it differently, so that it wouldn't interact with human biology in that manipulative way. I am not sure if that helps, but I am trying to answer the question.

Q3161 Rebecca Pow: Can I just ask one final question? Facebook have already made some changes to their policies and their privacy policies and all of that. Do you think that that is effective? Have they made enough changes? Should we leave it up to them, or should we interfere and bring in regulation?

Tristan Harris: Did you ask whether they have made enough changes and whether regulation is needed? The audio quality is not perfect.

Rebecca Pow: Yes. Facebook have already brought in some changes, on a voluntary basis, to transform their structure and business model. Do you

think that is enough?

Tristan Harris: This is a question about self-regulation. I would love for the industry to self-regulate to the full extent that it needs to. The challenge, which I think we have all noticed, is that, when given the opportunity to regulate themselves, they haven't done nearly enough.

It took years for them to make the simple changes, which they announced recently, around Time Well Spent to try to reduce the amount of time that people spend on Facebook. That involved a lot of people like me trying to raise awareness about the negative public health consequences that emerged. They don't do these things until there is a huge cultural pushback, and they generally wait until the last minute. For the Irish referendum, they only made changes on ads to protect that election about a week or something before the election. With the Russia investigation in the United States, they only released the number of people who were actually affected—126 million—the day before they had Congressional hearings.

In terms of being transparent and accountable, we don't seem to be able to trust Facebook, at least, very much with what they are doing. We need proactive responsibility, not reactive, which is why we might need regulations. Those would be very hard to figure out, especially for a general Government agency. We need more specific expertise areas to provide oversight of these companies. If they have that expertise, they can help to start crafting some of these regulations.

Q3162 **Giles Watling:** Thank you for coming and talking to us this morning, Tristan. In many ways, the evidence you have given us this morning is amazing. I would like a video on the sort of evidence you have been giving to be regularly shown on prime-time television, so that we can all clue in to exactly what addicts us to these things.

You said about joining something like Facebook or one of the other social media platforms that people want to join in and see their friends, and that they won't see the videos that they want to watch unless they agree to the terms and conditions. That is almost forcing them to agree to terms and conditions that they won't necessarily read; in fact, I don't believe that many people actually do. Am I right in saying that that is the way that it is designed to hook people in?

Tristan Harris: Yes. In the growth-hacking world, one strategy is to get someone's friend to tag them in a photo, so that someone who is not yet on Facebook receives an email saying that their friend has tagged a photo of them. They are not on Facebook yet, so to see the photo they have to sign up for Facebook. That is not something that Facebook specifically does, but it is a technique that is part of the growth-hacking literature on how technology companies grow their products. It is to try to create rewards that can only be accessed by signing up for the service and agreeing to the terms of service.

Again, in a moment where I most need to access, let us say, the address of a Facebook event I am going to tonight, and when I open up the app on



HOUSE OF COMMONS

that occasion I am in a rush, if that is the day that the GDPR with a long list of new policies comes up with a big blue button, I am not going to think about the blue button or what it said. I am just going to get to that address. We need a new understanding, from a regulation of consent, in which people have less and less time and are doing things in more and more urgent situations, especially on phones, to a regulation of default settings. We need default settings and default rules that work better for more people. Public education, for example.

Q3163 Giles Watling: Could we remove that in a regulatory way? Could we dismantle this system of set-up so that when you engage, you can just engage? You do not have to agree. You can just engage or go away. You do not have to agree. Is that what you are saying—a default setting?

Tristan Harris: Yes, I see what you are saying. There could be certain standards of access, almost like net neutrality or something where you can access core things without needing to sign up. I think that is the kind of thing you are getting at. Companies would certainly not be happy with that little drop in their sign-up rate and growth rates, but that would be something interesting to look at.

Q3164 Giles Watling: We have had the gambling analogy, and there is another analogy. I would like to know whether you agree with it. It is like the tobacco industry, which, in many ways, was kind of a perfect industry. Tobacco is really cheap to produce. Everybody gets addicted. You can then begin to charge what you like and then it takes generations to kick the habit. The western world has largely kicked that habit and we are moving towards a tobacco-free environment. Everywhere we go now is tobacco-free. Can you see that as a future for this kind of addiction online?

Tristan Harris: That's right, and I want to get to something else you say in that, which is that this is like a new species of tobacco. This is like a tobacco that has an AI engine inside of it that is updating how it works to perfectly match your specific human weaknesses and psychological vulnerabilities. It uses your social psychology and your dependence on your friends and your need for belonging to make it, so it is impossible. I have studied cults in my background in studying persuasion, and that is what cults do. You control the currency of belonging. Once I am in and my friends are all part of that group, I cannot derive my belonging from anywhere else; you have got me.

Let us say as an example that I am a teenager and I am listening to Tristan and I say, "Gosh, I realise I don't want to use Snapchat any more; it's trying to manipulate me. It's just addictive." I am a teenager, I am 15 years old, I live in the UK and I don't want to use Snapchat, but guess what? In the case of Facebook, for example, or in the case of the Snapchat people, if my friends are only talking about where they are going to hang out on a Tuesday night on Snapchat and they didn't use anything else, I have just cut off my access to belonging. I have cut off my access to social opportunities, sexual opportunities and development opportunities, in the same way that if I don't use LinkedIn. Once they own



HOUSE OF COMMONS

the currency of finding a job, I have to use LinkedIn. In the Facebook case, there are teachers who assign their homework through Facebook. If a user says, "I don't want to use Facebook", but then suddenly their homework is assigned through Facebook and that is where the classwork and sessions happen, then that person has to use it. So that is the challenge. Capturing our social psychology is what is new about these services rather than just being backup.

