

UC BERKELEY  
CENTER FOR LONG-TERM CYBERSECURITY



CLTC WHITE PAPER SERIES

# Tech Has an Attention Problem

AILEEN NIELSEN

**Cover Image:** “Opt Out,” created by Thomas Grimer, a freelance creative with a focus on new product and service concept generation based in the United Kingdom, as part of OpenIDEO’s Cybersecurity Visuals Challenge, in partnership with the William and Flora Hewlett Foundation. The image reflects the extent to which mobile devices have a hold on the attention of users, and how, through our continuous transmission of data, digital devices may be watching us as much as we are watching them.

CLTC WHITE PAPER SERIES

# Tech Has an Attention Problem

AILEEN NIELSEN

SEPTEMBER 2021



C E N T E R F O R L O N G - T E R M C Y B E R S E C U R I T Y

University of California, Berkeley



# Contents

## **Executive Summary 1**

## **I. HUMAN ATTENTION AND THE MAXIMIZATION OF KEY PERFORMANCE INDICATORS 3**

Outline 5

Attention as a pillar of technological fairness 5

Attentional costs and harms 7

Why now? 7

    Ubiquitous digital presence 8

    Enhanced science of human behavior 9

    Meeting markets on their own terms 11

Law and attention 12

    Attention as privacy 13

    Protected spaces and places 15

    Rights and wrongs of private parties as asserted under contract and tort 17

## **II. AN EXPERIMENT WITH ATTENTION METRICS 19**

Potential categories for attention metrics 19

Experiments to measure marketplace reaction 23

Experiment design & procedure 23

Results 25

Download response to attention warning label 25

Judged accuracy of the warning label 26

Specificity of response to the warning labels 27

Experiment discussion 29

Experimental conclusion 30

### **III. WHAT TO DO WITH AN ATTENTION METRIC? 31**

Attention by design 31

Labeling requirements 33

Audits 34

Concrete and particularized harms 35

Digital attention taxation 36

### **NON-MONETARY METRICS AND SOME CONCLUDING THOUGHTS 38**

**Appendix 40**

**About the author 45**

**Acknowledgments 45**

## Executive Summary

People increasingly spend their time in digital environments that are privately owned and engineered to maximize the utility of the entities that control those environments. Unsurprisingly, the interests of the designers and controllers of these digital spaces often run contrary to the interests of the people who use them, resulting in a range of potential conflicts of interest. This paper addresses a category of costs and potential harms that has not yet been adequately contemplated by law and policy communities: the pervasive, unregulated, undercompensated, opaque, and often harmful harvesting of human attention that results from design decisions made to optimize for the interests of owners of digital infrastructure, without regard to the wellbeing of consumers.

The discussion proceeds in three parts. In Section I, I briefly discuss examples of how human attention is under attack in digital products, and how resulting attention harms map onto various branches of law. In Section II, I propose some candidates for intuitive metrics — attention metrics — that could be used to quantify the cost of digital products on attention, and I describe the results of an online experiment designed to measure likely marketplace reactions to warning labels that reflect potential attention metrics. In Section III, I examine a variety of policy measures that could be implemented with attention metrics.

Legal and regulatory use of attention metrics could pave the way to more descriptive and accurate empirical methodologies for assessing consumer welfare in digital environments. Such methods could be applied to urgent and ongoing debates and legal investigations regarding consumer welfare and antitrust. For example, Section III explores several avenues of use for attention metrics in law and regulation, including:

- Labeling requirements, to make the attention costs of a product transparent before use;
- Inclusion of attention cost criteria in auditing requirements for digital products, such as those proposed for algorithmic impact assessments;
- Use of attention metrics to render harms from digital products “concrete and particularized,” as may be necessary to gain access to federal courts;
- Taxation of attention-harvesting activities, to disincentivize careless use of attention harvesting and to more accurately tax value-creation and income in modern attention economies.

## T E C H   H A S   A N   A T T E N T I O N   P R O B L E M

More generally, attention metrics could improve digital product quality and enhance marketplace transparency, creating incentives toward the design of digital infrastructure that responds to welfare concerns of the *users* of digital products, rather than solely the utility of the owners of digital infrastructure. With the right incentives, the technology of the future can improve along many axes of quality, including those that measure costs to consumers, rather than merely the gains to producers of digital technologies.

# I. Human attention and the maximization of key performance indicators

*I'm the dev that built Netflix's autoplay of the next episode. . . . When I worked there the product team at Netflix had two KPIs all new features were tested against: hours watched and retention. We would come up with all sorts of ideas to try out, and release them to small user populations. . . . It was great because you didn't have to debate much about whether a new feature was a good idea or not, you just built it and tested it. If the feature didn't increase hours watched or retention in a statistically significant way, the feature was removed.*

*Autoplay massively increased hours watched. I can't remember the exact numbers, but it was by far the biggest increase in the hours watched KPI of any feature we ever tested. . . .*

*As part of the autoplay test, we tested how long the countdown should be between episodes: 5 seconds, 10 seconds or 15 seconds. 10 seconds caused the biggest increase in hours watched. We thought that it gave people time to digest what they had just watched, but wasn't too fast (5 seconds) where it became jarring. Interestingly, Netflix recently changed the countdown between episodes to 5 seconds. That means they tested it out and found that people watch more if with a shorter countdown. This didn't use[d] to be the case. Netflix user[s] have become conditioned to expect autoplay.*

— HackerNews user rsweeney21<sup>1</sup>

Key performance indicators (KPIs) are quantitative metrics of a firm's financial health and growth. Many KPIs for tech companies are based on metadata generated and collected during the surveillance of people as they use digital products. There is a particular subset of KPIs geared to measure how much and how “well” users engage with a digital product. These fall under the heading of “engagement metrics.”

Producers of digital products understandably want people to engage with their products, both for direct profit motives but also, indirectly, because higher engagement metrics putatively

<sup>1</sup> rsweeney21 (2019), comment, “New bill would ban autoplay videos and endless scrolling” (2019), <https://news.ycombinator.com/item?id=20565141>, last accessed Aug 20 2020. Quoting user rsweeney21, who also links to the public Twitter profile of Robert Sweeney, <https://twitter.com/rsweeney21?lang=en>, the self-identified “Creator of Netflix's autoplay next episode.” The KPI-related practices described and the attitude about these practices are representative of the views and practices of many working in engineering and product roles at technology firms. See e.g. J Williams (2018), *Stand Out of Our Light: Freedom and Resistance in the Attention Economy*, Cambridge University Press.

suggest that humans like a product and find it useful. Hence producers of digital products often seek to maximize KPIs that target engagement as a standard industry practice. The logic and practice of maximizing engagement KPIs apply widely to digital products, as much to targeted ads as to video streaming services, and as much to physical mobile health devices as to the consumption of services offered through a digital interface. The logic of engagement maximization is sufficiently widespread to be viewed as a significant factor in shaping digital infrastructure, or the de facto digital environment “in” which many people spend large amounts of working and leisure time.

Yet in contrast to industry, lawmakers have largely ignored engagement metrics, even when considering policy interventions related to digital products. Consider for example that no explicit use of engagement metrics is made for determining which businesses are covered by data protection laws. Such a disconnect is noteworthy, given that most data-driven businesses both advertise their engagement figures and strongly emphasize these on valuations of company worth and forecasting of company profits. Commercial firms understand and respond to engagement metrics, which can also be understood as attention metrics, even if regulators and lawmakers do not.

There is a relatively small and esoteric genre of proposed laws that dance around the problem of engagement via “technology addiction,” a notion invoked commonly in the cases of video games or mobile apps, but also in cases of internet use more generally.<sup>2</sup> Such proposed laws (which, to date, have not been passed into law in any U.S. state) have done little more than grab headlines, in part because of a lack of scientific justification for the putative causal pathways these proposed laws purport to regulate. One contribution of the attention metrics proposed in this work is to provide scientific data about mechanisms that may connect attention harvesting to consumer welfare.

The loss of quality and utility to consumers — and to society at large — that may result from an undue emphasis on maximizing engagement is a primary motivation for this paper. Thus, the goal of this paper is to explore why the maximization of product engagement, and the attention-harvesting that occurs in many digital products, urgently raise the need for scientifically validated quantification. The development of attention metrics can serve as a first step to engage systematically with the widespread practice of attention harvesting,

<sup>2</sup> See e.g. Senator Josh Hawley’s proposed legislation to regulate certain in-app purchase practices, as described in B Lowry (2019), “Missouri Sen. Hawley finds a new target in his war with tech industry: Candy Crush”, *The Kansas City Star*, May 8 2019, available at <https://www.kansascity.com/news/politics-government/article230159634.html>. When Senator Hawley released the proposed legislation, he argued that “game developers shouldn’t be allowed to monetize addiction” and that children in particular should be protected from “compulsive microtransactions.”

contemplate potential welfare effects at the individual and social levels, and pave the way towards exploring, if justified, regulatory interventions to enhance consumer well-being.

This paper presents an experimental proof of concept for the utility of attention metrics. In an era in which people increasingly live their lives through the mediation of digital products, attention metrics can be informative to ordinary people, shape digital consumption behaviors, and enhance the welfare of humanity in a digital world.

## **OUTLINE**

This paper proceeds in three parts. In the remainder of Section I, approaches from law are discussed with respect to attention quantification and harms. More specifically, I survey how attention harms have been or could be recognized and addressed by diverse areas of law.

In Section II, I present an experiment to see whether attention-related information would affect consumer behavior in a common scenario: the downloading of an entertainment-related app from a smartphone app store. I find that some forms of information about the attentional harms of a smartphone app do indeed affect the likelihood of consumers to download the app. This provides experimental proof for further exploring the concept of attention metrics and launching attention metrics as a key research agenda for those interested in documenting and rectifying digital harms generally.

In Section III, I discuss what could be done if robust and externally valid attention metrics could be reliably measured for digital products. I look to both voluntary self-regulation use cases within industry and also the possibility of compulsory legal measures, each of which could be premised on robust and intuitive attention metrics. A wide range of interventions in favor of consumer welfare would be possible through attention metrics that accurately quantify the attention costs and legally cognizable harms associated with attention harvesting and engagement maximization.

## **ATTENTION AS A PILLAR OF TECHNOLOGICAL FAIRNESS**

Debates about fairness and technology have in recent years largely focused on issues related to data privacy and algorithmic discrimination. These widely discussed concerns are tied to the

widespread data collection and profiling of individuals by private firms, which employ a variety of techniques to engineer digital infrastructure use in their favor. Thus privacy and equality concerns are related to practices driven by the maximization of engagement metrics, but they do not encompass all the potential harms imposed by the practice of engagement maximization.

