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. Justin (Gus) Hurwitz, Associate Professor of Law, Director of the NU Governance and Technology Center, University of Nebraska College of Law
Director of Law & Economics Programs, International Center for Law & Economics, McCollum Hall

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?

This is an exceptionally broad question to which there are no simple answers. Internet platforms, as a category, are as varied as any market or market actors. In general, Internet platforms keep their users engaged by providing those users with content, goods, services, or other things that their users find value in consuming.

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?

This, too, is an exceptionally broad question to which there are no simple answers. Indeed, what "manipulates" one user may help another. Short of cases of abject fraud, it is difficult to call any given conduct manipulative. Indeed, even the effects of certain fraud may be relative, and beneficial to some individuals or in some contexts. It is well within the range of common human experience that lies, in certain cases, can be useful or appreciated. This same observation is true even for attempted manipulations that fall short of fraud. Consider the framing of this question, the phrasing of which presupposes the deliberate use of manipulation by Internet platforms and is designed to elicit responses that support that presupposition. Such manipulations are commonplace and largely accepted in civil society – both because their effects are generally innocuous and because it is exceptionally difficult (arguably impossible) to ask a question (or design an interface) that is not based on some assumptions and therefore will not "manipulate" some individuals.

Looking to existing legal standards, tools such as the FTC's authority to proscribe deceptive conduct are illustrative of tools for identifying harmful conduct that we may think of as manipulative. The key element in a deception inquiry is whether the conduct *materially* contributed to the harm.

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?

There is no accepted formal definition of algorithm. The meaning of the term is the subject of vigorous academic research and debate. At a general level, the term roughly means nothing more than a structured process of doing something. The "algorithm" to start many cars is "insert the key into the ignition; press the brake pedal and hold it in the pressed position; turn the key to the start position and hold it there until the engine has started; turn the key to the run position." Elementary school students learn "algorithms" for long division or calculating square roots. Computers use algorithms to convert a sound wave into an audio file, or to display an image file onto a screen.

Given what an algorithm is, asking "what are the benefits/harms to internet platforms using algorithms" is synonymous with "what are the benefits/harms to internet platforms existing." Algorithms are merely the instructions that tell the computers on which the platforms operate how to carry any – or, literally every – thing that they do.

Algorithms do not use artificial intelligence. Artificial intelligence uses algorithms to identify patterns or correlations in data and, in turn uses those patterns or correlations as an input into algorithms.

It is often the case that computer engineers cannot explain, or cannot easily, explain, the behavior of algorithms. This is the case both with many complex algorithms designed entirely by computer engineers or by algorithms that rely on machine leaning ("artificial intelligence") as an input into the operation of algorithms.

If it were easy to fully understand how an algorithm works, computer software would not have bugs. As anyone who has ever used a computer knows, all software has bugs. This is not because computer engineers are lazy or incompetent. It's because designing and implementing algorithms is exceptionally, incomparably, incomprehensibly, difficult to do. One of the first concepts that any computer scientists learns is the Halting Problem, which, in essence, states that it is possible to prove that *any* algorithm beyond a trivial level of complexity *can not* be fully understood without devoting an impossibly large amount of resources to it. Fully characterizing the algorithmic behavior of even the simplest of modern computer programs would take modern supercomputers a period of time longer than the Universe has existed.

The use of modern computers and algorithms is, in a sense, always a calculated risk – albeit one where the benefits generally outweigh the risks by thousands of orders magnitude. Any efforts to regulate based upon "algortihms" will be as effective as simplifying math by legislatively defining π to equal 3. Rather, regulation should focus on the effects of algorithms, not their design or inputs.

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?

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?

Comprehensive answers to questions 4 and 5 are outside of my areas of expertise, beyond general familiarity of the work of various individual researchers who do work in these areas. It is my general understanding that these are contentious issues subject to vigorous debate among experts in the field.

The following newspaper clippings, archived at https://twitter.com/PessimistsArc, however, provide useful cautionary context for approaching these discussions:

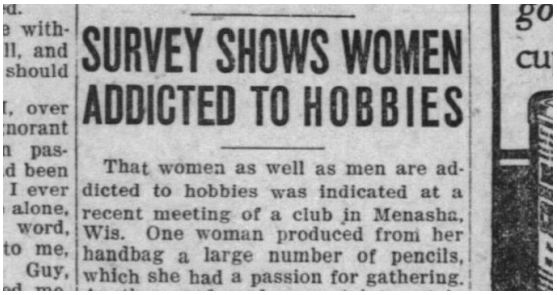

SURVEY SHOWS WOMEN ADDICTED TO HOBBIES

That women as well as men are addicted to hobbies was indicated at a recent meeting of a club in Menasha, Wis. One woman produced from her handbag a large number of pencils, which she had a passion for gathering.