Q3165 Giles Watling: I'm afraid I'm old enough to remember that it was a social thing to smoke. If you were outside the circle, you were not taking a cigarette, and you went into a social situation, you were always offered cigarettes. I'm afraid that is how it was and this seems very similar to me.

I want to ask one final question if I may. It was said about you that you were the closest thing that Silicon Valley had to a conscience. Do you think you are pretty much alone out there? Do you think there are other people moving towards improving the way these platforms work in Silicon Valley?

Tristan Harris: It's a great question. We are trying to create a movement here and we are trying to say that this has gone really horribly wrong. It is not because we are evil people. In fact, a lot of people in the tech industry have good intentions. We are trying to collect those who have a conscience and who really do see that we need to do things a different way and are willing to face the fact that it might require a different business model to do things a different way. That movement has been growing. Many insiders inside these technology companies are big fans of this work and the need for change. Chris Hughes, the co-founder of Facebook, recently became a supporter of our work—he has endorsed our work and is an adviser. Roger McNamee, who was Mark Zuckerberg's mentor and introduced Mark and Sheryl at Facebook, and I were part of starting the Center for Humane Technology. There is a set of characters that are increasingly coming out of the woodwork—because we are concerned about the long-term cost and damage that this is going to do if we don't fix it. So, I think this is a movement and it takes more people coming out. That has been happening progressively, the more people who have been coming out.

Q3166 Giles Watling: And those people now will work with us—will work with regulatory bodies, you feel?

Tristan Harris: I'm sorry—for some reason, the beginning of what you are saying is always very quiet—

Giles Watling: That group you were just describing who are coming out of the woodwork—who are developing a conscience, if you like—would be prepared to work with regulatory bodies?

Tristan Harris: That's right, yes. Our goal is to become helpful and a resource for companies that already exist to change their business models, design practices, develop standards—and to be a resource for Governments to develop sensible, common-sense policies and protections

that really deal with these problems. That is why I am here with you right now; that is what our goal is.

Q3167 Brendan O'Hara: Thank you, Tristan. It is fascinating testimony that you are giving. As a father of two teenage daughters, I have witnessed the emotional angst that is involved in the ending of a long-running streak on Instagram. It sounds ridiculous, but it is true, and I can witness that what you are saying is absolutely correct.

I would like to ask you a bit about Facebook. Every time you have spoken to Facebook, all the way up to Mike Schroepfer, they have always said—almost like a mantra—“We never thought, when we started Facebook, that we would end up in this position”. That may well be true, but if it is true, in your opinion, how far down the road was Facebook when they realised that this is exactly where they were headed? What choices and decisions did they take at that point, which took them to where they are now?

Tristan Harris: That is a great question. In terms of how much they knew, it is important to say again I did not work at Facebook. I know many who did. I think that in their addiction to growth and global domination, they have an almost fraternity—brotherlike, you know—achievement mindset. Like—kill it, grow, move fast and break things, done is better than perfect. When you go to the company and see the posters on the walls, those are the kinds of statements that you see. They very much have a growth-orientated mindset, more so than anything else. There is also an invisible libertarian bias, that it is not our responsibility how these products we make affect things, affect the world. We are just building a product that can come to you, the user: it is up to you, the society, to choose how you will use this neutral object that we have created. I have always thought that that was not true at all. The presentation, for reference, that I gave at Google in 2013—I made a slide deck that was viral. That is how I came to be doing this synapcistic sort of work. Basically, what I claimed was a moral responsibility that Google had in shaping 2 billion people’s attention and that we had an important moral responsibility to get that right and to think about the consequences of what would happen downstream from designing the products. In terms of when Facebook realised that, I think that it should have been obvious years earlier.

Here is the problem: as an engagement-based model, you just want to see the numbers go up. If you think about it, Facebook is 2 billion people doing billions of actions every day, billions of clicks. So the only way you have of understanding what is happening inside your 2 billion-person country called Facebook is by using metrics—by measuring how many things are happening. So, you measure a few things. What do you measure? When your business model is advertising, you care about how many accounts were created. You just want that number to go up, and you don’t really care whether those accounts are real accounts are not, or Russian trolls, or bots, or fake accounts, or sock puppets. You just want accounts to go up, and then you want content to go up: how much content is getting posted on the platform?



HOUSE OF COMMONS

YouTube just says, "How many videos were uploaded?" Facebook says, "How many posts were created?" You want to count how many friend connections are created and how much time spent was going up—time spent is the biggest indicator for engagement. You innocuously start to pay attention to those core, simple naive metrics. That is what caused the problem. They had all the user growth numbers go up, but they did not check whether they were real or Russian. They had all the advertiser campaigns go up, but they did not know that there might be nefarious bad actors, who might be using it deliberately to create culture wars.

In having all the content generated on YouTube, they do not know whether those are just conspiracy theories and people are trying to deliberately manipulate populations of people. They just want activity and growth at all costs. I really think that is a lapse of consciousness and responsibility, because to me it does not take an enormous amount of insight to see what kind of impact they will have downstream.

In fact, I do not know if this person is coming to your Committee, but Antonio García Martínez, who worked at Facebook, said that they used to joke that if they wanted to, they could steer elections. Facebook and Google probably both know before any election happens who will win that election, because they have so much data about what everyone has been posting before every election and the analysis of positive and negative sentiments that are posted around each word. They know before any of us do on the outside where elections will go.

Knowing that they knew that years ago says to me that they would know that there are things that they could predict about where culture, mental health for teenagers, elections and manipulation will go. They should have had a much deeper awareness from the beginning; especially after the 2016 US elections and Brexit, it should have been enormously clear that they had a role in that. In my mind, the biggest evidence to me that they have not had that consciousness is how long it has taken them to get to where they are now—only through huge amounts of public pressure, and not through their own sense of responsibility.