Solutions for problems related to privacy or algorithmic fairness likely can be deployed without fundamentally undercutting engagement KPIs or the level of use of digital products. In other words, an app can respect personal data protection norms and laws, not discriminate on the basis of an impermissible category, and even meet other responsible AI targets, such as providing explanation and accountability for decisions, all while continuing to consume levels of human attention that may impose unjustifiable attentional costs at individual or societal levels.<sup>3</sup>

Problems that may result from maximizing engagement are not addressed by existing concerns or proposed solutions put forward, for example, by the Responsible AI movement. Importantly, while concerns about engagement maximization, and the tightly related practice of attention harvesting, need to be eliminated to produce “fair AI,” problems of attention harms apply far more broadly than purely to the realm of algorithms and Responsible AI. Concerns about engagement maximization and attention harvesting apply in a far wider range of actions in the design and manufacture of digital environments and products, all the way from decisions about how to design a digital interface to decisions about where and how to deploy or market a digital product, regardless of the extent to which a digital product includes algorithmic or otherwise data-driven elements. There are real and ongoing harms resulting from overuse of some digital products, and even whole categories of digital products.

There is robust evidence that there is a science of increasing engagement, and that designers of digital products, among many others, use sophisticated manipulations to enhance engagement.<sup>4</sup> These harms are independent of privacy- and discrimination-related harms. Those working within the domain of algorithmic fairness must recognize and respond to the widespread phenomenon of attentional harms in existing digital products.

It is likely that many of attention harms are delivered with awareness (if not intention), are avoidable, and are the result of sophisticated industry players seeking to accrue one-sided advantages. These harms are unlikely to be corrected without regulatory action or legal redress for existing victims. Thus, attention costs as measured by the newly proposed concept of

3 Justification (or the lack thereof) for a particular digital design and associated attention costs could be assessed in a number of ways, including as a normative matter or as a cost-benefit analysis. That is beyond the scope of the present work.

4 See, e.g., A Alter (2017), *Irresistible: The Rise of Addictive Technology and the Business of Keeping Us Hooked*, Penguin Press.

attention metrics represent an important novel proposal for an element of digital fairness that requires urgent action from the Responsible AI community and other industry and academic groups concerned with digital fairness, including regulators and lawmakers.

## **ATTENTIONAL COSTS AND HARMS**

The tech world has many attention metrics, but few if any measure the attention cost to *consumers* associated with a digital product. Imagine that, rather than measure engagement for a social media platform, one could measure the resulting disengagement from work, social life, or family life. Such lost time could be reframed as a direct attentional cost, or even harm, resulting from the deliberate taking of time away from activities that people would engage in more when not facing product design choices made to maximize product engagement, possibly without any other welfare-enhancing effects.

Of course, it would likely be incorrect to attribute all, or even most, time spent with a digital product as an attention harm. Rather, this notion of attention harm necessarily interacts with notions of a private entity's legal and ethical duties to others. Such duties could encompass intuitive notions, such as the need to respect individuals' autonomy, or their rights to make thoughts and observations without interference. However, a full contemplation of the distinction between attention costs and harms is a complex topic beyond the scope of this paper.

Documenting the temporal and other attentional costs of using a product will likely bring more clarity to the highly opaque and manipulated nature of many digital products. Recognizing that attentional harms are likely commonly occurring — and pairing this recognition with a metric for attentional costs — will further add clarity. That is what is offered in this report. Later work can then develop methods, both theoretical and empirical, to systematically and scientifically distinguish attention costs from attention harms.

## **WHY NOW?**

Attention costs and harms are not new. Attention harvesting practices and technologies have existed far longer than the big data, AI, and ubiquitous computing revolutions associated with the digital products on which many people now spend large amounts of time and attention. Critics could fairly ask why it is worth revisiting attention costs now.

Briefly, it is a matter of scale and power. Today's digital product capabilities and deployment operate at scales, levels of granularity (i.e. personalization), and possibly at levels of effectiveness distinct from those of earlier forms of attention harvesting, such as newspaper or television advertisements. There are reasons why attention harvesting carried out through digital design for engagement maximization represents a newly potent manifestation of an old practice. These reasons are surveyed briefly below.

### Ubiquitous digital presence

Technology in the form of digital products is far more pervasive, personalized, and powerful than it has ever been in the past. Consider push notifications on smartphones, which can arrive at any time and which intrude into our attention at the time of the sender's choosing.

Not only is such a technology designed to allow attention harvesting based on the sender's preferences rather than the recipient's preferences, but these notifications have been found to have significant attentional costs.<sup>5</sup> It is hard to imagine that past attention harvesting practices can in any way compete with a push notification in the ability of a firm to reach anyone almost anywhere and at any time.<sup>6</sup>

Now, through their design of engagement-maximizing digital products, firms can reach people in their beds, in the bathroom, or when they are bored on a weekend morning at the breakfast table. Far more attention is available to harvest, in part because there are more people hooked

5 C Stothart, A Mitchum, and C Yehner (2015), "The attentional cost of receiving a cell phone notification." *Journal of Experimental Psychology: Human Perception and Performance*, 41(4), 893-897.

6 Of course, after many years of firms' pioneering push notifications, some countervailing measures have been offered to consumers, such as the ability to silence notifications, but it is notable how long it took for such options to become available. With attention metrics to empower marketplace transparency and a more holistic evaluation of digital product quality, consumers might have been empowered far sooner to control push notifications. Some details of iOS push notifications are provided here. Push notifications were first introduced in iOS 3.0, in 2009. However, controls for push notifications lagged. For example, users did not have the ability to mute local push notifications until iOS 8.0, in 2017.

Namraata Badheka, "The History of Push Notifications," medium.com, Jan 16 2017, available at <https://medium.com/the-pushcrew-journal/the-history-of-push-notifications-43343bdf2d85>, last accessed Jul 27 2021, "In June 2009, Apple introduced a new update for its iPhone series, iOS 3.0, with which Apple gave its customers a new user experience. This was Apple's version of push notifications, and they called it APNS, the Apple Push Notification Service. The users got alerts, much like Blackberry's, from Apple's kingdom, as well as from an ecosystem of little utilities Apple called 'apps'." Greg Weinger, 2017, How Apple's iOS 8 changes the game for brands and retailers, MarketingDive.com, available at <https://www.marketingdive.com/ex/mobilemarketer/cms/opinion/columns/19729.html>, last accessed Jul 26 2021, "Prior to iOS 8, users had no control over the display of local notifications. Some apps even took advantage of this feature to get around opt-in user preferences.

Because there is no way to visually distinguish a local notification from a push-based notification, this frustrated users who thought they had opted out of push but continued to see local notifications."

up to internet-connected devices at any given time, and for many individuals, seemingly at all times. Yet, at the same time, only a few social conventions<sup>7</sup> have arisen to limit the always-on state people have entered through new technologies. A related but distinct fact is that pervasive technologies permit ongoing surveillance by commercial entities, which in turn is used to collect data that is thought to better assist with engagement maximization and concomitant attention harvesting.

From the perspective of firms designing or making use of digital products for attention harvesting, this all combines into a virtuous cycle, one that gives and gives more from their perspective. They have greater access to attention harvesting, combined with lower costs to access that attention and greater observational opportunities to tune their pitch and make the most of the attention they capture.

### **Enhanced science of human behavior**

The science of human behavior has progressed significantly in recent decades, in many cases aided by the scaled-up and less intrusive means of collecting data and running experiments on human behavior enabled by increasingly small and ubiquitous sensing and recording technologies. In diverse circumstances, behavioral scientists have demonstrated that human autonomy can be easily defeated or manipulated<sup>8</sup> with small and unobtrusive design decisions, even in important decisions that would seem to invoke careful reflection and lower receptiveness to manipulation. Two simple examples can show the powerful ways in which human decision-making can be manipulated.

7 For example, in some but not all social situations, it is considered rude to look at one's smartphone during a meal. However, many such putative conventions seem to be violated at high rates, and in other cases it is not clear that there is a convention. For example, Smithers (2013) and Moscaritolo (2018) report on results of surveys about cellphone use during ceremonies or in locations that were traditionally robust against external distractions prior to the existence of cellphones and smartphones (funerals, using the toilet). Rebecca Smithers (2013), "One in six admits to using mobile phone at funerals," *The Guardian*, May 15 2013, available at <https://www.theguardian.com/money/2013/may/15/mobile-phone-use-at-funerals>. Angela Moscaritolo (2018), "Who Uses Their Phone on the Toilet? Most of Us," *PC Magazine*, Oct 8 2018, available at <https://www.pcmag.com/news/who-uses-their-phone-on-the-toilet-most-of-us>. Note that I do not offer a value judgment regarding whether smartphones should be used in such cases but rather point to the fact that a new technology enabled the circumvention of longstanding expectations of privacy (in the sense of freedom from interference or distraction) in such ceremonies (funerals) or spaces (restrooms).

8 It is beyond the scope of this report to define or explore the nuances of defining human autonomy. Here, for simplicity of discussion, if an intervention in a randomized controlled study changes outcomes between groups in a statistically significant way, that intervention is deemed to have defeated or manipulated human autonomy. Such showings are made in aggregate, and it may not be possible to delineate which people had their autonomy undermined and which made the same decision they might have done absent an intervention.

First, consider the case of opt-in versus opt-out mechanism design, here considered through organ donation studies. In Western countries with opt-in regimes, such as Austria, rates of registering to donate organs are high, even exceeding 90%. In contrast, Western countries with opt-out regimes, such as Germany or the United States, have registration rates of 15% or below. Thus, a serious decision with life-saving ramifications can have widely disparate rates, likely due to mechanism design rather than disparate cultural or social issues.<sup>9</sup> This simple way of modifying human decision-making — be it for good or otherwise — can and is often applied in digital settings, as in the case of website forms that often opt people into advertising emails or opt them into consent of privacy policies, with the designer’s choices designed to take advantage of a number of widely recognized digital “dark patterns.”<sup>10</sup> Mechanism design choices, such as opt-in versus opt-out, can and do substantially influence the exercise of human autonomy, and particularly so in digital platforms.

As a second example, consider the general question of digital device privacy. While Big Tech and the Supreme Court<sup>11</sup> alike agree that the privacy of personal digital devices is important, the reality is that a host of digital design choices<sup>12</sup> or seemingly slight social pressures<sup>13</sup> can result in people agreeing to privacy policies or intrusions into their digital devices that would normally be considered unreasonable.

Large technology companies employ thousands of data scientists, many of whom study how humans interact with or otherwise respond to their digital products. It has become common practice to continually experiment on users, unbeknownst to them, in ways that would

9 Davidai, S., Gilovich, T., & Ross, L. (2012). “The meaning of default options for potential organ donors.” *Proceedings of the National Academy of Sciences*, 15201–15205.

