KATHARINE ROLLER: Good morning, and welcome to "Bringing Dark Patterns to Light," the FTC's workshop on dark patterns. My name is Katharine Roller, and I'm an attorney from the Midwest regional office at the FTC. On behalf of the entire FTC workshop team, we're delighted that you're joining us today via our live webcast.

Before we begin our program, I have a few administrative details to cover. First, a video recording and transcript of these proceedings will be available on our workshop web page shortly after the event. Our intent is to create a lasting resource for anyone who's interested in this important topic.

Second, as with any virtual event, we may experience technical issues. If these occur, we ask for your patience as we work to address them as quickly as we can. We will let you know if there are going to be any significant delays.

Third, we'll be accepting audience questions through our dedicated email address, darkpatterns@ftc.gov. Due to time constraints, we may not be able to get to all of the questions. But we will review every question that we receive. We are also seeking public comments until May 29 on the topics discussed today. Take a look at our event web page for information on how to submit comments.

Finally, please join us on Twitter. Our Twitter handle is @ftc. And we'll be tweeting using the hashtag #darkpatternsftc. And now I have the great pleasure of introducing our first speaker, the acting chairwoman of the FTC, Rebecca Slaughter. Prior to joining the FTC as a commissioner, she served as Chief Counsel to Senator Charles Schumer of New York.

Earlier in her career, she was an associate in the DC office of Sidley Austin LLP. She brings to the commission more than a decade of experience in competition, privacy, and consumer protection. Along with advocating for consumers, particularly those traditionally underrepresented and marginalized, she strongly supports working families and work-life balance. Welcome, Acting Chairwoman.

REBECCA KELLY SLAUGHTER: Thank you, Katharine, and thank you for all the work that you and the rest of the team have done to put together this important workshop today. On behalf of the agency and my fellow commissioners, I'd like to welcome all of our audience to today's virtual workshop, "Bringing Dark Patterns to Light." As the title suggests, today we'll be examining, through a series of presentations and panel discussions, so-called digital dark patterns, user interface designs that manipulate consumers into taking unintended actions that may not be in their interest.

Dark patterns are part of a larger system of deceptive and manipulative practices that we see growing all too rapidly in the commercial digital surveillance economy. We're fortunate today to hear from a highly distinguished group of consumer advocates, academics, and designers, and also from Senator Mark Warner of Virginia and Representative Lisa Blunt Rochester of Delaware, who have introduced legislation in the Senate and House respectively that would further empower the FTC to address the harms of dark patterns. I'm pleased to welcome all of you.

We know that, in today's digital marketplace, data is currency. And we increasingly see companies using dark patterns to manipulate people into giving up their personal data, which is then sold, aggregated, analyzed, and used to target advertising and manipulate future purchases and behavior. This conduct can be particularly pernicious when orchestrated by large platforms who may use data collected through dark patterns to enhance their own market power.

We also see dark patterns used more directly to manipulate consumer behavior. Recently, the FTC has sued companies for forcing people to navigate a maze of screens and choices to cancel negative option subscription plans in ABCMouse, using inconspicuous dropdown links and autoscroll features to hide the material terms of virtual rent-to-own transactions in Progressive Leasing, and sneaking additional unwanted products into people's carts without their knowledge or consent in [INAUDIBLE]-- dark patterns all.

Whether it is expanded data collection or specific behavioral manipulation, the motivation behind the use of dark patterns is often the same-- enhance a company's bottom line and market position in ways that may be harmful to its users or contrary to their intent. We need to understand the technologies and information asymmetries that are driving the use of dark patterns, the myriad harms that dark patterns can cause to consumers and competition, and consumer's ability or inability to detect and avoid these harms. Fundamental to understanding all of these details will be an analysis of the business incentives that drive the development and deployment of deceptive and manipulative technological tools.

It's also crucial that we look at the impact dark patterns are having on different communities, especially those that have been and continue to be disadvantaged and marginalized in both the market and our broader society. Separately, we must consider how dark patterns affect people who may potentially be more susceptible to certain persuasive design elements, due to, for example, still-developing brains in the case of children, or lack of familiarity and exposure to digital technologies, which is an issue for some older folks.

