

Additional Questions for the Record

Subcommittee on Consumer Protection and Commerce

Hearing on

“Americans at Risk: Manipulation and Deception in the Digital Age”

January 8, 2020

Mr. Tristan Harris, Executive Director, Center for Humane Technology

The Honorable Kathy Castor (D-FL)

1. What techniques do internet platforms employ to keep their users engaged? What types of techniques are harmful? What types of techniques are harmless? Are there industry standards? If so, what are they? If not, what should they be?

Techniques:

- *Slot machine rewards (“pull to refresh”)*
- *Loot boxes in games*
- *Removal of stopping cues (Infinite auto-backfilling news feeds, auto play on countdown)*
- *Designing to maximize tightly interconnected “infinite worlds” (never-ending, tightly interlinked worlds of content, like on Twitter, where links to more and more content never end)*
- *Personalized recommendation systems and algorithms that know us better than we know ourselves (e.g. YouTube “Up Next” recommendations, Instagram or TikTok News Feeds). They are supercomputers pointed at our brains making increasingly accurate predictions about which posts or videos will keep each brain most engaged.*
- *Social tagging and puppeteering (“You’ve been tagged in a photo”, “X endorsed you, endorse them back?”, “Say hi with a wave!”, “Suggested users to follow”)*
- *Attention-seeking hacking (# of followers, # of likes, # subscribers, etc.)*
- *“Come back” emails to “resurrect” idle or dormant users (tech as “digital drug lords”)*
- *Beautification Filters that affect teenager’s self-worth and identity through positive social feedback limited to unrealistic standards of beauty. (“people like you, if only you look different than you actually do”)*

There are no industry standards on what practices are allowed or not. The “race to the bottom of the brain stem” to strip mine human attention knows no boundaries, except when there’s rare surges of public pressure that occasionally have them limit certain features – such as Instagram testing the removal of “Likes” in certain markets.

What should those standards be?

- *No auto playing videos in feeds*
- *No automated countdowns for the next video, by default*

- *No bottomless, infinite scrolling feeds without pauses, “speed bumps” or self-imposed random friction and slowdowns – either in the form of a “Load more” button or a breathing gap*
- *Consider limiting people under the age of 18 to 90 minutes of playing time on weekdays and three hours on weekends and holidays – as they do in China.*
 - <https://mashable.com/article/china-video-game-regulations-minors-under-18.amp>
- *No loot boxes. This mirrors policies already passed in Japan (see <https://www.adweek.com/digital/japan-officially-declares-lucrative-kompu-gacha-practice-illegal-in-social-games/amp/> <https://www.lexology.com/library/detail.aspx?g=9207df10-a8a2-4f67-81c3-6a148a6100e2>)*
- *Ban time-spent and engagement-maximizing business models and algorithms*

2. What techniques do internet platforms employ to manipulate their users? What types of techniques are harmful? What types of techniques are harmless? Are there industry standards? If so, what are they? If not, what should they be?

- *Micro-targeting unchecked, unregulated computationally generated messages and images of vitriolic content with no accountability*

All of the above techniques can be used in ways that lead to harms -- addiction, polarization (over personalization), distraction, depression and suicide, bullying, vanity culture, loneliness, micro-targeted disinformation, seeding conspiracies and extremism.

There are no common industry standards or practices to protect against these harms. The industry should be held accountable to a new body, a Attention Economy Agency whose job is to monitor for the standardized set of harms that must be reduced or eliminated from technology, create public pressure, and set quarterly targets for harm-reduction across these areas from participating tech companies.

New standards:

- *Ban micro-targeting*
- *Ban recommendation systems, which must be oriented around positive values, not engagement*
- *Ban time-spent and engagement-maximizing business models and algorithms*

3. The word algorithm is used frequently in discussions over internet engagement. What is an algorithm? Who designs the algorithms? What are the benefits/harms to internet platforms using algorithms? How do algorithms use artificial intelligence? Can internet platforms fully explain why an algorithm produces certain results? Do internet platforms have knowledge of all the information fed into the algorithms they Use?

An algorithm is, “a process or set of rules to be followed in calculations or other problem-solving operations” according to Oxford Internet Dictionary. Algorithms produce automated

decisions – and power everything from what content you see in news feeds, what order it appears in, which ads you see, when notifications on your phone get delivered, to which Uber car gets notified when you hit “Request Ride.” Some algorithms are backed by artificial intelligence, which are more advanced algorithms where the machine “learns” its own strategies and rules to most efficiently reach specified goals or outcomes.