10 Consider this widely cited taxonomy: <https://www.darkpatterns.org/types-of-dark-pattern>, where, among others, “forced continuity” and “sneak into basket” rely on mechanisms that aggressively make use of and even extend the influential power of opt-out mechanisms.

11 Sabrina McCubbin (2018), Summary: The Supreme Court Rules in *Carpenter v. United States*, <https://www.lawfareblog.com/summary-supreme-court-rules-carpenter-v-united-states>.

12 Ari Ezra Waldman (2020), “Cognitive biases, dark patterns, and the ‘privacy paradox,’” 31 *Current Opinion in Psychology* 105, <https://www.sciencedirect.com/science/article/pii/S2352250X19301484>. Behavioral manipulations resulting from digital design choices strongly undercut the notion that a legal formula of notice and consent is adequate to protect privacy in the digital domain.

13 In a behavioral experiment, 100 of 103 experimental subjects agreed to unlock their phone when asked to do so by an unknown experimental supervisor, although a majority of participants in a separate experiment predicted that people would not do so. This shows that people undervalue the degree of influence imposed on an ordinary person’s actions through light social pressure, such as an in-person oral request. Roseanna Sommers & Vanessa K. Bohns (2019), “The Voluntariness of Voluntary Consent: Consent Searches and the Psychology of Compliance,” 128(7) *Yale Law Journal* 1792.

be considered ethically problematic if carried out in academia.<sup>14</sup> This has led to a rapidly developing science of human behavior specific to digital products, for which experimental results may go unpublished as confidential in-house information. Given the above established facts that human autonomy is often subject to manipulation, and particularly so in digital environments, it seems likely that the efforts of these data scientists has been put to highly influential uses in the everyday digital environments in which people spend increasing amounts of time.

Thus, important features of social science as practiced in the era of current digital products and digital infrastructure reflect the possibility that digital products can be and are designed in a way that maximizes the interests of their designers, sometimes to the detriment of the ordinary people using such products. Yet to date, little in the way of regulatory action or law has recognized this possibility or the ramifications for human attention.<sup>15</sup>

### Meeting markets on their own terms

Attention increasingly drives the U.S. economy and represents a highly profitable commodity in global trade. For example, many of the largest companies in the U.S. (as assessed by market capitalization) offer services related to attention harvesting.<sup>16</sup> New efforts in many European countries to tax digital services might be understood in part by reflecting the economic value of attention collected by digital products.

Relatedly, some scholarly dissatisfaction with the current shape of digital markets and digital services has been channeled into critiquing how existing regulatory bodies and legal standards rely too heavily on monetary metrics, rather than on other measures of market behavior and consumer welfare. For example, a growing body of antitrust scholarship argues that antitrust law in its current form, driven by consumer welfare tests based on monetary pricing, may be inadequate for addressing harms to consumers that occur in markets associated more with attention costs rather than financial costs, such as could the case for social media. In other words, regulators are ignoring the currency of attention and likewise ignoring the attention

14 R Meyer (2014), “Everything We Know About Facebook’s Secret Mood Manipulation Experiment,” *The Atlantic*, June 28 2014, available at <https://www.theatlantic.com/technology/archive/2014/06/everything-we-know-about-facebooks-secret-mood-manipulation-experiment/373648/>.

15 While the Federal Trade Commission in the U.S. or the European Union’s Competition Authority have at times invoked behavioral facts regarding the influence of digital design in various enforcement decisions, such legal bodies have yet to pursue the development of a rigorous set of metrics that can measure relevant quality metrics for digital products so as to quantify the degree to which products may be deploying unfair designs.

16 For example, Apple, Facebook, Google, Amazon, and Microsoft all sell targeted advertising options.

costs that consumers may be forced to pay due to anticompetitive practices and/or a lack of competition in markets where attention harvesting can be conducted by dominant firms without suffering a loss of customers or market share.<sup>17</sup> More scholarly and regulatory consideration is needed for non-monetary metrics of wealth and power, including attention metrics.

## LAW AND ATTENTION

Some notion of attention costs and harms is not entirely new in legal scholarship. For example, the broad use of attention as used in this paper relates to a subset of privacy defined in Professor Daniel Solove's (2006) taxonomy of privacy. Specifically, Solove identified one branch of privacy as relating to "invasion," and within this branch he demarcated "intrusion" and "decisional interference," both of which relate to attention costs imposed by digital design. As a digital example, push notifications could be seen as an intrusion into a person's thoughts or private space, but likewise could be seen as decision interference if they deliver a message or invite an action with a design engineered to produce a specific behavior. Imagine, for example, if Facebook's experiment<sup>18</sup> on prompting their users to vote had been successfully conducted via a push notification.

Likewise, attention costs and harms overlap with "dark patterns," which are deceptive features added to user interfaces to manipulate users into certain behaviors. Dark patterns have been investigated more recently by legal scholars<sup>19</sup> and technologists<sup>20</sup> as examples of digital design patterns that engineer behavior against a user's own interest. Some dark patterns increase the time a person must spend to get to their desired goal when using a digital interface, a temporal attention cost. Other dark patterns increase the amount of mental effort a user must invest to

17 Future work is urgently needed to address the extent to which attention costs may be exacted as monopoly rents, or, in the alternative, as a form of anticompetitive behavior. There is an empirical question as to whether firms may seek to maximize engagement, not for pro-competitive reasons of product quality, but rather for the anti-competitive reason of reducing the amount of consumer attention available for potential competitors, or for the reason that they enjoy a monopoly position in a market with significant barriers to entry.

18 Bond, R. M. et al (2012). "A 61-million-person experiment in social influence and political mobilization," *Nature* 489, 295-298.

19 J Luguri and L Strahilevitz, "Shining a Light on Dark Patterns" (August 1, 2019). *U of Chicago, Public Law Working Paper* No. 719, University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 879, Available at SSRN: <https://ssrn.com/abstract=3431205>, last accessed Aug 24 2020.

20 A Mathur, G Acar, MJ Friedman, E Lucherini, J Mayer, M Chetty, and A Narayanan. 2019. "Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites." *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 81 (November 2019), 32 pages. <https://doi.org/10.1145/3359183>

avoid unwanted outcomes, such as reading pre-checked options and opting out of them, rather than opting in. Dark patterns often impose attention costs on a user in order to further the interests of a digital product owner. Importantly, while most scholarly focus on dark patterns has emphasized their use to defeat human autonomy, dark patterns are also problematic because they impose additional attention costs on users without a compensating benefit.

Thus, notions related to attention have been treated in the legal academic literature, but largely by implication rather than explicit emphasis. “Attention” as such has received little coverage in academic research to date.<sup>21</sup> Thus this paper briefly surveys some legal entries into the topic of attention to show the many ways in which attention harms and costs could be recognized and reasoned about under existing U.S. law.

### Attention as privacy

A classic case of attention as privacy is *Public Utilities Commission v. Pollak* (1952),<sup>22</sup> a Supreme Court case that addressed the decision by a street railway company in the District of Columbia to broadcast radio programs through loudspeakers in its streetcars and buses. The legal basis for the constitutional review upon which the case relied came from the fact that the streetcar company was so heavily regulated that its federal regulator could potentially be liable for constitutional infringements if any were found.<sup>23</sup>

The plaintiffs challenged the streetcar company’s playing of the radio on the basis that the practice infringed on their freedom of expression and right to privacy. In the instant fact pattern, the Supreme Court found that being forced to listen to content while making use of a public utility, the railcar service, did not rise to being a constitutionally cognizable harm with respect to freedom of expression or with respect to privacy even if the plaintiffs did not have the ability to opt out of public transport.

The Court reasoned that the radio programs did not infringe these rights because the factual record indicated that the radio did not interfere with conversations between passengers or with other rights of communication constitutionally protected in public places. In the case of

21 Importantly, this absence of scholarship is finally being addressed by a handful of scholars. The author notes some recent publications and working papers in 2020 and 2021 finally circulating to address this important and urgent topic explicitly.

22 *Public Utilities Commission v. Pollak* (1952), 343 U.S. 451.

23 A note for non-legal readers: U.S. constitutional protections apply only against the government. For example, the First Amendment protects speech only against government intrusion. Thus, the Constitution does not prohibit a private employer or private business from restricting speech.

Fifth Amendment privacy protections, the Court reasoned that the streetcar was a public place where the right of privacy was not as extensive as would be expected in more private spaces, such as the home.

Notably, the Court did not reject the notion that external stimuli *could* infringe on privacy or freedom of expression. Rather the Court rejected that idea in the specific facts presented by the plaintiff's case. Its reasoning for not finding an infringement in the case suggests some ways in which future cases could have a different outcome.

The Court noted that, unlike in one's home, where it would clearly be unconstitutional for the government to force an individual to listen to content, an individual has less expectation of quiet on a railcar. The Court also noted that the content was of general interest and not designed to influence the listeners or inculcate them with propaganda in the government's interest. This suggests that content that was judged not to be of general interest or that was judged to have too great a persuasive element could be constitutionally problematic, all else equal. Also, the Court noted that only a relatively low percentage of time was used for public service announcements (5%) or advertisements (5%), far lower than the portion of airtime (or screen space) nowadays devoted to advertisements or other forms of attention harvesting in modern digital products.

The *Public Utilities Commission v. Pollak* decision thus leaves open important avenues to distinguish levels of engagement maximization, or attention harvesting that could be constitutionally or otherwise legally problematic. The Court's reasoning could be applied to recognize that some such practices are harmful and could even infringe on privacy or freedom of expression if practiced by the government. We can therefore wonder whether the Court would have distinguished the case of technologies that now follow us not only into the homes, but into our most private moments, where they can influence our seemingly private decisions.<sup>24</sup>

There is, of course, a key limitation to how far this reasoning applies. Any protection the Supreme Court might grant to human attention on the basis of the United States Constitution (as the plaintiffs sought in *Public Utilities Commission v. Pollak*) would apply only against the government because the Constitution provides protection only against governmental infringements on rights.

<sup>24</sup> See e.g. S Wolfson (2018), "Amazon's Alexa recorded private conversation and sent it to random contact," *The Guardian*, May 24 2018.

Nonetheless, the reasoning in *Public Utilities Commission v. Pollak* identifies legally salient bases that could be applied more widely if enacted into legislation or regulatory rules. The decision illustrates a principled basis for distinguishing reasonable from unreasonable attention-related infringements on ordinary people.

### Protected spaces and places

Law has traditionally recognized special spaces where people are vulnerable or entitled to higher expectations of privacy and quiet. A special degree of legal protections for privacy and tranquility are reserved for intimate spaces, such as the home. This is true in both public and private law, as seen in Fourth Amendment Constitutional protections against government searches<sup>25</sup> and in the common law torts with respect to privacy.<sup>26</sup> This concern for particularly protected places was already evidenced in the discussion of *Public Utilities Commission v. Pollak* above, with the Supreme Court finding that public transport was not a place with a special expectation of privacy.