Finally, although the FTC can and should continue to aggressively pursue companies that use deceptive, unfair, or otherwise unlawful design tactics, it's important that we consider the challenges that the scale, sophistication, and personalization of dark patterns pose to our enforcement efforts. How can the FTC best target and prioritize our enforcement resources? Could our efforts be strengthened by the adoption of additional rules, guidance, or other policies?

In grappling with these questions, the FTC has worked and will continue to work closely with our state and international partners, as well as with the advocacy community, all of whom are doing important work in this area, as you'll hear later today. And we are, of course, grateful to the leadership of elected officials, like Senator Warner and Congresswoman Blunt Rochester, who are working hard to ensure that our statutory framework keeps pace with and enables our enforcement efforts to keep pace with swiftly moving developments in the market. My hope is that today's workshop will help the FTC chart a path forward on these pressing questions.

The time for us to get this right is now, as design interface is being fully optimized. We need comprehensive understanding and a comprehensive plan of action. If enforcement and regulation are unable to check the power of unsupervised algorithms, consumers will face deception by design. And for companies who may be listening today, I have this to say. If you have to resort to design tricks to obtain consumers' consent or to tracking or charges, you ought to think twice about your business model and how your design choices align with the principles of honesty and fairness. With that, I want to welcome Senator Warner, who will be delivering his remarks via video.

MARK R. WARNER: Hi, I'm Virginia Senator Mark Warner, and I want to thank the FTC for holding this very important workshop today. It's addressing an issue that's very critical to me and something I think is absolutely essential that we get right. My background before I came to the Senate was in technology. And I am a big believer that technology can empower our lives.

But I think we've also seen, recently, particularly when it comes to social media, that there is a dark underbelly in social media, on Facebook, on Instagram, on YouTube, on Twitter. And we have unfortunately seen that consumers are too often taken advantage of. It is absolutely critical that the FTC step up and figure out ways to make sure that consumers online are protected.

One of the ways that we see this abuse take place is with deceptive tactics, which have been called dark patterns, where consumers, basically, looking for information or looking for an answer, are basically greeted with screenshot after screenshot, where virtually the only option is to agree. You have to go through a series of hoops and hurdles to ever kind of get out of that page, to say no, to say, I don't want to opt in and opt out. So I've set up legislation, bipartisan legislation, called the Detour Act, that would prohibit major social media platforms from using these deceptive patterns, these dark patterns.

We've seen certain states start to adopt these rules. Now, I think bipartisan legislation like this is extremely important to take place. But I think, in the interim, and if Congress, which sometimes has not been all that good at getting its act together and acting, I believe, under Section 5 authority, that you can put in place a regulatory structure to try to protect consumers and prohibit this level of deceptive practices. I know some of the social media firms will fight this, because this goes at the heart of how they can gain more eyeballs and ultimately garner more information about all of us, that they can then sell and market.

I have nothing-- I think there's nothing objectively wrong with that. But we need to make sure that consumers are not tricked into giving up that information. And too often, these patterns are the kind of very deceptive activities that trick consumers into giving up information that then these platforms can then market. We've got to sort this out.

This is an issue that is extraordinarily complex for many of my colleagues. And while I'm hopeful that our Detour legislation, bipartisan, can be advanced into law, I think you at the FTC and consumer privacy experts can help us sort this out and make sure, whether we do this through legislative action or Section 5 authority, we make sure that consumers' experience on the web is one that safe, secure, and allows people to protect their privacy. So good luck with this workshop. I look forward to hearing the results and how I can incorporate some of your ideas into our legislation.

KATHARINE ROLLER: Thank you, Senator Warner, for your remarks. Now we'd like to welcome representative Lisa Blunt Rochester.

LISA BLUNT ROCHESTER: Thank you. Good morning, everyone, and thank you especially to Acting Chairwoman Slaughter for the work that you are doing, for the introduction, to the entire FTC for putting together today's event, and to my colleague, Senator Mark Warner, for his tireless work in the Senate to help combat dark patterns. I'm really thrilled to be joining so many policy experts for today's conversation about bringing dark patterns to light. I was sharing with the chairwoman earlier that I actually was googling, and came across some dark patterns, saw that this actually was happening, taking place today, and really wanted to be a part of it. Because I know that, together, we can come up with real solutions.