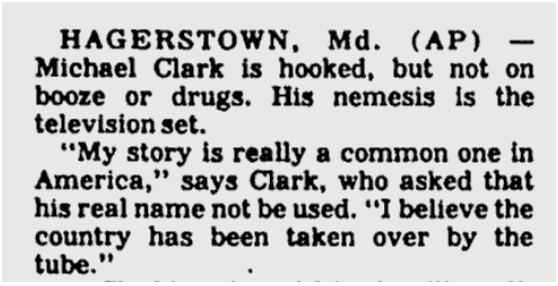

HAGERSTOWN, Md. (AP) — Michael Clark is hooked, but not on booze or drugs. His nemesis is the television set.

"My story is really a common one in America," says Clark, who asked that his real name not be used. "I believe the country has been taken over by the tube."

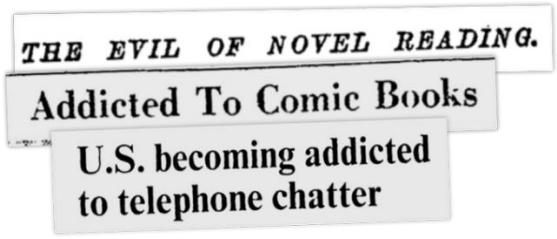

THE EVIL OF NOVEL READING.

Addicted To Comic Books

U.S. becoming addicted to telephone chatter

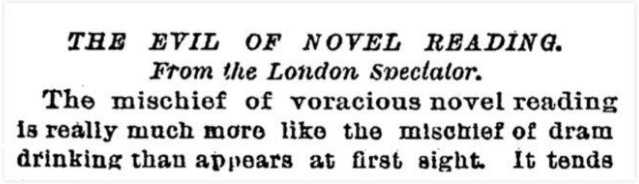

THE EVIL OF NOVEL READING.

From the London Spectator.

The mischief of voracious novel reading is really much more like the mischief of dram drinking than appears at first sight. It tends

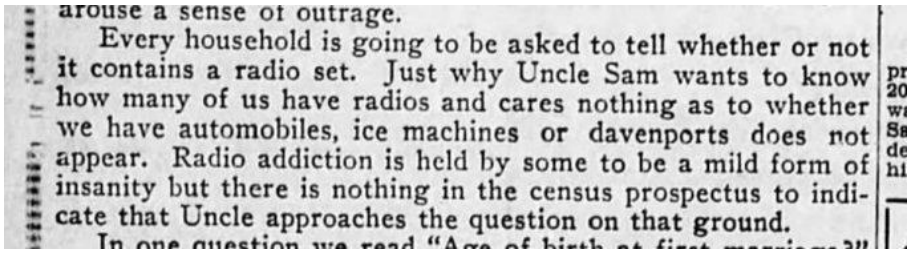

arouse a sense of outrage.

Every household is going to be asked to tell whether or not it contains a radio set. Just why Uncle Sam wants to know how many of us have radios and cares nothing as to whether we have automobiles, ice machines or davenports does not appear. Radio addiction is held by some to be a mild form of insanity but there is nothing in the census prospectus to indicate that Uncle approaches the question on that ground.

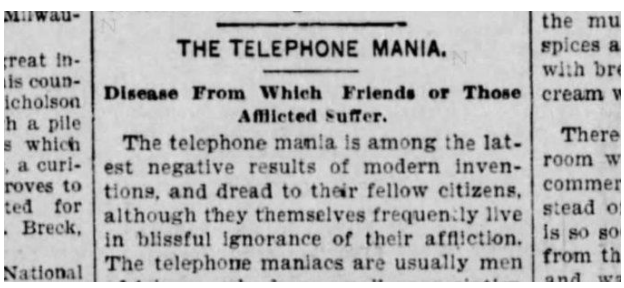

THE TELEPHONE MANIA.

Disease From Which Friends or Those Afflicted Suffer.

The telephone mania is among the latest negative results of modern inventions, and dread to their fellow citizens, although they themselves frequently live in blissful ignorance of their affliction. The telephone maniacs are usually men

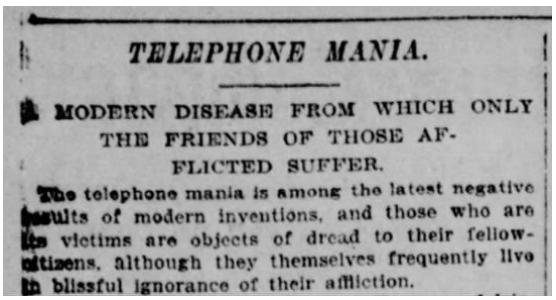

TELEPHONE MANIA.