Some legal contemplation or recognition of the problems of engagement maximization and attention harvesting could occur through branches of law where there are special protections for intimate places. Courts could, in theory, reason by analogy that certain techniques, i.e. digital designs used to impose attention costs, are not in conformity with the legal duties imposed by law in particularly sensitive cases, akin to certain forms of privacy invasion by a private party or by a government search being unacceptable in areas where there is a special and reasonable expectation of privacy. Alternatively, legislators could write specific laws to clearly delineate the bounds of acceptable forms of attention harvesting where this occurs in private spaces or in particularly intimate moments if they found that enforcement under existing laws was inadequate or insufficiently on point when looking at the legal formalities of establishing privacy invasions.

In recognizing attention harms under existing or new law, including statutes, courts need to find that these harms rise to a level sufficiently serious to gain access to a court. This is an issue generally treated under the rubric of standing, which addresses whether a plaintiff is a proper person to bring a case and in particular whether the plaintiff has raised a harm that

25 SM Stern (2010), “The Inviolable Home: Housing Exceptionalism in the Fourth Amendment.” 95 *Cornell Law Review* 905 at 906. “The notions of the inviolable home and the paramount importance of constraining government search of the home are cherished tenets of constitutional law and scholarship.”

26 AJ McClurg (1995), “Bringing Privacy Law Out of the Closet: A Tort Theory of Liability for Intrusions in Public Places,” 74 *North Carolina Law Review* 989 at 990.

is addressable and sufficiently important to merit the attention of a court. The doctrine of standing can be a particular barrier in the case of gaining access to federal court.

Nevertheless federal courts have recognized conceptual connections between traditional torts, such as inclusion upon seclusion,<sup>27</sup> and harms resulting from new technologies, such as robocalls or unsolicited text messages.<sup>28</sup> In litigation under the Telephone Consumer Protection Act, various federal courts have found that unsolicited text messages constitute a sufficiently concrete injury to establish standing due to their connection to both intrusion upon seclusion and nuisance.<sup>29</sup> This trend is interesting because both intrusion upon seclusion and private nuisance law are commonly associated with a special place where the plaintiff has some reasonable expectation of a right not to be disturbed.

### **Rights and wrongs of private parties as asserted under contract and tort**

Contract law and tort law are major sources of law for private individuals to assert rights and to seek compensation for wrongs perpetrated against them. Specifically, contract law is the law of enforcing agreements made between private parties. Tort law is the law of correcting violations of legal duties that private parties have towards one another as members of society, independent of any specific agreement between them.

Contract law and related legal logic offer some ways to think about how attention metrics and attention costs could be recognized in legal processes. Consider laws designed to govern the sale of goods. In some cases, an unwritten warranty of merchantability can be implied into the contract for a sale of goods, requiring the purchased goods to conform to some minimal standard of quality as could be expected by a reasonable person. If attention metrics come into ordinary use and attention costs are recognized as important in the design and quality of goods, products that engage in unjustified attention harvesting could be found to violate implied warranties of merchantability.

27 The tort of intrusion upon conclusion generally covers an intentional intrusion into the private space or private affairs of someone in a way that would be offensive to a reasonable person and that resulted in mental anguish or suffering for the victim.

28 This is important because relatively new jurisprudence has served to put a check on the judicial recognition of some harms alleged under new technologies through a requirement of concreteness, assessed in part by looking to whether a harm maps onto a traditional tort.

29 J Erpenbach (2019), “A Post-Spokeo Taxonomy of Intangible Harms,” 118 *Michigan Law Review* 471 at 486 citing *Van Patten v. Vertical Fitness Grp., LLC.*, 847 F.3d 1037, 1043 (9th Cir. 2017) (“Unsolicited telemarketing phone calls or text messages, by their nature, invade the privacy and disturb the solitude of their recipients.”) and *Susinno v. Work Out World Inc.*, 862 F.3d 346, 352 n.3 (3d Cir. 2017) (“[I]ntrusion upon seclusion is a well-recognized subset of common law invasion of privacy.”)

Relatedly, but distinct from contract law, the Federal Trade Commission (FTC) investigates and penalizes unfair and deceptive business practices. One standard for assessing such practices is to note when companies fail to disclose important elements of design in their digital products or associated business practices. If the FTC were to find that certain digital products impose enormous, but hidden or otherwise deceptive or unfair attention costs, it could investigate or penalize such activities as a form of regulatory consumer protection. The model of the FTC is distinct from contract law but relies on notions, such as notice and consent, that are inspired by a contract-like model of respecting agreements so long as those agreements or pieces of information are clear and honest.

Tort law could offer another avenue for the legal system to account for attention costs imposed by digital products, although the legal arguments would be more difficult for a plaintiff to make, given existing doctrines that often bar compensation for psychological or economic harm (as tort law is primarily oriented around physical harm). There has been a longstanding and widespread skepticism of tort liability in the case of harms regarded to be nonphysical. However, the rationale for this in part rests on the inability to measure such harms, which would not apply if robust attention metrics were created. Also, in some cases, attention costs could lead to physical harm, such as a physical accident resulting from distraction due to a digital product, such as the classic example of texting while driving resulting in an automobile accident.

Attention harms could potentially be recognized under a products liability theory of tort law. This form of tort law was meant to be a form of strict liability whereby any producer or seller of a product along the course of manufacture of that product could be liable for damage caused by that product. This form of liability recognizes the intrinsic complications of modern-day products and the modern day marketplace, where a *prima facie* case is established if a plaintiff shows that a defendant was involved in the manufacture of a product, that the product was defective, and that the defect caused harm. In the case of attention harms, the defect would relate to the way in which a product imposed unduly high attention costs that led to injury. There are many ways in which it can be complicated to establish products liability cases, particularly for software, but products liability could in theory offer a route to address attention costs.

An alternative form of torts liability for attention harms could arise under the more general but more demanding theory of negligence liability. To succeed on a theory of negligence, plaintiffs would have to show that the owner of a digital product had some kind of legal duty to the user

of the product and had violated that duty, resulting in harm to the plaintiff. Attention metrics might serve as a form of evidence for multiple elements necessary to establish negligence and potentially lead to liability.

Finally, a third form of tort liability for attention harms could potentially be established under tort-like rights of action derived from statutory law, possibly asserted on behalf of the public. Such suits for the government to address widespread attention harms might take their cues from the cases some states currently are pursuing in the wake of the opioid epidemic, which are grounded in theories of law where public health and addiction are legally cognizable.

Future legal scholarship could look in detail at contract and tort theories that could potentially lead to redress for unjustified attention costs. The routes discussed here would be promising avenues for further investigation.

## II. An experiment with attention metrics

In this section, I present exploratory work to look at whether affixing attention labels to a digital product description affects likely behavior with respect to that product. The experiment attempts an initial answer to the following: if given information premised on a potentially viable attention metric, will consumers adjust their behavior?

The discussion proceeds in two parts. First, I describe some broad categories of attention metrics one could imagine being sufficiently well-defined to be both intuitive and measurable. Second, I present an experiment in which a sample of U.S. residents were studied to infer the likely marketplace reaction to information premised on attention metrics applied to an imagined digital product.

### POTENTIAL CATEGORIES FOR ATTENTION METRICS

Several important factors related to attention are explored below, with respect to how each could translate into a quantitative indicator specific to a digital product, along with potential benefits or downsides. Ultimately, the selection of an appropriate mechanism would involve complex considerations of accuracy, transparency, ease of use, resistance to strategic gaming, and external validity, among other considerations. The factors presented here thus represent only a brief initial investigation into these possibilities.

#### Time

Time may be one of the best proxies for attention, as it is a fundamentally limited resource that can be straightforward to quantify. Using time as an attention metric would directly address an important potential attention harm, and particularly one related to the much discussed concept of “technology addiction.” Use of time for attention cost would also be quite suitable because time spent using a product is already widely used as an engagement KPI, providing one example of where existing data could be repurposed.

On the other hand, time has some potential disadvantages. First, it does not capture the difference between time well spent and time lost. A digital product that teaches mathematics in a way that improves cognition and a sense of accomplishment might on average take the same amount of time as a product that leads to addiction or decision fatigue. If the average time spent is the same, the attention metric would fail to distinguish obvious and important characteristics of these products. Nonetheless, this metric could be combined with people's own judgments about whether time was well spent, and could provide concrete information as they make this assessment.<sup>30</sup> A second potential problem with time as an attention metric is that there is evidence that people are already aware of the amount of time spent on at least some digital products, suggesting that a metric based on time might not convey novel or hidden information in all cases.

It could be that people become aware of the temporal attention cost of a digital product only after using the product (one example of what economists call an “experience good”). With an attention label premised on time, this information could be frontloaded so that consumers could have it at a relevant decision point, rather than when they have already paid the attention cost so as to gain knowledge of a product's attention cost. What's more, a time-based attention metrics label could help thwart some dark patterns, particularly those designed to make opt-outs time-consuming, by alerting users to this possibility before they begin using a digital product.

### **Advertiser pricing**

An attention metric could take advantage of the long experience of advertisers, who measure attention, price it, and sell it based upon monetary metrics. While perhaps not applicable to all use cases, connecting use of a digital product with advertising prices could offer some indication of the monetary value of the attention harvested. There could be many ways to implement this, including measuring how owners of a digital product price exposure of ads to their users, or how they value users' personal data. This would give users some quantification of how their use of a product is perceived to influence the value (or amount) of their attention for commercial messaging.

Of course, such pricing information is sometimes closely guarded as a trade secret, covered by non-disclosure agreements and other legal mechanisms to protect the information. Thus an obvious downside of such a metric would be that it might impinge on existing legal rights

<sup>30</sup> Also, such a problem is not limited to attention metrics. For example, even the caloric content of food, a physical quantity, does not indicate whether a particular caloric intake is warranted or not, and yet this fact does not defeat the utility of providing such information on food labels.

and could trigger fierce resistance from many industries, who could perceive it as undermining their very business model.

Pricing would also run into some of the problems faced when using time as an attention metric. While a pricing metric would provide some uniformity and objectivity, the use of ad spend would not necessarily be a good proxy for distinguishing value-creating and value-destroying uses of human attention. It would still be up to recipients of this information to use contextual cues or even personal experience to determine whether an advertising pricing metric is more likely reflecting a benefit or a cost of using a given digital product.

### **Observational capacity**

An alternative metric of attention that could be relevant to some digital products is the extent to which engagement with the digital product brings about disengagement with the non-digital environment, as proxied for by some measure of observational capacity. Observational capacity could indicate to what extent attention is captivated during the time of use. For example, one could imagine metrics geared towards highly technical and conceptually rigorous notions of attention, such as employing typical laboratory setups to measure *filter attention*<sup>31</sup> while someone uses a digital product. Or, in quite a different example, one could imagine an empirical study of whether accidents due to distraction are associated with a particular product.<sup>32</sup>

Metrics related to observational capacity could be particularly advantageous for considerations of physical safety as affected by distracting digital products. Likewise, such metrics might be particularly interesting to those seeking to protect vulnerable populations, such as children, from unduly influential or addictive digital products.