Q3168 Brendan O'Hara: One of my younger, trendier colleagues informs me that Snapchat has streaks, not Instagram, as I said. Going back to Facebook, Sandy Parakilas said when he gave evidence to the Committee that Facebook collected personally identifiable data on people who had not specifically authorised it. They then intentionally allowed that data to leave Facebook without any controls once it had left. Is that a fair assessment in your opinion?

Tristan Harris: I would trust Sandy's judgment on that. I am not an expert on exactly what happened with the Cambridge Analytica and Facebook app platform, but the entire premise of Facebook's app platform when they designed it was to enable third-party developers to have access to people's friends data. That was the purpose of it; and that is what is missing in that debate about Cambridge Analytica. It is not that Facebook is great and the app platform was great but there was this one really evil actor called Cambridge Analytica, which took that data and they are so

bad; the premise of the app platform was to enable as many developers as possible to use that data in creative ways, to build creative new social applications on behalf of Facebook. Theoretically, that would be helping to the world to be better, using Facebook's data. They wanted to make that better.

For the reasons I described at the beginning about the attention economy, the app platform turned into a race for engagement. It turned into a game of garbage about how you basically manipulated people into poking each other back and listing their top friends; in this case, Cambridge Analytica filling out a personality profile to manipulate them.

We tend to focus on privacy and data a lot, and where data goes and whether people know where their data is going. We have to assume at this point that the data is out there—that there is a dark market where you can buy data on anybody. That information is out there. The most important thing that we do now is protect the means by which it can be used.

Just assume that there is some corner of the dark web where I can buy 2 billion people's political profiles and understand the sentiments that they have ever expressed about key topics such as immigration or something like that. That is just available now. It is not going to go away; it is now just out there. Now the question is: how do we protect the ways that that gets used?

Q3169 Brendan O'Hara: What was in it for Facebook to build a platform that made data abuse so easy, and what benefits did it get from developing platforms that allowed developers to access user data in this way?

Tristan Harris: In the case of the Facebook app platform, my understanding was—it's back to the attention economy—that it was creating and enabling all sorts of other people to invent reasons for people to come back to Facebook. If you think of how it is now, there are thousands of popular apps on Facebook that you have to use Facebook to do. That forces people to come back to Facebook and creates more engagement for Facebook, more time spent on Facebook, and more data collection by Facebook. So much of those application platform APIs, meaning the programming interfaces that people use to build Facebook apps, would usually require that those activities and events would come back, and Facebook would see all those events.

Suddenly, they get to capture even more data. They get to capture even more activity and time spent, which is good for advertising and good for telling Wall Street. They get to capture more users. This whole thing turned into, "How can we make Facebook grow?" Again, I think that was really the motivation, because the app platform came out in 2008 or 2007. I think that was really about growth. They were in very early stages, and they wanted to grow the platform to be much bigger.

Q3170 Brendan O'Hara: Finally, how important is surveillance of the internet to Facebook's business model?

Tristan Harris: How important is data to Facebook's business model? Is

that what you are asking?

Brendan O'Hara: Does someone else want to try asking the question? It may be this mic—or this accent.

Q3171 **Chair:** Brendan was asking about the importance of surveillance of the internet as a whole to the Facebook business model. I suppose that would include the way that Facebook monitors what people do off of Facebook, and monitors their other browsing history and sites that they visit, to support the data that they gather about those users—and non-users as well, obviously.

Tristan Harris: For the record, it was not just the accent; it was the volume. The question I heard from the Chair was: how important is the data that is gathered offsite to Facebook's business model in general? The thing about data is that it lets you do more and more prediction. I think that these companies—Google, Facebook and YouTube—should be seen as AI companies. They are not social networks; they are artificial intelligence systems.

The fuel for artificial intelligence systems is information that lets you do better prediction. They are specifically interested in the kinds of data that would enable them to do better predictions, and knowing that you browse certain kinds of websites is very helpful in knowing how to predict what you will be vulnerable to, what messages you will be influenced by or invulnerable to, and what advertisers might want to target you.

A simple example of this that emerged in the last few years is the lookalike model—I am sure you are all now aware of lookalike models. If I am an advertiser and I want to find an audience and I know only about 100 people who are interested in a thing I am interested in, Facebook can calculate based on those 100 people who looks like them in their behaviour. The more data I have about their behaviour across the internet, the better I can find them on Facebook, and the more alternative lookalike users—those who look like that initial seed group of 100—I can give you. That is one clear example of how having more data gives you access to be more and more effective for advertising.

What is really dangerous about this is that you can increasingly predict the perfect ways to influence and manipulate someone. That is what Cambridge Analytica was about, right? It is about saying that in addition to all the behaviour data you can collect, if you know people by personality traits you can influence them to an even greater degree.

I wanted to bring up the point that data has been important, but there will be other creative ways to learn things about people, even if they do not give you that information explicitly. There is a study by Gloria Mark at UC Irvine saying that you can actually predict people's big five personality traits—the same personality traits that Cambridge Analytica were seeking to obtain from that Alexander Kogan personality survey. You can predict people's big five personality traits based on their behaviour and click patterns alone, with 80% accuracy. This is important to understand,



HOUSE OF COMMONS

because in a way in the future there will be very little ability to hide who you are, what you are or your political beliefs—a friend of mine, Poppy Crum of Dolby, says this is “the end of the poker face”—because there are enough signals or signatures coming off your behaviour itself, based on how your eye movements work, your body language and your looks, that we can predict this information about you without you ever consenting to give it to me explicitly.