And while this is an issue of technology, privacy, consumer advocacy, it's also personal to so many Americans. Just this weekend, I was reminded how it touches us on a day-to-day basis and how prevalent dark patterns are in our everyday lives. So we say this thing in Congress, I'll take a point of personal privilege.

I'd like to take a point of personal privilege and just share with you that one of my hobbies that I've developed over the years is birdwatching. And my mother knows this. She knows it's a hobby that I like and that provides me some respite. And so earlier this week, she sent me an app that supposedly helps to track and catalog birds.

Maybe, like most people or unlike most people, I'm a little suspicious these days of apps and websites and all of that. And it could be the work that we're doing. But I decided to go back and just read the reviews on this app, and found that, once you download it and sign up for the app, you begin a seven-day free trial. Once the trial expired, you're then charged a $39 monthly fee.

Now, free trials are a common enough experience that we've all had. But the real issue, as I continue to read into the reviews, was that canceling your trial wasn't nearly as easy as signing up for it. There was no obvious button to cancel the trial, no prominent link to navigate to. And the users were left guessing if deleting the app altogether would cancel the subscription as well.

That's dark patterns. That's what we see in face every day. And that, for me, is why addressing the issue of dark patterns is so important. It's about protecting our consumers, protecting our children, and protecting our communities, whether it's that birdwatching app that locks consumers into a subscription they did not ask for and do not want, autoreply features on certain websites that could expose children to content they shouldn't see or keep their attention for hours at a time, or cookie collection interfaces that are intentionally difficult to navigate.

We know that all of these practices, whether targeted at consumers, our kids, or entire communities, are at the core of a simple principle-- collecting as much data as they possibly can. And we know that, for these tech companies, data equals dollars. And while there's nothing inherently wrong with companies making money, there is something wrong about those companies intentionally manipulating users to extract their data.

That's why Senator Warner and I introduced the Detour Act last Congress and plan to do so again this Congress. What the Detour Act does is prohibit online platforms from using those intentionally deceptive interfaces to extract data from online users. The data that these companies want belongs to the user, and it should be up to the user to make informed decisions as to whether or not they're willing to give that data to these companies.

I also believe at the core of this issue is the need for tech companies, especially the largest platforms, to take meaningful responsibility for their practices. It shouldn't be the case that companies, like Facebook and Twitter, are manipulating users in a way that we wouldn't permit any other company of good standing to do to their consumers. And it shouldn't be the case that, if these companies continue to operate in this manner, that the federal government does nothing about it.

Before the pandemic, I had a chance to travel to Europe as a part of a bipartisan congressional delegation and talk to EU Parliament members who are grappling with the very same questions about how to address these deceptive practices as a matter of policy. So we know that this issue stretches beyond our borders and is one that the international community is looking to address as well. When the CEOs of Facebook, Google, and Twitter came before my Energy and Commerce Committee last month, I told them why I am so passionate about combating these dark patterns.

It's because our children, seniors, veterans, people of color, even our democracy is at stake here. And I asked them to commit to ending these harmful practices and being open, honest, and transparent with their users. Yet, despite all of the evidence that we've seen of these dark patterns and despite all that is at stake if this issue isn't addressed, those CEOs wouldn't provide a meaningful commitment to ending these harmful and manipulative practices.

So in the absence of responsible action from these companies, it is my firm belief that Congress should act. There's simply too much at stake. And that's why working with allies like Senator Warner, I am continuing to fight to bring dark patterns to light and fighting to end them once and for all. I thank all of you for your willingness to come together and address this. And I look forward to working with you. And thank you for having me here today. Let's get this done.

MIRY KIM: Thank you very much, Congresswoman Blunt Rochester, for your remarks. And good morning, everyone. My name is Miry Kim, and I'm an attorney in the division of Marketing Practices. My colleague, Min Hee Kim, is an investigator in the Office of Technology Research and Investigation. Min Hee and I be moderating the very first panel of the day, "What Are Dark Patterns and Why Are They Employed?"

Our panel will explore the definition of dark patterns, the characteristics of dark patterns, the prevalence of dark patterns, and the factors and incentives that give rise to dark patterns. As time permits, we will try to answer any questions you might have. So please feel free to submit your questions on dark patterns ftc.gov. I'd like to begin with a brief introduction of our esteemed panelists.