It is impossible for internet platforms to fully explain why an algorithm produces certain results. Recommendation and news feed algorithms powering services like Facebook or YouTube often use millions of variables to determine what we see – let alone the fact that there are trillions of possible combinations or orderings of content they could show.

There is also a direct tradeoff between what’s called “algorithmic explainability” – the degree to which an algorithm can be explained, and how accurate it is at producing outcomes or predictions. An algorithm that uses a trillion invisible parameters is harder to explain, but more accurate at producing the intended goals, than an algorithm with 10 parameters that is easier to explain, but less accurate at producing the intended goals.

4. What should be considered healthy engagement with an internet platform? Is healthy engagement defined differently for children? If so, what should be considered healthy engagement with an internet platform for children?

It’s tempting to define healthy use as simply spending less time on today’s existing internet platforms like YouTube, Facebook, etc. However, that is deeply insufficient. That would be like a government advocating that every citizen have a “healthy daily habit” of cigarettes and limit it to three per day. We can’t recommend “healthy use” of cigarettes when they are designed to be addictive, just like we cannot recommend “Healthy engagement” when commercial interests from asymmetrically powerful systems know how to manipulate each user and externalize cultural harms that affect everyone. So long as platforms’ business models are based on extracting and mining attention at all costs, and billion-dollar profits depend on it, their design decisions won’t have the best interests of society, or stakeholders in our society, in mind.

We shouldn’t be aiming for “healthy engagement” as much as we should aim for the notion of “humane technology” that does not asymmetrically override and manipulate users for the best interests of business. Humane technology does not prey on human vulnerabilities for commercial interest.

A phone app that lets children make audio and FaceTime calls to each other, for example, is not manipulating children with news feeds, beautification filters, social feedback and ratings (e.g. Likes and Followers). A plain telephone is like a tool, like a hammer, waiting patiently to be used. Tools or hammers aren’t harmful. Engagement-maximizing services, driven by the corrosive, unbounded business model of advertising that seeks to consume and extract human attention, is what is existentially harmful. Given their business model, a human being is simply worth more if they are addicted, isolated, outraged, narcissistic, voyeuristic, polarized and disformed (because of ads), than if they are a sovereign human being.

5. When does engagement with an internet platform turn into addiction? How are classic

signs of addiction measured in the digital context? Is the addiction connected to internet platforms similar to manifestations of addiction in other situations? How is it similar? How is it different?

Defining a clinical threshold for addiction is a distraction from the existential issues at stake. Tech companies will happily stall that debate for decades as governments enter never-ending debates about clinical thresholds and measurements everyone can agree on. What matters is that business models of asymmetrically powerful, and ever-growingly more powerful tech companies, who know more about each user than they know about themselves, have a runaway unbounded incentive to exploit that asymmetric understanding of your vulnerabilities.

Imagine a patient sitting unconscious in the operating room of a doctor. The patient's life is in their hands. They are vulnerable to whatever the doctor chooses to do. The patient also trusts the doctor who knows more about medicine and their personal information and differential vulnerabilities than they know about themselves, not to exploit their vulnerability as they lie on the operating table, unconscious. We have to redefine the asymmetric relationship between users and tech companies into a "fiduciary" relationship. That means technology platforms cannot be allowed to operate with business models based on exploiting their users and the societies in which they operate.

6. Why are repeat engagement with or addiction to an internet platform harmful to the individual or society as a whole? What are the costs? For example, what are the economic costs?

Addiction is the least of the harms. Platforms that create individualized, addiction-maximizing "Truman shows" with never-ending AFFIRMATION instead of INFORMATION, due to the business model that reinforces what keeps each of us clicking, destroys shared truth and facts, drives polarization, and is existential to the function of democracy and policy-making.

*If no one agrees on what any given policy is, or what it will do, then no consensus can exist. For existential issues like climate change that can *only* be addressed by passing policies that bind market forces to reduce and reverse emissions, that's game over for all of humanity.*

7. Why are manipulative techniques employed by internet platforms harmful to the individual or society as a whole? What are the costs? For example, what are the economic costs?

The complete breakdown of democracies all around the world. The inability of humanity to address climate change, immigration, democracy deficit, fake news and many other important issues.

8. How do internet platforms monetize repeat engagement or addiction? Why does this model benefit internet platforms? What are the benefits?