The dangerous thing in where this is going, and why I think some of the people I work with lose sleep—along with me sometimes—is that, because AI will allow you to predict things about human animals that they do not know about themselves, that can be used in dangerous ways if we do not protect against it.

A relevant question for regulation that someone might have asked Mark is: what are you willing not to learn about people? And what business model are you willing not to engage with that would override human agency? Because we have got human animals over here and we have got supercomputers pointed at them. Literally the biggest supercomputers in the world are inside two companies, Facebook and Google, and we have pointed those supercomputers at people’s brains to predict what will get them to do, think and feel things. That is a dangerous situation. We should have a new relationship with businesses that are based on that system. We could point AIs, artificial intelligence systems and supercomputers at climate change, drug discovery and solving problems, but I do not think we should be pointing them at human beings, especially not with the goal of selling access to that prediction and that ability to influence to a third party: the priest who sells access to the confession booth.

Q3172 Paul Farrelly: I am conscious that we have been focusing on Facebook to the exclusion of everyone else. I want to come back to other companies and technology, but, staying on Facebook for a moment, their answer could be, “Well, you get this for free—just delete your account.”

I do not want to delete my account because I find it very nice to keep in touch with friends, but it was my choice to open it up to a wider audience—possibly for political reasons—so in that sense I have made a pact with the devil. To the limits of my technology, I have tried to make a conscious choice to get away from Facebook Messenger, and, having deleted the app, it will not allow me to do that. People are still contacting me, and I do not want them to contact me on Facebook Messenger. That seems to be one simple choice that Facebook could respect internally, but there seems to be no full escape. Are there any other simple choices that Facebook could grant to people, so that the individual level may then moderate the need for a policy response?

Tristan Harris: It is a great question. Like you are saying, when you are a magician you are always offering people choices. Every choice looks like a free choice, but the way you have structured the menu means that you win no matter what they choose. Facebook and all tech companies are in a position where they want to structure the menu in a way that looks like you are making a free choice, where you don’t have to use it, you can just



HOUSE OF COMMONS

use Messenger. But if you use Messenger and you don't let it track your location, for example, it will notify you and put a message on the screen each time saying, "Please allow us to track your location." They are always rearranging the menu, so that they will get access to the things that they want, which is more access to locations, more access to your contacts and giving you reasons to upload your contacts, so that you can invite your friends. I am speaking more in terms of the problem than solutions. I have not thought about how to divvy up Facebook in the same way.

One example is that it used to allow you to do what you are talking about—Messenger—on Facebook. You could tap "messages" and then right from the web interface you could type a response back. Now they do not allow you to do that. They force you to download Messenger. One of the things that happened in the 1990s with AOL Instant Messenger—I think this happened through regulation—is that it was forced to become interoperable, so that you could use AOL Instant Messenger, which used to be a closed system under AOL's system, with other messaging clients. You could imagine a world where there is more interoperability—not just data portability, but fundamental interoperability. I can take my phone number and say that I don't want that number to be on AT&T or Sprint anymore, but I want to move it to Vodafone.

With Facebook I can't do that. I can't take my entire group of friends and say that I want to move my ability to communicate with them and everything I have ever said with them in the past—that entire stack of stuff—to this other place. They theoretically let you click download on this very messy zip file that archives all of this information in this unusable format, but they don't really make it interoperable in a clean way. If there was interoperability, there would be more competition and Facebook would be forced to improve its product, and there would be possible competitors that might be more ethically or humanely designed. I don't have a clear answer for you in terms of how to split it up just at the moment.

Q3173 Paul Farrelly: I just use that as one example of where the users or the product—we humans—make a choice to download an app for a particular feature, such as Messenger, and then make a choice to delete the app, because we don't want to be contacted. In my case, I can be contacted by email, phone and face to face. With that feature, I find that people come from all directions and, in particular, do not know how to moderate their language when they are addressing you. I simply don't want to use it. At the micro level, that is one example of where you try to make a choice, but the choice isn't fully respected, so the only answer is the nuclear answer, which is just to delete the account.

Tristan Harris: We call these all or nothing choices. You have to use the app and they do whatever they want to—a good example of this, going back to our previous discussion, is Snapchat. You cannot use Snapchat without them doing the streaks. You can't flip a switch off and not do that. There is another example from Snapchat—another manipulative thing that they do. When your friend just starts to type to you the first two characters of a message, it actually buzzes into the other person's phone—it interrupts the other person, the other kid—saying, "Your friend is



HOUSE OF COMMONS

starting to type to you,” and it builds anticipation for that. Snapchat does not allow you to turn off that feature, at least it didn’t the last time I checked. In general we need much more accountability for these all or nothing choices. Frankly, that is something our organisation should be more on top of, but many more people should be looking at these agency-inhibiting or disempowering architectures. This is an area where more oversight and pressure needs to be made, especially when it comes to the more vulnerable populations, such as children.

Q3174 Paul Farrelly: Clearly the concern here is not just about political manipulation but about making these platforms a more enjoyable, safer and more responsible place for people to be. Politicians clearly cannot specify individual changes to individual models and algorithms, or whatever you call it. Do you see, overall, that the competition between these actors to maximise face time and eyeball contact is such that there will have to be a policy response, because everything that they will do will just be tinkering at the edge?

Tristan Harris: One thing that just came to mind as we were just talking was having third-party reviews of metrics and things. We have this problem where there are not enough choices on the menu when people have real problems. Like you said, there are ways to rearrange or re-rank news feeds, which they are not doing themselves.