There are potential disadvantages to attention metrics that emphasize observational capacity. For example, it could be that most observational measures of attention only capture effects during product use in a way that does not reflect the negative attention externalities of using a product (such as examples of digital engagement leading to harmful forms of disengagement with other domains). Also, observational capacity, while important from a safety perspective, does not necessarily provide information necessary to a cost-benefit analysis of using the product. Being immersed in a digital product, even at the expense of reduced observational capacity, could be judged as beneficial in some cases, but not in others.

31 Filter attention is concerned with how individuals filter out some sensory inputs while processing others.

32 Such studies have led to the widespread banning of texting while driving.

## Cognitive performance

Cognitive performance would be an attractive attention measure *if* it could be measured in a manner that was robust and externally valid. It would not be enough to know that certain digital products affected performance on highly specific cognitive tests at a particular moment in time. On the other hand, if long-term cognitive indicators — or short-term metrics that were predictive of long-term “real world” performance — were found to be impacted by use of digital products, this would be important for consumers and policymakers alike.<sup>33</sup>

Such a potential measurement could address a consistent disadvantage discussed in earlier categories of potential attention metrics. If an attention metric were found that could measure cognitive impairment, this would necessarily entail a negative aspect of a digital product rather than a metric that could have either a positive or negative meaning, depending on the circumstances. There are, however, potential disadvantages. For starters, it could be that claims of “cognitive impairment” would either not be taken seriously by the public even if they had a substantial scientific basis, negating their value. On the other hand, it could be that the general public or policymakers in particular might overreact to such labels, in the sense of modifying their behavior or policymaking more than was rational of welfare-optimizing given the actual effects.

## Closing comments

The above enumeration is by no means an exhaustive listing of potential attention metrics, but merely an initial set of suggestions drawn from lived daily experience, anecdotal evidence, and related regulatory experiences. Also, there need not be a single kind of attention metric. The proposals in this discussion are largely complementary to one another, and even within single categories, diverse related attention metrics could be implemented more than one way.

In any case, if shown to be robust and externally valid, attention metrics for digital products could be used in a number of practical ways. For example, the attention metrics documented here could be similar in function to warning or information labels already familiar to ordinary people, such as those found on packages regarding the use of certain substances, such as warnings on alcohol packages but also nutritional labels on prepared foods. Such possibilities will be discussed further in Section III.

<sup>33</sup> It is important to recognize that such findings have not yet emerged from extensive study of the cognitive effects of technology, but this is in part due to the impossible conditions of studying the impact of technology on cognition. In particular, it is nearly impossible to do randomized controlled studies, and technology changes so quickly that the typical timescale of longitudinal study is too long compared to the rate of evolution of technical products.

Ultimately, this short list of potential attention metrics points the way forward to a rich experimental and theoretical space in which social scientists and digital designers alike could experiment with metrics to produce accurate and effective documentation for the attention costs of digital products.

## EXPERIMENTS TO MEASURE MARKETPLACE REACTION

An online behavioral experiment was undertaken to measure the preliminary feasibility of attention-based warning labels and the potential effectiveness of such labels. An online *vignette study*<sup>34</sup> was the methodology of choice. Participants were exposed to the description of a mobile app and asked whether they would download that app onto a new smartphone. The smartphone app was described on the basis of an extremely popular app.<sup>35</sup> At the bottom of the app description was an experimental treatment reflecting an attention related warning label, as drawn from one of the categories described previously.

The experiment was inspired by the likely critique that attention metrics are not necessary or useful for consumers. As demonstrated in this experiment, consumers likely would respond to the information provided by attention metrics, and might potentially avoid certain digital products if alerted to attention-related costs of such products.

## EXPERIMENT DESIGN & PROCEDURE

In this online survey experiment, participants were randomly exposed to one of five possible warning label treatments. There was a control treatment, i.e. no warning label, and alternate treatments corresponding to four potential attention metrics: *time*, *ad-aware*, *distraction*, and

34 “A vignette in psychological and sociological experiments presents a hypothetical situation, to which research participants respond thereby revealing their perceptions, values, social norms or impressions of events.” Vignette (psychology), Wikipedia, [https://en.wikipedia.org/wiki/Vignette\\_\(psychology\)](https://en.wikipedia.org/wiki/Vignette_(psychology)), last accessed Aug 5 2021. While vignette studies can seem highly stylized and far removed from real world conditions, they are often successful in predicting and explaining real world behavioral phenomena.

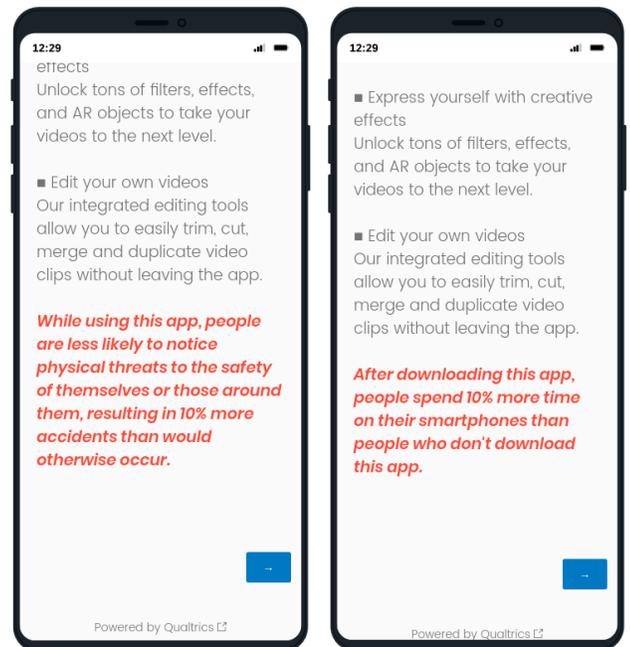
35 The app description was taken from the TikTok app’s description on the Apple app store, which as of the time of running the experiment held the international record for most downloads in a quarter, with 315 million downloads in the first quarter of 2020. M Sing (2020), “TikTok tops 2 billion downloads,” *TechCrunch*, Apr 29 2020, available at <https://techcrunch.com/2020/04/29/tiktok-tops-2-billion-downloads/>. After participating in the experiment, all research participants received a communication that the warning label did not describe an actual app.

cognition treatments.<sup>36</sup> The text correlating to each experimental treatment condition is shown in the table below.

TREATMENT NAME	TEXT
no-label	
ad-aware	After using this app, people are 10% more susceptible to advertising than are people who do not use this app, and as a result advertisers pay a premium to identify and advertise to people who use this app.
cognition	After using this app, people score 10% lower on a test of reasoning than do people who do not use this app.
distraction	While using this app, people are less likely to notice physical threats to the safety of themselves or those around them, resulting in 10% more accidents than would otherwise occur.
time	After downloading this app, people spend 10% more time on their smartphones than people who don't download this app.

Participants were randomly assigned with equal probability to one of the conditions. In each case, participants read the same app description and saw an app warning label in bold red text at the bottom (except in the case of no warning label).

Participants were asked to rate, on a 5-point Likert scale, their likelihood of downloading the app onto an empty new smartphone. They were then asked in counterbalanced<sup>37</sup> order about the utility and accuracy of the warning label, also on a 5-point Likert scale (and with the option to indicate that there was no warning label). Finally participants answered an attention check to identify what kind of warning label they had seen, if



36 The experiment also included a “blackbox” experimental treatment, which has been eliminated here for reasons of brevity.

37 Counterbalanced ordering means that all possible orders of a series of questions (that is, permutations of those questions) were presented, on a random basis, for all treatment groups. This technique wipes out potential ordering effects. This is important because in behavior experiments it can often be the case that the average answer to question A will depend on whether it is asked before or after question B.

any. This was followed by a series of exploratory questions relating to participants' smartphone use, including whether they were satisfied with their current amount of time spent using their smartphone. All question text and screen sequencing is shown in the appendix.

The study was conducted in August 2020 via a Qualtrics survey. A convenience sample of several hundred (N = 283) U.S. residents over age 18 was collected via Prolific.co.

The full *experimental protocol* is included in the appendix. Results related to demographic distribution of participants and performance on the attention checks are available in the appendix. The experimental data is available from the researcher upon reasonable request.

## RESULTS

### DOWNLOAD RESPONSE TO ATTENTION WARNING LABEL

After reading the vignette information, participants were asked how likely they were to download the described app. Below, the mean reported likelihood<sup>38</sup> to download the app is compared according to the treatment. Only in the no-label treatment does the mean self-reported likelihood to download the app exceed the midpoint (that is, the neutral point on download attitude). These results suggest, overall, a disinclination across the participant pool to download the described app. Yet not all treatments resulted in the same self-reported likelihood to download the app.

TREATMENT	MEAN REPORTED LIKELIHOOD TO DOWNLOAD (LIKERT SCALE 1–5)	N
ad-aware	2.45	60
cognition	2.51	57
distraction	2.63	57
time	2.79	56
no-label	3.1	53

<sup>38</sup> All references to the likelihood to download the app are self-reported stated likelihood to download the app. A much more complicated experiment would be necessary to measure the actual likelihood of downloading the app, as in this case the app description would have to be presented in the actual app store *and* researchers would need access to participants' phones to verify whether they had or had not downloaded the app. This demonstrates why the vignette method is a desirable alternative where field experiments are difficult or impossible.

Within the spectrum of mean reported likelihoods, there was significant variation, however. A *post hoc*<sup>39</sup> Wilcoxon signed-rank test<sup>40</sup> indicated that participants who read the ad-aware label had a significantly lower reported likelihood to download the app than did participants in the no-label treatment group ( $Z = 1170, p < .05$ ). A post hoc Wilcoxon sign-rank test also indicated that participants shown the cognition label were also significantly lower in reported likelihood to download the app than those in the no-label condition ( $Z = 1160, p < .05$ ).<sup>41</sup>

These preliminary results suggest that attention-oriented warning labels would affect consumer behavior at a relevant decision point in the consumption of a common kind of digital product, specifically a smartphone app. Thus warning labels, or alternate forms of conveying information about attention costs, should be considered both by technologists (such as large tech companies that run app stores for smartphones) and by regulators (such as the FTC) as they consider what kind of information is relevant and important to disclose to consumers in the interests of transparency, fairness, and even public health.