Because their business model is based on showing as many ads as possible to users, they make money the more time they can manipulate users into coming back and spending ever more time

on their platform to see more ads, which allows the platforms to make even better predictions with their virtual models of each user. This business model has made these advertising-based technology platforms the most profitable corporations in history.

9. How do internet platforms monetize manipulation? Why does this model benefit internet platforms? What are the benefits?

*Manipulation is the very basis of each design decision of advertising-backed platforms – not because of the *advertisement* as manipulation, but because each design decision surrounding the delivery of the ads is based on manipulating the user and strip-mining their attention in order to serve the user up to advertisers. Follow the money, and you will understand how and why they manipulate users.*

10. What tools are at internet platform users' disposal to stop repeat engagement or addiction? Should companies provide or fund those tools?

They lack sufficient tools. "Screen Time" features like on iPhone only allow you a theoretical limit. Research suggests this hardly works, and very few people opt into these tools. It is equivalent to telling someone who is a smoker, "well you shouldn't keep smoking, just limit yourself to two cigarettes a day." A product designed with scientifically precise addictive capacity cannot be dealt with through serving size limits. The better solution is to re-align the business model and incentives of technology to align with the interests of users and society. In other words, acting like "attention utilities" who, like any public utility, must be regulated to operate for the public interest and come with certain safety standards.

11. What tools are at internet platform users' disposal to stop manipulation? Should companies provide or fund those tools?

They do not have access to tools to stop manipulation. They cannot opt out of micro-targeting or advertising or the ways that the news feed content, groups, events, etc., themselves can be manipulated by nefarious actors. If they did that, it would kill the goose that lays the golden eggs, from the company point of view. Manipulation is core to these platforms' money-making machines; it is core to their business model. Regulation and rules from the government are badly needed, because these companies cannot and will not regulate themselves.

12. What role should Congress play in combating repeat engagement with or addiction to internet platforms?

Congress is way behind the need to put rules and regulations around the business models of these companies. These companies have been able to establish their own rules for many years now, and as a result have created a business model that is destructive to children, families and society at large. I can't emphasize this enough – it is the BUSINESS MODEL of these companies that is the problem, not the technologies themselves. These companies use engagement and manipulation tricks because they are trying to keep us glued to their websites, because the longer we are glued, the more ads we see and the more money they make. So, you have to fundamentally break the connection between the money-making machine and the engagement.

There are several ways to do that. You could limit or even forbid advertising on these platforms. Just as the Highway Beautification Act of 1965 forbids the construction or showing of billboards on almost every piece of private property in the US, digital platform companies would not construct or show any advertisements on a user's web page on that platform. The digital platform may, however, provide a link to an advertising page that any viewer may visit if she/he wishes.

Alternatively, Congress could use legislative encouragement to push these platforms towards a subscription model. We should closely examine whether a business model based on free services in exchange for our personal data, which then is harvested for advertising, persuasion and the spread of disinformation, is the appropriate one for these platforms which have essentially created the "new digital infrastructure" for the 21st century. A subscription model, like the one used for cable TV in which companies charge a monthly fee and must adhere to a digital license of conditions, would be a better match for this crucial infrastructure sector. European Commission Vice President Margrethe Vestager, who is also the Commissioner on Competition for the EU, has called for a Facebook based on a subscription with "no tracking and advertising and the full benefits of privacy." Keep in mind a paid, subscription version of Facebook wouldn't just be the same service we get today – without the ads. It would be an entirely different kind of service built around helping us get the most out of our lives with friends, like Mark Zuckerberg's original 2005 description of Facebook as a "social utility" to help us connect with our friends.

13. Should Congress fund more research studying the techniques utilized by internet platforms to increase engagement and manipulate users and their effects? Should some of that research focus on the effect techniques utilized by internet platforms to increase engagement and manipulate users have on children?

*Yes, Congress should fund more research, not only concerning children but also the impact of these techniques on polarization, fake news, democratic discourse, elections and more. Internet platforms represent an existential threat to our democracy. But the need for ongoing research should not be an excuse for inaction. And research aiming to study the existing harms would simply take too long to make the necessary changes. Research would provide us with a more accurate record of how the social fabric, truth and mental health melted in front of our eyes, instead of funding research immediately into new platforms that would reverse and prevent this process. There already have been a number of studies done about the techniques utilized by Internet platforms, at this point the effects are well-known among researchers. See the book *Surveillance Capitalism*, which references a number of studies on the impacts of these Internet platforms and their toxic methods.*

In fact, these platforms and their techniques should be subject to ongoing "Attention Impact Analyses" (AIA). -- akin to an environmental impact analysis – which assess potential impacts on fake news and info-sharing, democracy, social polarization and mental health before deployment of new techniques. An AIA should apply a "precautionary principle" – a kind of Hippocratic oath of "first, do no harm" – to their business model. This would put society on a healthier footing for the safe use and enjoyment of these technologies.