Right now there is a problem that, first, there is no transparency. We on the outside do not know how news feeds are ranked. There is also no accountability, so when we want to change them or we encounter a problem, there is no way to force them to do it. There could be a different system by which they are forced to have expert reviews with outside bodies. You have seen Facebook start to make that kind of change around elections recently, when they set up a research group.

I say this because there is a need for external oversight by trusted actors who can basically help them to make those choices. You can’t accommodate everyone’s choices, but some of this work will happen through forcing them to look at those things with outsiders. But again, like you said, so long as those outside changes might fundamentally inhibit their business model—if they would lose 30% of their revenue because it turns out that people don’t like spending so much time on the screen and would much prefer to spend that time in person; Facebook previously made that happen by prioritising Facebook events and in-person ways of being with each other—Facebook will not do that, because its primary goal is still funded by the advertising business model.

It comes back to regulation around the advertising business model and ways of making it easier for people to pay. By the way, there is a contradiction in their argument about it being elitist to suggest that people should pay for Facebook. People say, on the one hand, “How dare someone say that people should pay for Facebook? That creates inequality in society”, while on the other hand they say that they are trying to build programmes so that people fund journalism and have subscriptions with different journalistic publications, which people should pay for. On the one



HOUSE OF COMMONS

hand they want Facebook to be free—"How dare you say that it shouldn't be free?"—but on the other hand they say that people should pay for good-quality content.

Maybe it should be the other way around. People should get great journalism for free, because that has a deeper epistemic and positive impact on society, and maybe Facebook should cost money. I say that because the business model will always come up against the kinds of concerns that we are raising. Usually, the thing that is best for people would require a sacrifice in what is best for engagement, growth or time spent looking at the screen.

Q3175 Paul Farrelly: We access social media through hardware that runs on proprietary operating systems. You've made some suggestions as to how the main companies in those fields might respond to and address this problem. Could you explain those to us?

Tristan Harris: Thank you for bringing that up. One thing that everyone should know is that Apple and Google—who make the devices, not the apps—are kind of like the government of the attention economy. The US, UK or EU Governments don't set up the rules for what apps can do or not do, or what business models apps can or cannot have. Apple and Google, who control the iOS App Store and the Play store, and the home screens and the programming interfaces and the rules, are like the government of the attention economy.

Think of it like a city. Apple and Google are the urban planners of that city, and they are funded by slightly different business models. One urban planner, Apple, is paid by you to live in that city. You pay Apple for a phone, and then you live in the Apple city. In the Google city, you pay them by searching for things and asking for things you specifically want, plus spending time in that city looking at things, so they are funded by advertising. Both Apple and Google are the governments or the urban planners of the attention economies that they govern.

Then there are no zoning laws. It is completely unregulated, so you could build casinos everywhere. It is not like you have a residential zone and a casino zone, or a commercial zone. It is that right next to your bedside table, you can build casinos, as loud as possible, and those casinos can be as manipulative as possible and there are no rules.

Apple and Google are in a position to impose, frankly much faster than world Governments, rules that include things like changing the default settings for notifications, or creating different zonal laws for communications applications and standard communication—meaning messenger. Snapchat, Facebook Messenger, iMessage and WeChat could be governed by specific different rules that involve communication, which are different from generic application rules. This is looking into the weeds a little bit, because these are more technical and design discussions, but the point is that they are the ones that are the rule-making bodies. They do not have a strong incentive to change the rules, except when there is large public pressure, as happened recently when Google announced its



HOUSE OF COMMONS

wellbeing initiative at Google I/O about two weeks ago, and hopefully what Apple will announce at WWDC coming up in about a month.

I am not sure if I am answering your question enough, but that is the kind of thing. Potentially, Governments can put pressure on Apple and Google to make those changes more urgently. They could, for example, change taxes. They could tax the advertising-based business model. They could say, if you are going to be an advertising-based business model in our store, you have to pay a certain amount of your revenue out. There are different ways to do that and there would be different issues that people would have with the way that they executed that, because it puts too much power in Apple and Google's hands, but those are the kinds of things that I think are possible.

Paul Farrelly: Rather than delay us here, if you could possibly provide us amateur gardeners with a professional's guide as to how to cope with the weeds, that would be really useful.

Tristan Harris: Sure.

Q3176 **Paul Farrelly:** I presume that, having singled out Apple and Google, you might include other big players like Microsoft and Samsung in that as well.

Tristan Harris: That's right. I am sorry; I didn't meant to single out the others. Microsoft is also in a good position to do this. It is just that when you think about the mobile phone market—I specifically spent my time there—Apple and Google have such a dominant share of that market. But Microsoft and Samsung are also both in positions to do much more around this, and they have different sets of incentives that could be used to push them faster and further in the right direction.

The key thing is also future device factors. Whether that is eyeglasses or watches, these are devices that again are the governments of whatever occurs or shows up, or whatever the rules of competition for attention are. Right now, if you take conspiracy theories on YouTube, maybe there are more rules that say you cannot compete with conspiracy theories.

There are different places where moral authority is a sign that governance is happening. Right now YouTube is being forced to try to strip off some content, but Apple could be putting pressure on YouTube to do more. There are multiple layers of accountability that could happen, depending on how you see the problem.

Q3177 **Paul Farrelly:** You mentioned the App Store and similar sorts of marketplaces. What problems do you see arising from how those sorts of downloads and purchases take place at the moment, and how might they be remedied?