A MODERN DISEASE FROM WHICH ONLY THE FRIENDS OF THOSE AFFLICTED SUFFER.

The telephone mania is among the latest negative results of modern inventions, and those who are its victims are objects of dread to their fellow-citizens, although they themselves frequently live in blissful ignorance of their affliction.

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?

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?

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

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

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

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

In general, the response to questions 6 through 11 is that overwhelmingly Internet platforms approach all of these issues largely in the same way as other businesses, technologies, and platforms have approached them in the past. There are technological and economic differences between all of these platforms – to the extent that there are meaningful differences between Internet platforms and past businesses, technologies, and platforms, it is unclear whether these effects ultimately militate for or against the need for regulatory intervention. Just as platforms may have some greater ability to act in ways that are ultimately harmful (or beneficial) to consumers, consumers or competitors may similarly have greater ability in the modern technological era to protect themselves from or take action against such potentially harmful conduct.

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

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?

In response to questions 12 and 13, any Congressional response to the concerns evinced above should be based in empirical assessment of effects on consumers that compare the relative costs and benefits to consumers to plausible counterfactual worlds. Congress should absolutely fund, or encourage funding of, significant research into these areas. Importantly, the framing of that research is important – lest Congress itself be engaged in the practice of dark patterns. For instance, soliciting research "to study techniques used by Internet platforms to manipulate their users and the effects of these manipulations" *will* produce results that find that platforms do

manipulate their users and that effects of these manipulations are adverse to users' interests. Such a study – and the funding behind it – would be political legerdemain (or, since relatively transparent politicking, merely manipulation)

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?

The FTC's Section 5 authority is an exceptional model for how to approach these issues. Critical to its deception authority, in particular, is the requirement that for any conduct to be deceptive it have a material adverse effect on consumers.

Questions about differential regulation of platforms for adults and children is an exceptionally difficult subject. As the Supreme Court unanimous said in *Reno v. ACLU*, "the Government may not reduce the adult population to only what is fit for children." In general, the First Amendment requires that we not limit a forum intended broadly for use by adults to content and form suitable for children.

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?

Section 230 has only the barest and most attenuated of relevance to these issues. Section 230 has two functions: it immunizes platforms from liability for content created by its users, and it allows platforms to moderate that content, if they so elect, without assuming liability for content created by its users. To the extent that Congress is concerned about the conduct of Internet platforms, the first provision is wholly irrelevant. To the extent that Congress is concerned about users of platforms engaging in problematic conduct, the second provision facilitates action by platforms to curtail that problematic conduct. Both of these provisions are fundamentally sympathetic to any concern Congress may have, and supportive of Congressional efforts to curtail problematic conduct.

This is not to say that Section 230 is perfect. I have proposed narrow revisions to the statute that would bring it more into line with todays technological realities by enabling individuals engaging in harmful conduct to be identified, and subject to legal process, by parties harmed by their conduct. But proposals such as this are narrow and designed not to disrupt the fundamental operation of an exceptionally important law that has been overwhelmingly beneficial to American consumers and industry alike.

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?

Universal design is a laudable design goal and valuable is a principle that should be incorporated into design practices. However, it cannot be – as a technical matter – be reduced to or implemented as a design requirement. Requiring any firm to abide by the principles of universal design under penalty of law is tantamount to requiring that firm to successfully balance – in effect solve – all of society's competing tradeoffs and to face legal sanction for failure to do so. Researchers have studies questions related to these issues for hundreds of years and widely understand that there is no single, stable, equilibrium that maximizes the myriad competing values required by the principles of universal design.

This is not to say that universal design is not a good idea. Rather, it is necessarily aspirational. Firms that abide by it should be lauded and rewarded in the marketplace. Congress may choose to require firms to follow specific aspects of universal design, or, more likely, to meet narrow prescriptive goals required by those aspects under certain conditions. But the idea of requiring a firm to follow universal design concepts cannot be reduced to enforceable law – and any effort to do so would be a textbook example of a law that was unconstitutionally void.

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?

This is a generational effort – and not one that Congress can address on its own. There is no simple answer to this question. My recommendation is to create (that is, fund) more opportunities for interdisciplinary engagement between the fields of law, business, and engineering. To whit, I am currently in the process of establishing a new center at the University of Nebraska, the Nebraska Governance and Technology Center, that does precisely this.