In that light, the US badly needs to create a watchdog Attention Economy Agency (AEA). Like the Environmental Protection Agency was created in 1970 to oversee and consolidate the watchdog function of protecting the environment, the Attention Economy Agency would play a regulatory and watchdog role for this sector that is increasingly becoming central to every aspect of our economy, culture and society. This AEA also would facilitate other federal agencies engaging in a “digital update” and “harms audit” of how the new attention economy is impacting the legal and regulatory frameworks under its purview.

For example, there are restrictions on violence and advertising for Saturday morning cartoons and other programming for children, resulting from laws like the Children’s Television Act passed in 1990. Yet Google’s YouTube/YouTubeKids violates these regulations and norms of decency on a regular basis. The Federal Communications Commission should examine how to apply existing law to the online digital platforms. Other federal agencies, as well as state governments, should do the same. The AEA would help facilitate this kind of re-examination and update.

14. Section 5 of the FTC Act prohibits “unfair or deceptive acts or practices in or affecting commerce.” What types of manipulation should be considered unfair or deceptive? What types should not be considered unfair or deceptive? Should a different standard be developed for manipulative techniques used by internet platforms? If so, what should that standard be? What manipulative techniques should be allowed for adults but not for children?

Their entire business model is “unfair and deceptive.” How many people realize that every app on their cell phone, every platform company they interact with, is tracking them, grabbing their data, monitoring their email, and generally using powerful methods of surveillance and manipulation of their every next move that would have made the Nazis or the Stasi envious? And then they use that information to build psychographic profiles on each and every one of us, and then advertisers pay for access to those profiles. These platforms are pointing powerful supercomputers at our brains in order to more perfectly predict, more and more successfully, how to get us to think, feel and do things – click on the link, agree with a personalized micro-targeted message, read a popular but fake or extreme news article (because the algorithms cannot distinguish between “popular” and true). On and on and on. This entire practice is “unfair and deceptive.”

Sure, we could break it down more minutely, look at which forms of engagement, recommendations, micro-targeting, auto feed, social tagging and more are the most destructive. But by singling out any one or two or three of these techniques, we missed the 800-pound gorilla in the room – the entire business model is unfair and deceptive.

15. Does the application of section 230 of the Communications Decency Act (section 230) enable increased manipulation and repeat engagement/addiction? If so, how does section 230 enable increased manipulation and repeat engagement/addiction and what are the potential fixes?

Yes, undoubtedly. Section 230 has made it so that these platforms are not legally or financially responsible for any of the content that their billions of users are spreading with the assistance of these platforms' global reach. Because they are not legally or financially responsible, they make very few efforts to stop bad actors. Look at Facebook, which has violated multiple consent decrees agreed to with the FTC and the Department of Justice, showing that even when it agreed to certain rules, it broke them because they were not afraid of strong enforcement. The New York Times ran a series about how live streams on Facebook, as well as on other platforms like Zoom, are being used by bad actors to live stream outrageous activities such as the Christchurch killer who live streamed his mass murder of Muslims in New Zealand; or Zoom and Google, being used to live stream sexual child abuse in real time, with participants observing and shouting encouragement to the abusers; or Buddhist extremists in Myanmar who used Facebook to whip up anti-Muslim hysteria against the Rohingya minority. This is horrific and when confronted about it most often the platform companies try to evade both action as well as responsibility. They get away with this because of section 230.

The harsh reality that everyone needs to come to grips with is that the same technology that is used for posting our children or puppy photos, or finding an old college roommate, or for live streaming important meetings, is also being used to do terrible things. Are we supposed to simply throw up our hands and say, "Oh well, that's the price for being able to post cute puppy photos?" Or can we find a way to regulate these extremely powerful technologies?

Section 230 became law in the mid-1990s, when all of these platforms were still rather small. It seemed like a good idea at the time to encourage the growth of the Internet. Now, we see That section 230 is part of an "Attention economy" ecosystem that is harming individuals, families, communities(Like the Rohingya in Myanmar), societies in general, our democracies. Section 230 needs to be modified in such a way as to rein in the harms and retain the good. This can be done, following some of the proposals we have recommended above.