Tristan Harris: This has been changing, a little bit as a result of the awareness that we have been driving. It used to be that app stores sorted what was at the top of the menu. They rewarded the top 10 or top 100 apps based on downloads, ratings and revenue. That produces an



HOUSE OF COMMONS

incentive. Let's say you make it by downloads. Then what you want to do if you are an app developer is manipulate how many people download the product. You create growth-hacking invitation mechanisms that force users to invite 10 of their friends just to use the app, and that gets more downloads, and suddenly you are at the top of the app store. That is not a good way to incentivise them.

Instead, you could incentivise them based on time well spent, which is to say, minimal regret. What are the things that people find lastingly, durably valuable, and net positive when making positive changes in their lives? You would have a totally different set of winners and losers at the top of those app stores. I think Apple has been going a little bit in that direction, and those are some examples.

You could change the rankings; you could change how much Apple and Google tax the business models of the apps they are promoting if they are in an advertising or a freemium model, or if they are manipulating people with virtual currencies—a lot of the apps that make a lot of money right now are virtual games where people have to pay money to put virtual credit into the game, and their parents look at their credit card bill and suddenly it's enormous. Apple can take a moral stance, a normative stance, and tax some of those business models with a preference for things that people find lastingly valuable, and align business so that people pay for a product in a way that is durable—more like subscriptions. The best way to solve the problem is to align incentives. When people pay for the ongoing benefits that they receive in a lasting way, like subscriptions, that is the most aligned business model that we know of right now.

Q3178 Jo Stevens: Thank you, Tristan, for your evidence and for joining us today. I was going to ask you questions about future regulation and what you would be recommending, but you have talked quite a lot about that so far. Picking up on your evidence, you said that we need a ban on micro-targeted advertising. We need legal reclassification of the platforms, whether that is a contract model or some other model, and you mentioned taxing the advertising-based business model. Are there any other recommendations that you would make on future regulation?

Tristan Harris: That is a great question. The first one is a ban on micro-targeting or long-tail personalised advertising—ways that things can be personalised to people although they will not know or realise that they are being personalised to. There is going back to a public form of impersonal advertising.

One thing that needs to happen—we are kind of working on this—is what we call the ledger of harms. It is basically identifying the economic externalities on society that come from these products. Think of it like a pollution tax. You price how much you are externalising harm into the balance sheets, so in the same way that we tax carbon at \$40 per tonne, you can ask how much you are polluting the belief in truth environment. How much are you polluting the mental health environment? Based on the number of users who use your products, and the average cost of mental

health issues that show up—it is about how we start putting a price tag on those externalised systems.

Fundamentally, at a structural level, because of the finite amount of attention there is, I see this as an extraction-based economy. It is a win-lose economy. I am extracting attention out of society—privately for profit and publicly harming. For a society to thrive, certain loops of attention have to exist. A parent has to be able to spend attention on their child; we have to spend attention on ourselves and in conversation with each other. We need to be present with each other. If attention does not go to those places, society would collapse: people do not believe in truth; kids do not get their developmental needs met—all those kinds of things.

Technology companies are pointing their attention extraction machines at human minds, and they are sucking the attention out of all the places where it used to go. In the same way, someone might see a tree as \$10,000 worth of lumber, and they do not see all the complex ways that it is feeding the ecosystem because it is way more profitable simply to cut the tree down than it is to leave it there and realise its value in those other ways. Just like the Environmental Protection Agency, we might want to protect certain forms of attention where it needs to go, and tax places where we are harming those different parts of society.

That is an abstract answer to your question. There are ways of making it more concrete. We need further research, and we need economists to work on trying to name what some of those concrete costs are per platform. For example, YouTube drove 15 billion impressions to a popular conspiracy theory channel. We know that YouTube drove hundreds of millions of viewers to disturbing videos of children with the YouTube Kids app. If you can quantify those views—they took down those accounts and we don't have transparent external access to how many people saw those videos.

Q3179 Jo Stevens: That is the problem, isn't it? Facebook is so secretive and does not have the level of transparency that is needed for the research to be done. In order to do the research you are talking about, wouldn't they have to be forced to co-operate and disclose what is needed?

Tristan Harris: Sorry, I am having problems with the audio quality, but I heard the last bit of what you said, about forcing them to co-operate. It would be great to be able to get these data points out of them and to force them to co-operate with getting that information out. For example, we had to bring them to Congress to know that Facebook had influenced 126 million Americans. We would not have got that number had we not brought them to Congress, because until then they claimed it was only 1 million or 2 million people. Similarly to Brexit, maybe there should be a fine for underestimating some of these costs.

Q3180 Jo Stevens: What do you think the United States will do in terms of regulation? Do you have a feeling about what approach they will take?

Tristan Harris: I don't know. In some ways, we depend on Europe right now to lead the way. As you know, the United States is a little bit of a

circus right now when it comes to governance, and we do not have agreement on the different sides. I think there is a power and antitrust issue, and there could be a case around that. I am not sure, though, that specific things will be done in the US now. In the same way that GDPR originated in Europe and they have led on privacy protections, this might be an area where Europe can continue to lead.

I have a couple of other ideas to run by you. The first is more transparency for researchers. We should simply regulate so that there is more transparent access to trusted research groups. Right now, there is no way for outsiders to know what the living, breathing crime scene, as we like to call it, was with Brexit and the US elections in 2016. Only Facebook and Twitter have the data. In some cases, it has been released after a lot of external pressure, but we should force that access. Until they can fix these problems, which, as we know, are long-term issues, they should fund public awareness campaigns about the costs and harms of these problems and the fact that they cannot control what is true and what people see, in the same way that cigarette companies could be fined to fund the public awareness campaigns around cigarettes.