## JUDGED ACCURACY OF THE WARNING LABEL

After reading the vignette information, participants were also asked how accurate they believed the warning label to be. The judged accuracy of the warning label also varied significantly between different warning labels. Interestingly, the judged accuracy of a warning label was not necessarily correlated to the likelihood to download the app. In particular, the two labels that induced the strongest change in stated likelihood to download the app, the *ad-aware* and *cognition* conditions, had the highest and lowest mean judged accuracy ratings respectively. This result suggests that the influence of the labels on stated likelihood to download was not merely a function of perceived accuracy.

The perceived accuracy between the two labels that significantly reduced the stated likelihood to download (the cognition and ad-aware treatments) were significantly different (post hoc

39 No hypotheses were pre-registered and all comparisons are post hoc. This means that the analyses conducted were not committed to by the researcher in advance of running the experiment, as is typical in the case of exploratory research.

40 This is a non-parametric statistical hypothesis test, which means that it is a statistical test that does not assume a particular distribution of data that can test a specific hypothesis, namely whether two populations have the same means. In this case, the hypothesis testing is applied to whether two treatments have the same mean reported likelihood to download the app.

41 Reported p-values are not corrected for multiple comparisons. No comparisons other than those presented were made, in part due to the low number of research subjects..

Wilcoxon sign-rank test,  $Z = 871$ ,  $p < .001$ ), with the ad-aware label judged to be most accurate and the cognition label judged to be least accurate among all the treatments.

TREATMENT	MEAN REPORTED ACCURACY OF LABEL (LIKERT SCALE 1-5)	N <sup>42</sup>
cognition	2.42	55
distraction	3.19	57
time	3.41	51
ad-aware	3.47	60

More work is needed to understand this difference, but as a preliminary matter, this raises a number of important questions and considerations. First, it may be that attention warnings based on cognitive effects will be viewed with skepticism, even if they are scientifically validated and robust. If regulators or technologists were to pursue such a label, they might benefit from being particularly careful to document the scientific basis of such warnings to overcome prior skepticism from ordinary people. Second, the warning label judged to be most accurate, the ad-aware label, presents a category of information that could be particularly difficult to obtain because that information would require reporting of information from ad sellers or buyers (a logistical challenge)<sup>43</sup> and because such information may be regarded as a form of intellectual property (a legal challenge). However, the strong belief by research participants that such information was accurate suggests that such information could be the most effective if it benefits from higher perceived credibility by consumers.

## SPECIFICITY OF RESPONSE TO THE WARNING LABELS

Finally, the lifestyle question asking research participants whether they would like to spend less timing using their phone provided an opportunity to test whether reactions to the warning

42 Participants were given the option to indicate that they had not seen a warning label, as would be the correct response for participants in the no-label condition. Participants who indicated they had not seen a warning label (whether correct or not in that assessment) were not included in the participant count for this table. Interestingly the rate at which participants failed to realize they had seen a warning label varied slightly between treatments, suggesting possible differences in recall between the treatments. This should be investigated in future work.

43 This would present a logistical challenge to handle the reporting and processing of such information. This would also involve requesting information from parties whose interests are, at least in part, in conflict with the interests of consumers, by definition.

label treatments showed a link to consumer preferences. Specifically, by combining the lifestyle question about participants’ preference to spend more or less time using their smartphones and participants’ likelihood to download the app, it is possible to see whether participant response to warning labels about time is related to their stated desire to spend more or less time with their phones.

Thus we next examine, *within each treatment group*, the mean difference in likelihood to download between those participants who would like to reduce the amount of time they spent using their smartphones and those who would not. Interestingly, we can see that the biggest difference in mean likelihood to download the app between those who want to reduce their smartphone screen time and those who do not is precisely within the *time* warning label treatment condition.<sup>44</sup> This result suggests that people *do respond* to the *specific* content of the label in a way that correlates with their actual desires and hopes for lifestyle decisions related to use of technology.

TREATMENT	MEAN DIFFERENCE IN DOWNLOAD LIKELIHOOD BETWEEN WANTING AND NOT WANTING LESS TIME SPENT USING SMARTPHONE
time	-.68
no-label	-.17
cognition	.05
ad-aware	.13
distraction	.14

These results suggest that attention-related warning labels can affect consumer behavior at a relevant decision point *and also* that the warning labels affect behavior in a way consistent with a consumer’s own expressed preferences. This is important because it shows the warning labels work not because they inspire a generalized fear but because they correspond to elements of digital product quality that affect consumer decision-making because of how a consumer wants to live her life.

44 No additional statistical tests of significance were conducted due to the small sample size. A larger, representative experimental sample is left for future work.

## EXPERIMENT DISCUSSION

At least two categories of attention-related warning labels for a smartphone app inspired by attention metrics show a statistically significant reduction in interest in downloading an app geared toward entertainment.<sup>45</sup>

The cognition-related warning labeling makes use of existing and potential future research by psychologists who study the effects of technology on cognition. The effectiveness of a cognition-specific label in modifying likelihood to download a smartphone app suggests that cognition-oriented attention metrics could provide effective warnings, where warranted, to properly and holistically warn consumers of the dangers of certain digital products. The cognition label has this effect, despite receiving lower perceived accuracy ratings than the other warnings. It remains for future research to better delineate the mechanism by which a cognition-based warning reduces likelihood to download while being perceived as inaccurate. One possibility is that participants are reacting rationally because they find the potential harm quite worrying, even if they believe the likelihood of that harm actually occurring (that is, the likelihood of the label being true) to be quite low. In this case, they would be evincing a rational response to *harm in expectation*, which is a combination of the possible harm with the probability of the possible harm.

The ad-aware labeling treatment makes use of an existing and well-known market for attention from advertisers. This metric was effective in reducing the likelihood to download the app and also achieved the highest perceived accuracy. Hence, an ad-aware set of attention metrics appears most promising both in affecting market behavior and also in benefiting from high levels of credibility.

The results also show that people respond to the specific aspect of attention emphasized in a warning label. Specifically, the effect of a warning label related to the time consumed by an app was quite different for those people who wanted to reduce their smartphone screen time and those who did not.<sup>46</sup> This suggests that underlying individual *heterogeneity* can modify the effectiveness of different attention labels, as it should if this information is primarily designed

<sup>45</sup> As stated earlier, no additional statistical tests were conducted. Due to the small sample size, the researcher will expand the research in future work so as to enable meaningful statistical reporting for all proposed attention labels rather than only those that appear most promising in these initial results.

<sup>46</sup> Research into dissatisfaction with digital products that may relate to attentional harms should be surveyed and expanded. Such research — while seemingly only tenuously connected to attention — could in fact serve as inspiration for other forms of attention metrics by highlighting the concerns that would drive consumers to care about attention-related warning labels on digital products.

to inform consumers and enable them to determine their own best interests with fuller information about the quality of digital products.

## **EXPERIMENTAL CONCLUSION**

Attention metrics as described in this paper do not yet exist for use in academic research or policymaking. Yet, the experiment described here has shown that if such attention metrics were developed, they could significantly impact marketplace behavior, and, in some cases would benefit from surprisingly high levels of credibility among ordinary people even without the need for “consumer education.” Given the power of attention metrics to educate and warn consumers, there are many good reasons to push forward in developing such metrics for use in industry and academic research alike.

## III. What to do with an attention metric?

*Over the coming century, the most vital human resource in need of conservation and protection is likely to be our own consciousness and mental space. —Professor Tim Wu, *The Attention Merchants*<sup>47</sup>*

It seems likely that the attention metrics suggested and tested in Section II can and should be developed and validated.<sup>48</sup> It is also likely that the academic and industry communities can come up with other attention metrics beyond those contemplated in this work. Attention metrics are of interest to ordinary people, and their development should be pursued. We next move to the question of what could be done with such attention metrics.

In this section, a variety of potential uses of attention metrics from a law and policy perspective are briefly surveyed. This non-exhaustive list is organized in increasing order of intrusiveness, such that voluntary schemes associated with self-regulation are discussed earlier in the list, while some forms of direct governmental intervention are addressed later in the list.

As the reader will see, there are many potential use cases for attention metrics. These use cases could serve a number of high-level goals. The existence of attention markets could serve to increase transparency regarding digital product quality and so, ultimately, improve product quality by enabling consumers to be more efficient and selective in their choice of digital products.<sup>49</sup> Relatedly, attention metrics could serve as useful information for social policies, informing areas as diverse as public health, marketplace competition, or protection of vulnerable populations. Attention metrics could also serve as a measure of economic activity, itself useful to understand what is happening in markets and in society generally.

### ATTENTION BY DESIGN

One goal of attention metrics is to bring attention costs and attention depletion into greater focus, both for consumers and designers of digital products. For consumer welfare and

<sup>47</sup> T Wu (2017) *The Attention Merchants*, Vintage Publishing House.

<sup>48</sup> Candidate communities to develop such metrics include psychologists, economists, and technologists.

<sup>49</sup> This seems particularly likely since the experimental results in the previous section show that consumers likely would reduce their downloads of products based on certain kinds of information about the attention costs of those products.

practical digital interface design guidelines, the privacy by design movement offers inspiration for a movement that could fashion itself, in parallel, as “attention by design.”<sup>50</sup>

Attention by design would encourage a design paradigm in which user attention is treated as a scarce and cherished resource that should only be captured to the extent necessary to serve the legitimate purposes of a digital product, much like privacy by design encourages design choices that emphasize a user’s privacy right from the start of product creation and construction. In design, explicit choices must be made about attention, but from a welfare point of view it is likely undesirable that such choices always favor the capture of *more* attention, analogous to the way in which it is seldom in a consumer’s best interest for a product to capture more personal data. Attention by design would provide pushback against the incentives established by commercial impulses towards ever increasing attention harvesting and engagement maximization.

When looking at how attention by design might be operationalized concretely, one can consider some principles of the privacy by design movement. For example, the privacy by design movement recommends that design be “preventive, not remedial” and that there should be “respect for user privacy.” These and other principles associated with the privacy by design movement would directly translate to the separate concern of user attention.

The combination of robust attention metrics with an attention by design ethos would seem to provide numerous opportunities to ensure digital product quality improvements. Attention by design could promote a design process and quality standard that are holistic and preventive. Also, attention by design could benefit from association with privacy by design, with the latter already widely adopted by many organizations<sup>51</sup> and some legal systems,<sup>52</sup> thus benefiting from established institutional guidance and advocacy.

Of course, there are potential downsides, as with any potential intervention. An obvious concern with attention by design is that it offers rather nebulous guidance, and no obvious route to enforcement by parties concerned that principles are not adhered to. Also, in contrast to privacy, where more is always assumed (by certain advocates) to be better, no one can

50 Wikipedia, “Privacy by design”, [https://en.wikipedia.org/wiki/Privacy\\_by\\_design](https://en.wikipedia.org/wiki/Privacy_by_design), last accessed May 3 2021.