In general, Congress always has been and always will be reactive to technological change. That is the nature of technology. The solution is not figure out how to bring greater technological knowledge into Congress. By the time problems created by any new technology reach the level of Congressional attention it will be too late for Congress to be anything but reactive -- the horse will have already left the barn. Rather than bring greater understanding of technology into the legislative and policy process, we need to bring greater understanding of the legislative and policy process into the engineering and business sides of technology development.

The Honorable Brett Guthrie (R-KY)

1. There is clearly a spectrum of business practices as it pertains to influencing consumer choices. On one end, these practices are legitimate and on the other, such practices have the potential to harm consumers. Professor Hurwitz, how do we draw the line between

legitimate business behavior designed to influence users and exploitative "dark pattern" interfaces that may harm consumers?

The best approach is to focus on the effects of these practices on consumers, not on the practices themselves. This approach has long been central to the Federal Trade Commission's authority to proscribe "deceptive" acts or practices. In order to be deceptive, an act or practice needs to have a *material* effect on consumers. In other words, and without getting into the details of FTC regulations, consumers need to be harmed because of the act or practice. It is not enough that an act or practice *could* conceivable cause harm – there needs to be some causal relationship between the conduct and actual harm. This requirement, of course, is not unique to consumer protection law – causation is a basic element of most areas of law.

It is often the case that design practices that appear likely to harm consumers have little, or even beneficial effects – and, conversely, that design practices that appear likely to benefit consumers may actually harm them. For instance, it is widely believed that supermarkets stock their checkout lines with "impulse" purchases that consumers are unlikely to buy unless "tricked" into buying them. While there is likely some truth to this, supermarket layout is intensively studied. Supermarkets often stock their checkout lanes either with products that consumers are likely to have forget to put in their carts (a benefit to consumers) or with curiosities that consumers are likely to engage with but not buy (e.g., tabloid magazines), which improves the customer experience. Or consider "ban-the-box" legislation, intended to give individuals with criminal records a better chance at getting jobs by preventing employers from asking them to indicate (check "the box") whether they have a criminal record on job applications. While a laudable goal, the result of these efforts has often been to reduce the likelihood that African American men get jobs at all – unable to ask about criminal history on applications directly, employers instead assume that men with names that "sound" African American are more likely to have criminal records and simply don't interview any such individuals. This is an example of a simple design intervention (ban the box) with a laudable goal (give more people opportunities to get jobs) that in many cases has had the opposite effect (even African Americans without criminal records now have a harder time of getting a job).

The only way to understand whether a design practice is beneficial or harmful is to focus on the actual effects of that practice.

2. Professor Hurwitz, how is product design used to attract consumers? If the federal government were to regulate how companies may or may not design their products, what effect do you expect that to have on the free market and competition?

Product design is used to attract consumers in myriad ways – most often by demonstrating of highlighting the value of products to consumers. Sometimes this value is superficial (consider a flashy but low-performance sports car), but sometimes even these superfices are valuable to consumers (consider the driver who enjoys having world think he owns a fancy sports car). Importantly, product design is often used precisely to *attract* consumers – to get them in the door, not to close the sale. Most products in the economy are not commodities, where every firm

sells identical products. Rather, they are differentiated products, where different firms sell products that are similar but not identical. Consider cars: Honda, Ford, Toyota, BMW, and many other companies sell sedans – but no to companies sell identical sedans. Companies will highlight aspects of their products, sometimes trivial or irrelevant ones, in order to get the attention of consumers. Importantly, this practice generally makes it more likely that consumers will compare more products from differentiated firms, which increases competition between those firms and decreases prices paid by consumers. That is, a cute advertisement of a dog driving car actually lowers the price that consumers pay for all cars, even though the advertisement communicates nothing of substance about the actual product.

In general, the best response to concerns about design is to rely on competition to address them. Poor design decisions create opportunities for competitors to enter the market with better products.

Regulating product design runs two parallel risks. Fist, as discussed in my response to the previous question, it is very likely that the government regulation will get things wrong and will make consumers worse off. Such regulations may proscribe designs that are counterintuitively beneficial to consumers or may mandate designs that are actually harmful to them. And, second, by specifying practices that firms must or cannot use, it reduces opportunities for innovation, experimentation, and competition. This is likely both to harm consumers today as well as to deprive future generations of beneficial technologies.

This is not to say that Congress and regulators should not be concerned about potentially problematic design practices. But any decision to regulate should be narrowly tailored to address design practices that can be demonstrated to have material harmful effects on consumers.