We need to deal with human impersonation. I do not know whether that should be done through fines, but we need to force companies to anticipate the impersonation of alternative identities. I could start an account called "damiancollins2" right now and respond to a big tweet from the Associated Press. That just happened—I have evidence of it from a week ago. Twitter does nothing to take down these impersonations of accounts. Theoretically, people could thwart them, but they need to have a proactive responsibility, not a reactive one. That is one of the biggest problems: these companies do something, but they usually do it way too late. I don't blame them—it is very hard. They are basically now accountable for impacting more than 180 countries in languages they cannot speak, but they have to be able to dedicate more resources to this. One way or another, we have to force them to dedicate more resources to these problems.

Another thing is reinstituting a lot of the election protections that exist in different countries—equal-price campaign advertising online, for example. In the United States, there was a 17:1 difference in the cost of election ads. Those are the kinds of things that I am worried about and some asks from a regulatory perspective.

Q3181 Giles Watling: A thought occurred to me when listening to Mr Farrelly's questioning about Facebook. We hear about Cambridge Analytica harvesting data from Facebook. I am not on Facebook. Would Facebook be holding data on me?

Tristan Harris: Yes.

Q3182 Giles Watling: And that is from other internet searches I do?

Tristan Harris: Well, it is through, first, the fact that your name is showing up on other people's posts in contexts that you do not control.

They have that data whether they want it or not, and they associate it with other clusters of information whether they want to or not.

Q3183 Giles Watling: I am probably being quite dim, but I do not see what use my data could be to Facebook, if I am not using Facebook. How could they use it?

Tristan Harris: How could they use it? It is a great question. I do not actually know the answer to that, because I do not know their practices. All I can say is that, as we all know, there are trackers on every website that we go to. Data is acquired from lots of different places. Sometimes Facebook acquires it, and there are also third-party independent trackers. You can be identified because your patterns of behaviour show up everywhere. Again, there are ways that other people who are on Facebook might reference you, and they could collect that information and you would not really have a say in that, because it would involve changing the speech of other people who are referencing you. What would it mean to protect your information in that case? That is one of the issues that we have to face.

Q3184 Giles Watling: It is kind of scary. One final thing: Tristan, do you think Mark Zuckerberg should appear before this Committee?

Tristan Harris: I do think he should come before the Committee.

Q3185 Chair: To go back to your call for a ban on micro-targeting users with advertising, in which specific ways would you introduce a policy such as that on to the Facebook platform?

Tristan Harris: One other thing on Zuckerberg coming to your Committee, because I simply said "Yes": after 9/11, the CEOs of all the major airline companies in the United States appeared before Congress. Even though they were just the vehicles that had been used for a terrorist attack, they appeared before Congress. Even if social media platforms are just vehicles for massive influence campaigns that have geopolitical implications, it is clear now that in Brexit and the 2016 elections, there was a psychological attack on two countries in a certain way. Specifically in these circumstances, that would make it justified for Zuckerberg himself to appear, and there is precedent for that. But your question was about micro-targeted advertising.

Q3186 Chair: Yes, what would you do to introduce a policy such as that on to Facebook? Should certain tools be banned or disabled? For example, should there be restrictions on people loading custom audiences on to Facebook to target users individually? Would you ban Facebook Pixel from gathering data about users who visited sites off the platform so they can be targeted just for having visited those sites? I would be interested in your view on the implications of introducing a ban on micro-targeting.

Tristan Harris: It is a really good question. As you said, there are multiple angles. There is the collection of micro-targeting data, the use of it on other websites, the use of it on Facebook—so custom audiences using that data on other websites and the Facebook advertising engine, which could have used it on Facebook—an audience selection process, which is



HOUSE OF COMMONS

how much you can select audiences, and then how much you can computationally generate 60,000 variations of an advertising message to test it on different people. That is part of the micro-targeting process—the fact that I am letting a supercomputer calculate for me the perfect words I should say to you to have you nod your head.

All those things are dangerous in different ways. We have said that one of the challenges is that I am not the person who runs the Facebook advertising team. The unfortunate situation is that the people who most know where the dangers are, are inside the company. The question is, how willing are they to examine where they do not have control over the platform—where there are circumstances where bad things can continue to happen, as they have been happening, in the Irish referendum and other places. They cannot control that.

It is important to note that Facebook decided—they acted to try to reduce some of the things that were happening in the Irish referendum, but Google, on YouTube, realised that it could not solve the problem, so it simply banned all ads about the Irish referendum altogether, because it realised that it was too complicated for it to solve the problem.

There is the model of Tylenol in the 1980s. There was a problem and the company said, “We have to take our product off store shelves until we fix it.” I think that we have to ask ourselves, “What is a safe advertising environment?” This is not about being naively anti-technology or anti-advertising; it is about asking, “Where is it contributing to the best possible social goods for society?” For example, advertising is a great way for small businesses especially to grow. Many small businesses in the world would not exist if it were not for Facebook advertising. I think what we want to do is to find those safe examples—small business advertising and certain types of product advertising or marketplace—find out where this is working really well, and then slowly add back, brick by brick, the safe parts and make sure we are not including the dangerous parts.

Right now, Facebook are taking the entirety of what they are doing and they are trying to chip off some of these edges and say, “Well, we will make our advertising transparent.” That is like saying, “We’re a sugar company and we’re dosing the entire world with sugar, but now we have this programmatic interface, so people can try to figure out where all the sugar is going.” We just do not want it to go to all those places in the first place; we have to be careful about where it is going.

So I think that more than an approach of transparency, which is their current approach, we need to ask, “What are the safe forms of advertising?” Just a total micro-targeting ban might be too aggressive; maybe it does not have to be a total micro-targeting ban, but this is the conversation that needs to be had. Only Facebook have the deep inside information to know what the threats are and where they can build it back safely, but my deep concern is that they are in denial and think that they can retain the majority of their business model and the exponential benefits of long-tail targeting to any keyword, to any person, of any age,

gender or whatever. Instead, they need to take a very conservative approach. That is what I most want them to do.