The Honorable Lisa Blunt Rochester (D-DE)

1. At the January 8, 2020 hearing, you indicated that you were familiar with the concept of universal design. Do you think online service providers, like Facebook, should follow universal design concepts as a best practice on all of their platforms?

2. I am concerned that sometimes our laws and regulations are too reactive and do not anticipate future developments in technology and their societal impacts. Frankly, it often seems that technology has outpaced people and policy. We need to be more proactive. How do you think we in Congress can develop a more agile and effective response to these concerning trends on the internet?

New technology frequently drives change that exceeds that capacity of government to respond. The nature of this evolutionary process is that government is always somewhat behind the curve. But the more exponentially powerful technology becomes at shaping more and more consequences – controlling what billions of people believe, think and do – the less we can afford any errors. Technology has made us too powerful to be negligent. Technology is the new infrastructure for society. At a certain point, governments must step in to create the "guardrails"

for this new infrastructure. Like in 1982, when AT&T controlled 90% of the nation's telephone market, it was necessary to break up that company. Like our utilities - the US devised a regulatory structure in which they are given a quasi-monopoly, but in return they must agree to a number of "duty of care" rules and regulations.

We now find ourselves in a similar situation with digital technologies. The digital platform companies have created the new digital infrastructure for the 21st century. So now we must begin the crucial process of crafting the rules, regulations and "digital licenses" that these companies must abide by in order to operate. To do this the right way, we can look to our successful past. The Environmental Protection Agency was created in 1970 to oversee and consolidate the watchdog function of protecting the environment. 1977 saw the creation of the Department of Energy because we recognize that the landscape and the technology around energy was becoming sufficiently complex that we needed a federal agency that focused on that portfolio.

We have reached a similar moment now with the digital technologies and the "attention economy." We need to create a watchdog Attention Economy Agency to oversee a new classification for these companies as "attention utilities." The Attention Economy Agency would play a regulatory and watchdog role for this new sector that is increasingly becoming central to every aspect of our economy, culture and society. This AEA also would facilitate other federal agencies engaging in a "digital update" and "harms audit" of how the new attention economy is impacting the legal and regulatory frameworks under its purview. The EPA and the DOE functioned in that way as well, when they were first created, helping other federal agencies to incorporate that perspective into their own portfolios.

For example, there are restrictions on violence and advertising for Saturday morning cartoons and other programming for children, resulting from laws like the Children's Television Act passed in 1990. Yet Google's YouTube/YouTubeKids violates these regulations and norms of decency on a regular basis. The Federal Communications Commission should examine how to apply existing law to the online digital platforms. Other federal agencies, as well as state governments, should do the same. The AEA would help facilitate this kind of re-examination and update. And this agency would make sure that the government does not fall too far behind in its need to keep up with the fast-changing digital technologies.

The Honorable Robin L. Kelly (D-IL)

1. When discussing deception online, one thing in particular that concerns me is phishing schemes used to obtain individuals' personally identifiable information, such as credit cards and health data. This is of particular concern for vulnerable populations such as those for whom English is a second language, as well as seniors or other populations with lower tech literacy rates. What opportunities exist to provide consumers with the confidence that the website requesting their information is legitimate? Twitter has the blue check mark. Is there a way to provide an equally easy and identifiable verification of websites requesting PII?

*This question - how to protect vulnerable populations and their PII - illustrates exactly why we need privacy (and other ethical and humane aspects of technology) by design, not by requiring individuals make burdensome, heavily-researched choices in every moment. The world is getting more and more complicated. Everyone is busy. No one has time to research the basis of every supply chain or every button they push, especially when technology bombards them with more and more communication and choices. Users should be able to trust that technology has their best interests in mind – and have platforms proactively protect users. **The default should be settings that we can trust.***

Today's tech platforms are like a car that only goes 0 mph or 100 mph. When it gets into a metaphorical crash – mental health problems like anxiety or loneliness, or political problems like polarization – then, technology companies blame the driver. But it's not a driver issue, it's a design issue. Cars should be safe to drive at most speeds. Technology platforms should be safe and harm-minimizing for all of us, most of the time, by default.

As I wrote previously:

I believe in a world where the technology industry is remade in a manner that becomes a more empowering tool -- something that serves humanity and life again. Where it is built around servicing our needs and strengthening the fabric of our society, not parasitically extracting value from the most vulnerable organs of society. Where technology strengthens our capacity to see multiple perspectives, nuance and complexity – where there are no black and white answers.

We need technology to aid us in these endeavors for our civilization to survive.