51 See e.g. IBM, “IBM Security and Privacy by Design”, available at <https://www.ibm.com/trust/security-spbdd>, last accessed May 1 2021.

52 FTC Report, Protecting Consumer Privacy in an Era of Rapid Change, March 2012, available at <https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf>, last accessed May 2 2021. Likewise, “data protection by design” is the topic of Article 25 of the European Union’s General Data Protection Regulation.

say the same for attention. Attention minimization (or maximization) will not always be the welfare-enhancing design choice, meaning that an attention by design scheme will face an even murkier set of objectives than does privacy by design. However, any ethical or legal framework is bound to have some shortcomings, and some use of attention by design would surely be an improvement compared to the current regime, in which designers are often encouraged to enhance user engagement without weighing the attention costs imposed on consumers against any potential benefits to consumers or producers of digital products.

## **LABELING REQUIREMENTS**

Product labeling is another possible intervention based on attention metrics, as was illustrated through the experiment presented in Section II. There is already a recent precedent for using labeling to clarify opaque quality issues related to a common and important case of digital products, specifically mobile apps. A privacy labeling requirement was adopted by the Apple App Store in December 2020, with the Google Play Store following suit in May 2021.<sup>53</sup> This example shows that labels are feasible and also desired by some portion of consumers.

There are numerous potential advantages to the use of attention labels on products as a source of product transparency and consumer education. Such labeling requirements could serve multiple purposes, including to directly educate consumers at points of sale.<sup>54</sup> Labeling requirements could potentially raise salience so that even facts already known to consumers would be presented with emphasis, at the time and place most likely to influence action. Labeling could also potentially introduce greater competition into markets that may currently lack desirable levels of competition. In this latter case, labels could serve as a transparent and objective basis of comparison for consumers who wish to compare digital products. Thus, labeling requirements could benefit individuals who otherwise face the toll of attention costs,

53 N Statt (2020), “Apple launches new App Store privacy labels so you can see how iOS apps use your data”, The Verge, Dec 14 2020, available at <https://www.theverge.com/2020/12/14/22174017/apple-app-store-new-privacy-labels-ios-apps-public>. Several months later Google announced that it would follow suit with its mobile app marketplace. S Perez (2021), “Following Apple’s launch of privacy labels, Google to add a ‘safety’ section in Google Play”, Tech Crunch, May 6 2021, available at <https://techcrunch.com/2021/05/06/following-apples-launch-of-privacy-labels-google-to-add-a-safety-section-in-google-play/>.

54 In the attention economies that characterize many digital products of interest for this work, “sales” is necessarily a broad concept that has more to do with the choice to use and thus give attention to a digital product rather than to a monetary transaction.

and could also potentially benefit digital markets as a whole by increasing transparency and therefore possibly increasing competition on important product quality attributes.

Of course, there are potential disadvantages or weaknesses of a labeling requirement. Labeling requirements have a much-debated history as to whether they are informative and whether they modify behavior in the desired direction.<sup>55</sup> Another issue is that the presence of labels does not guarantee the truth of the information they report. For example, even prior to introducing privacy labels, both the Apple and the Google mobile app marketplaces required some mandatory labeling with respect to apps for children, including whether advertisements (and what kind of advertisements) were shown to children.<sup>56</sup> Unfortunately, investigations have found that the reporting was dishonest in a large percentage of cases.<sup>57</sup> Finally, another concern is that, if mandated by the government, attention labeling would likely face challenges premised on freedom of expression, which has sometimes acted as a bar to governmental mandates on labeling. None of these challenges necessarily precludes the use of labels, but these concerns would have to be addressed in potential use cases.

## AUDITS

Much proposed U.S. legislation related to privacy and algorithmic accountability has included the use of auditing to look at the impact of certain digital technologies. While such audits have not recognized attention costs, attention costs could and should be integrated into any audit requirements that may eventually enter into law.

The value of an auditing requirement would be multifold. First, an auditing requirement would encourage producers of digital products to define the consumer-side attentional costs of the product by their own metrics, precluding the possibility of complaints about mismeasurement. This in turn could lead to voluntary efforts for standardization, as different technological sectors might take audit requirements as an opportunity to come to consensus definitions. Second, an

55 See e.g. D.E. Ho (2012), “Fudging the Nudge: Information Disclosure and Restaurant Grading,” 122 *Yale Law Journal* 574. But see S Shangguan et al (2018), “A Meta-Analysis of Food Labeling Effects on Consumer Diet Behaviors and Industry Practices,” *American Journal of Preventive Medicine* 56(2).

56 A Siddiqui (2019), “Apple issues warning to developers to stop incorporating ads and third party trackers in kids’ apps,” [digitalinformationworld.com](http://digitalinformationworld.com), Jun 4 2019. S Salim (2019), “Google reinforces new play store policies to safeguard children from inappropriate app downloads,” [digitalinformationworld.com](http://digitalinformationworld.com), May 31 2019.

57 A Siddiqui (2019), “Apple issues warning to developers to stop incorporating ads and third party trackers in kids’ apps,” [digitalinformationworld.com](http://digitalinformationworld.com), Jun 4 2019. S Salim (2019), “Google reinforces new play store policies to safeguard children from inappropriate app downloads,” [digitalinformationworld.com](http://digitalinformationworld.com), May 31 2019.

auditing requirement would also enhance transparency, both for consumers and regulators, contributing to diverse social policy areas, such as public health and antitrust law alike.

The disadvantages of auditing as a mechanism to regulate attention harvesting are related to those discussed for previous interventions. An audit would likely be a nebulous process potentially subject to manipulation, and auditing might prove effective only with buy-in from a firm's leadership or from a particularly powerful regulator. It might also focus parties too much on particular metrics to be fulfilled, rather than emphasizing a holistic scheme. It is worth recognizing that such concerns are related to a more general problem of “ethics washing” that some worry will occur with AI fairness issues.

As in the case of previous interventions discussed, none of the potential disadvantages of an auditing requirement negate the possible utility of such an intervention. Rather, as with all regulation and tech interventions, it would be necessary to carefully study this intervention both prior to and after it had been deployed to calibrate the intervention so as to ensure benefits to consumers, regulators, and participating firms alike.

## **CONCRETE AND PARTICULARIZED HARMS**

Attention metrics could also enable greater judicial notice of attention harms. Harms associated with digital products have faced difficulties in obtaining recognition in judicial proceedings, in part due to being considered insufficiently concrete or particularized for purposes of standing in federal court.<sup>58</sup> If attention metrics gained wide currency, they would provide quantifiable and routinized measurements that could prove instrumental in convincing the relatively change-averse U.S. judiciary to recognize non-traditional harms, such as those alleged in digital environments, as legally cognizable and sufficient to justify standing under current doctrine.

If attention metrics enabled sufficient documentation of attention costs so as to render them tangible (or otherwise sufficiently concrete even if judged intangible by the judiciary, as required under the relevant legal doctrine), this would pave the way for a wide variety of potential private actions by individual or class-action plaintiffs where legal harms occurred as the result of unjustified or otherwise illegal attention costs. In this way, the full availability

<sup>58</sup> A Nielsen (2021), “Technological Tangibility: A Route Back to Federal Court for ‘Intangible’ Harms”, draft available from author upon request.

of common law, statutory law, and typical remedies applied to other wrongful conduct could apply to engagement maximization and attention harvesting where appropriate, according to the facts of the case and the legal duties or rules at issue in a particular case.

The advantages of this system would be that attention harms, which may be commonly occurring in digital environments, would enjoy legal cognizability, and thus the externalities that may currently be imposed by certain digital design practices would have to be accounted for by firms that would otherwise face legal penalties in the form of individualized litigation. This would also enable legislators to pass laws about attention costs and harms that could include private rights of action, without fear that the judiciary would find statutory violations insufficient to recognize standing (a necessity for access to judicial remedies) for private parties.

There could of course be disadvantages. Anecdotally, many people seem to think the U.S. is already far too litigious a society, so that judges and ordinary people alike might find it undesirable to allow novel forms of harm into court. Also, a private right of action and private enforcement may be insufficient to bring about a desirable level of care for attention costs in digital design unless large class actions proved effective.

Again, the potential advantages likely outweigh the potential disadvantages. Having judicial notice of ongoing personal and social harms would undoubtedly be helpful in bringing attention costs to the fore, as private judicial action has proven so instrumental in so many other areas of consumer protection. What's more, attention metrics could help draw legislative focus to this problem, which has to date not received any coverage in proposed legislative solutions.<sup>59</sup>

## **DIGITAL ATTENTION TAXATION**

Taxation is a particularly useful regulatory tool because it can translate non-monetary costs and externalities, usually ignored by firms, into monetary metrics that matter very much for firms' ultimate objectives. Given a sufficiently robust and descriptive attention metric, it could

<sup>59</sup> For example, recently proposed federal legislation does not explicitly recognize problems associated with the use of algorithms for engagement maximization or attention harvesting. Instead, much focus in proposed legislation is devoted to problems of algorithmic bias or lack of efficacy. For example, neither the proposed legislation represented by the Algorithmic Accountability Act of 2019 nor that of the Algorithmic Justice and Online Platform Transparency Act of 2021 includes any measures that would rectify attention costs and harms widely documented in various algorithmic products and digital infrastructure design decisions. While these pieces of legislation do address real and crucial problems, they are far from complete in recognizing the various harms that need to be rectified in commonly used digital environments.

be possible to tax undesirable forms of engagement maximization and attention harvesting through the construction of attention-aware tax policy.

Taxation of attention harvesting activities could be seen as a form of consumption tax. The consumption of attention by digital products could be the act formally taxed, at a rate determined by an attention metric. Interestingly, attention taxation is not as unorthodox as it may appear under that name. In fact, new and emerging digital services taxes proposed and enacted in several European countries can be understood as attention taxes. In particular, the current schemes call for some approximation of how much value is contributed by users in a certain country<sup>60</sup> — value that would necessarily be created in relation to the attention that is “paid” to digital products. Attention metrics could provide an objective and robust way to operationalize this. Thus, using attention metrics to rationalize attention taxation could be a powerful regulatory policy that would make taxation more accurate and responsive to the economic realities of digital products, which often respond to the logic of an attention economy as much as to direct monetary incentives.

Of course, as with any policy intervention, there are potential disadvantages to such taxation (and it bears recognizing that design of taxation policy is notoriously difficult). Any kind of taxation is often perceived as being bad for business and also bad for innovation, leading tax reform to be a particularly wrought area subject to heavy lobbying. Another disadvantage is that taxation is a form of regulation, rewarding desirable behavior with tax expenditures and discouraging undesirable behaviors with tax penalties. Some distinction between desirable and undesirable behaviors would be needed for an effective development of rationales for decisions to tax digital products’ attention harvesting and engagement maximization activities, meaning that regulators would have to think carefully about what kind of attention metric to use for taxation and in what circumstances. Finally, taxation is an area where highly paid professionals are particularly deployed to find loopholes, such that unintended behavioral distortions are a possibility. Attention taxation would have to be carefully designed to ensure that such regulation did not backfire from a welfare perspective.