Q3187 Chair: Talking of that sort of conservative approach, do you think there would be a case for saying that political advertising should not be allowed in the news feed on Facebook?

Tristan Harris: To be honest with you, I would not know whether that would be the thing that matters most, because you could say, "They used to have more ads in the sidebar; now they have ads in the news feed." Obviously, there is the difference on mobile versus on desktop. These are complicated issues. I would love to come back to you when we have more specific ideas about what that could look like.

Q3188 Chair: One of the issues we covered with Mike Schroepfer, when he gave evidence to the Committee, was whether users should have the right to opt out of political advertising. At the moment, they can't stop receiving it.

Tristan Harris: We were discussing earlier whether it should be legal for priests to have the business model of selling access to the confession booth. I kind of think that is one of the situations on Facebook. However, if you take that approach, at the very least there should be such a thing as priests that allow you to pay them, as opposed to the current situation, where the only business model is for priests to sell access to the confession booth. That is kind of where we are now. At the very least, as you are saying, Facebook should have an option for you to pay—to not just not see the ads. That is what people miss about this conversation.

If it is pitched that way—that I am paying so the ads do not show up—people will say, "Well, I don't really mind the ads." In fact, Facebook will say that if you do a study, people actually prefer the version of Facebook—or the version of a newspaper or magazine—that has these big, beautiful, compelling ads in it; it makes it interesting. So the problem is not the advertisement; it is the fact that there is this unregulated automated system that is matching 2 billion people's eyeballs to products that could be from anybody, and that includes divisive or incendiary speech that they cannot control. So what we are paying for is control. It is for safety. We are not paying to remove the ads completely; we are paying for a safer world. I think that, more than that being the voice of the individual, we want to create standards that raise the bar for everybody, and that is the challenge.

I think this is very much like climate change. We have found this business model of advertising that is like coal. Coal is the most efficient way to get lots of energy at a really cheap cost. It propped up our economy for a long time, and that was great; we couldn't have gotten here without it. Now, we need to invent, and we did invent, renewable energy alternatives. That was more complicated. It took more science, more math, more development, but we got the prices to shift. I think the same is true of advertising. Advertising is like coal. It is this cheap business model that was very efficient at generating a ton of money for a lot of people—well, a



HOUSE OF COMMONS

handful of people. It has worked really, really well and it also is polluting the fabric of society because it tears the fabric of truth apart. It corrodes democracy; it corrodes the mental health of children. We now say, "Okay, we're really happy we got to this point where we made all that money from advertising. Now we need to develop an alternative." The alternative will take ingenuity, social science, philosophy, ethics and a much more sensitive view of what it means to impact 2 billion people's lives and the structure of people's culture and society.

I think that is where we are. We know where we came from. We are starting to wake up to how bad that was, and we have not yet invented the alternative business model that is just as economically profitable—and it may not be. That is kind of where we are.

Q3189 Chair: In the short term, a lot of our concern has been about political messaging—of disinformation being spread through political messaging, driven by advertising as the agent to disseminate it. We know that political advertising is a relatively small part of Facebook's business. It is a relatively small part of the ad market. Do you think there would be a case now to say, "If people don't want to receive paid-for political messaging, they should at least be able to opt out of that, even if they can't opt out of advertising on the platform in general."?

Tristan Harris: Were you saying that if people do not want political advertising they should be forced to not use the platform at all, or were you talking about carve-outs? Sorry, I missed—

Q3190 Chair: No. Say you do not want to receive political advertising—let's say you are a user who is worried about Russian fake news and other things that you might get—and you say, "I want to opt of that. I want to change my settings so that no one can target me or send me political advertising." At the moment, you cannot do that, because Facebook can still sell you to an advertiser, and you cannot stop receiving it. Given that a lot of focus has been on this sort of disinformation, spread through advertising on Facebook, and the value of this is a relatively small part of Facebook's advertising revenue, should we now, as a kind of first step, say, "Let's do something about that, and at least let people have the option of turning that feature off if they don't want to receive it."

Tristan Harris: Instead of just giving you the option of it, why not just make it the default until we know that it is safe? Let's just turn political advertising off altogether. That is where I would stand. I would say, though, that the challenge is that what Russia did, for example, was less about messaging to political keywords and instead just using the keywords that certain kinds of divisive groups would already use.

If I want to magnify conspiracy theories in your population, all I have to do is use completely non-political keywords, which are just the words that that conspiracy theory group use. I can just use the word "vaccine", or something like that, and I suddenly have access to the anti-vaccines or "Warrior Moms". I can access a group and start to stoke them in a certain direction. That is not political. One of the challenges is that this issue



HOUSE OF COMMONS

defies the normal labels and boundaries that would be convenient. It is exponentially complex territory in terms of the diversity of things that can be targeted against.

That is why I am saying that I think we have to be really careful about which areas of that landscape we want the switches to be turned on. You can advertise, for example, products and merchandise. Do we really need to enable anyone to simply put messages into people's minds? Maybe there are safer ways to do it. I am hoping so.

Again, I think people inside Facebook have the most knowledge about what might be safe and what is not safe. That is one of the issues: transparency. From the outside, it is hard to make recommendations that they would not roll their eyes at. They would be right to roll their eyes at me, because I do not have access to the data. On the other hand, they are in denial because they want their business model. I think that is the real issue that we are in.

Chair: That covers our questions. It has been really great talking to you again, and a really interesting session. Thank you very much.