As with any regulatory intervention, potential downsides do not in and of themselves provide a reason not to act. In fact, by one interpretation, many European countries have in fact adopted some form of attention taxation in the form of new digital services taxes. The U.S. may very well consider following suit, in part aided by accurate attention metrics that can help respond to the realities of a digital economy.

60 Jim Stewart (2019), “User Value and Taxation of the Digital Economy,” available at <http://kluwertaxblog.com/2019/05/09/user-value-and-taxation-of-the-digital-economy/>.

## Non-monetary metrics and some concluding thoughts

The possibilities explored above represent only some potential opportunities to incorporate attention metrics into markets, regulations, and lawmaking for digital products. No doubt, there are additional options not explored here, as well as additional benefits or disadvantages beyond what this brief survey allowed. The optimal calibration of such interventions is beyond the scope of this paper. Regulation is difficult to get right, and sometimes regulators or consumer advocates must develop experience when moving into a new domain. Yet the need to do so should not defeat calls for recognition of an important form of potentially widespread consumer harm and manipulation. Rather, scholars, advocates, and digital product designers alike should recognize the importance of attention as a limited and highly personal resource, and work together to develop robust technological safeguards and transparency practices to benefit all stakeholders.

Attention metrics and the recognition of attention harms can also contribute to making progress on another problem that is dogging legal scholarship related to technology, namely that many efforts to quantify harms and preferences are grounded in financial metrics. There is a growing consensus that monetary valuations are highly manipulable with respect to measuring preferences and values that affect the human-technology relationship. For example, scholars have noticed that the monetary indicators used to “value” privacy may not make sense to ordinary people as a way of conceptualizing or valuing privacy.<sup>61</sup>

Likewise, there is a growing consensus that in the markets for digital products, monetary metrics may likewise not measure adequately the state of a market or consumer well-being, as has been noted in recent antitrust scholarship. It is time to move beyond monetary metrics when assessing the state of the world, most particularly for digital products and markets.

This work has been an exercise in imagining what could happen if we could take a very general construct in human culture and turn it into something actionable. The goal in doing so would be to match the analytical might and sophistication currently deployed in the service of

61 See literature review in, and results of, Aileen Nielsen. 2021. “Measuring Lay Reactions to Personal Data Markets.” In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3461702.3462582>

harvesting human attention and turn some of that to contemplating the cost of doing so.<sup>62</sup> Attention metrics could empower consumers and regulators to evaluate and respond to one form of technological harm currently entrenched in the design of many digital products. With the right incentives, the technology of the future can improve along many axes of quality, including those that measure costs to consumers, rather than merely the gains to producers of digital technologies.<sup>63</sup> Attention metrics are not the solution to all that ails tech, but they offer a promising route forward for bringing more humanity back into digital products.

62 D Baer (2013), “Why Data God Jeffrey Hammerbacher Left Facebook To Found Cloudera,” *Fast Company*, Apr 18 2013. “The best minds of my generation are thinking about how to make people click ads . . . That sucks.”

63 It’s worth recognizing that there are some efforts taken by industry to do this. Consider, for example, that TikTok introduced measures to help users recognize when they had spent long amounts of time on the app, although the specific details of when such warnings are triggered are not public. C Burke (2020), “Ever Spent Hours On TikTok Without Realizing It? The App Is Trying To Fix That,” *Bustle*, available at <https://www.bustle.com/p/tiktoks-new-screen-time-prompts-remind-users-to-take-a-breather-21816527>. I thank David Stein for suggesting this example.

# Appendix

## DEMOGRAPHIC DISTRIBUTION AND ATTENTION CHECK PERFORMANCE

Here are reported results relating to the demographic distribution of participants and their performance on the attention checks used to determine inclusion in the analysis of the data set.

The purpose of reporting the demographic distribution is to give the reader the opportunity to consider how representative the sample is to a given population of interest. As the reader will note, the sample was not a representative sample of the United States but did include a diversity of age groups and a relatively well sampled group by gender. Unfortunately, information about ethnicity was not collected.

The purpose of reporting the results of the attention checks is to establish that the attention checks were not unduly onerous and likely did not drive the results reported in the paper. As the reader will see, the rate of passing the attention checks was high and indicates that most participants were reading the vignette information reasonably carefully and so were included in the analysis.

### Demographic Distribution

The demographic indicators followed a distribution similar to what is commonly reported for online convenience samples. 54.3% of the sample identified as female and 44.1% as male. 1.3% of participants identified as non-binary and one participant preferred not to self-report a gender identity. 72.5% of participants identified as white. The age distribution is reported on the table below.

AGE RANGE	PERCENTAGE OF SAMPLE
18 - 24	29.7%
25 - 34	38.1%
35 - 44	20.2%
45 - 54	6.5%
55 - 64	4.2%
65 - 74	1.2%

## Data Selection

Of the 400 participants, 335 passed the comprehension checks and were so included in the data analysis, resulting in a passing rate of 83.5%. The analyses presented below were conducted on the portion of the sample that passed the comprehension checks, but the results are also robust to inclusion of all data. All analyses presented below are post hoc.

## EXPERIMENT FULL TEXT AND DESIGN

### [Screen 1]

With a billion users around the world, NikNok<sup>64</sup> is THE destination for mobile videos. On NikNok, short-form videos are exciting, spontaneous, and genuine. Whether you're a sports fanatic, a pet enthusiast, or just looking for a laugh, there's something for everyone on NikNok. All you have to do is watch, engage with what you like, skip what you don't, and you'll find an endless stream of short videos that feel personalized just for you. From your morning coffee to your afternoon errands, NikNok has the videos that are guaranteed to make your day.

We make it easy for you to discover and create your own original videos by providing easy-to-use tools to view and capture your daily moments. Take your videos to the next level with special effects, filters, music, and more.

- *Watch an endless amount of videos customized specifically for you*  
A personalized video feed based on what you watch, like, and share. NikNok offers you real, interesting, and fun videos that will make your day.
- *Explore videos, just one scroll away*  
Watch all types of videos, from Comedy, Gaming, DIY, Food, Sports, Memes, and Pets, to Oddly Satisfying, ASMR, and everything in between.
- *Pause recording multiple times in one video*  
Pause and resume your video with just a tap. Shoot as many times as you need.

64 The app description was taken from the Google play app store's description for the actual TikTok app.

- *Be entertained and inspired by a global community of creators*  
Millions of creators are on NikNok showcasing their incredible skills and everyday life. Let yourself be inspired.
- *Express yourself with creative effects*  
Unlock tons of filters, effects, and AR objects to take your videos to the next level.
- *Edit your own videos*  
Our integrated editing tools allow you to easily trim, cut, merge and duplicate video clips without leaving the app.

[ATTENTION LABEL HERE. TEXT SELECTED FROM TABLE BELOW DEPENDING ON WHICH RANDOMIZED EXPERIMENTAL TREATMENT A PARTICIPANT WAS ASSIGNED TO.]

TREATMENT NAME	TEXT
no-label	
ad-aware	After using this app, people are 10% more susceptible to advertising than are people who do not use this app, and as a result, advertisers pay a premium to identify and advertise to people who use this app.
cognition	After using this app, people score 10% lower on a test of reasoning than do people who do not use this app.
observation	While using this app, people are less likely to notice physical threats to the safety of themselves or those around them, resulting in 10% more accidents than would otherwise occur.
time	After downloading this app, people spend 10% more time on their smartphones than people who don't download this app.

[Screen 2]

Imagine you have just received a brand new smartphone without anything pre-loaded onto the device. How likely would you be to download this app?

Extremely likely — Extremely unlikely

**[Screen 3a]**

How useful did you find the **warning text** provided in the app store description you just read?

Extremely useful — Not at all useful, There was no warning text

**[Screen 3b]**

How accurate did you find the **warning text** provided in the app store description you just read?

Extremely accurate — Not at all accurate, There was no warning text

**[Screen 4]**

What did the warning label about the NikNok app indicate?

- Use of this app tends to decrease performance on cognitive reasoning tests.
- Use of this app is associated with increased risk of accidents due to digital distraction.
- This app consumes a high rate of “attention units” out of the quantity available each day to the average person.
- Use of this app is associated with an increase in time spent using a smartphone.
- People who use this app are a desirable audience for advertising companies.
- There was no warning label.

**[Screen 5]**

Are such warning labels provided in the app store?

- Yes
- I'm not sure
- No

*(If answered yes)*

What kinds of warning labels are provided in the app store? Please briefly describe.

**[Screen 6]**

How many hours a day do you spend using your smartphone?

**[Screen 7]**

Please upload a screenshot indicating your hourly usage of your phone.

You will receive a monetary performance bonus for doing this, but completion of this task is not required. Also, please be sure to block out any identifying information in the screenshot.

For sample Apple device instructions, [click here](#). For sample Android device instructions, [click here](#).

**[Screen 8]**

Which statement below best describes your feelings about your smartphone use?

- I would like to spend more time using my smartphone.
- I would like to spend less time using my smartphone.
- I am happy with the amount of time I spend using my smartphone.

**[Screen 9]**

What kind of phone do you have?

- iPhone device
- Android device
- I don't have a smartphone

**[Screen 10]**

How old are you?

What option(s) best describe(s) your ethnic identity?

What best describes your gender identity?

**[Screen 11]**

Anything you wish to share about the study?

## About the Author

Aileen Nielsen is a Fellow in Law & Tech, at the Center for Law & Economics, ETH Zurich. She can be reached at [aileen.nielsen@gess.ethz.ch](mailto:aileen.nielsen@gess.ethz.ch). She is also the author of the author of *Practical Fairness: Achieving Fair and Secure Data Models*, the first trade book on mathematical methods and computer coding techniques to enhance algorithmic fairness in real world products.

## Acknowledgments

The author thanks workshop participants at UC Berkeley's Center for Long Term Cybersecurity and the Information Law Institute's Privacy Research Group at NYU School of Law for helpful comments on this work. The author particularly wishes to thank Ido Sivan-Sevilla, Shir Etgar, Sebastian Benthall, and David Stein for helpful comments and suggestions on earlier versions of this work, and Steve Weber, Ann Cleaveland, and Chuck Kapelke for insightful feedback specific to the white paper manuscript. Finally, the author thanks Marine Jorio for excellent research assistance.



**CLTC**

Center for Long-Term  
Cybersecurity

---

UC Berkeley

Center for Long-Term Cybersecurity

[cltc.berkeley.edu](http://cltc.berkeley.edu)

@CLTCBerkeley