But I would encourage everyone to review their bios in more detail, as they have very impressive backgrounds. So in order of their presentations, which we will get to in just a minute, I'd like to begin by introducing to you Dr. Harry Brignull, who is joining us all the way from the UK. Harry is a user experience specialist, or UX design expert, who coined the term dark patterns. He runs darkpatterns.org to raise public awareness to manipulative and deceptive design practices.

Next, we have Dr. Arunesh Mathur, who is a post-doctoral research fellow at the Center for Information Technology Policy at Princeton University. His research examines the societal impact of technical systems through an empirical lens. Next, we have Miss Johanna Gunawan, who is a doctoral student in the Cybersecurity program at Northeastern University. Her research investigates user privacy, consumer protection, and tech policy issues.

Then we have-- and joining us all the way from Stockholm is Miss Kat Zhou. Kat is a product designer at Spotify and the creator of the Design Ethically project, which helps teams forecast the consequences of their products. And last but not least, we have Dr. Katharina Kopp, who is the deputy director of the Center for Digital Democracy, where she explores how commercial data practices and technologies adversely affect individuals, groups, and society.

Thank you, panelists, for joining us today. But before we begin, I have to note at the outset that, on behalf of myself, Min Hee, and the panelists, the views that we're about to express today are our own views and do not necessarily reflect the views of the commission or any one particular organization or company. So with that, I'm going to turn it over to Min Hee Kim for presentations.

MIN HEE KIM: Thanks, Miry. So first up, we have a brief presentation from each of our panelists to sort of help the stage for our panel discussions. So Harry, if you would, please tell us about dark patterns.

HARRY BRIGNULL: Thanks. OK, we might as well get started. The first thing people normally ask me is, what is a dark pattern? So I might as well start there. A dark pattern is a manipulative or protective [INAUDIBLE] in software that gets users to complete an action they would not otherwise have done if they had understood it or had a choice at the time. For example, if you have a button that functions as a yes when it's clicked, but through the use of placement, [INAUDIBLE], and trick wording, it appears to say no, then many users are going to be caught out.

If you're wondering where the term came from, it's based on the idea of a design pattern. A design pattern is a common reusable solution to a problem. There's also something called an anti-pattern, is a common mistake. In 2010, I realized there was another type of pattern, which is [INAUDIBLE] deception [INAUDIBLE]. So I gave it the term dark pattern. It's meant to communicate the unscrupulous nature and also the fact it can be shadowy and hard to pin down.

So people can end up purchasing things, sharing things, or agreeing to legal terms without intending to. People can try but fail to do the thing they set out to do because it's been designed to be difficult. For example, trying to cancel a premium subscription when you're called to call a phone line during working hours, to then have a rep try and talk you out of it for 10 minutes before you're finally allowed to leave. This is sometimes referred to as sludge, a high friction [INAUDIBLE] experience that, by it's nature, causes people to become fatigued and give up.

Another type of trick involves persuading people to do things by taking advantage of cognitive biases. For example, if a website tells you that other people are currently looking at the same item as you, that's tapping into your social approval bias. If it tells you that the item is almost sold out, that taps into your scarcity bias. It if it shows a discount, that taps into something called anchoring.

And of course, if all those things are actually true, then that's honest marketing. In fact, it's really useful for [INAUDIBLE] to know that information. But if any of those things are lies, then it's a dark pattern. But of course, the problem is that language can be very ambiguous.

If a retailer says they have a low stock, but they have a trailer [INAUDIBLE] that [INAUDIBLE] get unload it, are they telling the truth? But what about if the trailer hasn't yet arrived, but it's about to pull into the parking lot, or if it has already been unloaded, but the store manager hasn't yet signed the receipt or [INAUDIBLE] barcode [INAUDIBLE]? The definitions matter. And human language is very well-suited to vague definitions of things.

The recipe for dark patterns involves a conversion of applied psychology, A/B testing, and user interface design. [INAUDIBLE] is evolving really rapidly. So if we look back, say, 20 years, this sort of knowledge was only known by a privileged few. And now it's commonplace. So the genie is out of the bottle now. We can't put it back in.

When I first defined dark patterns, I was quite naive. I thought they could be eradicated by shaming the companies that use them and by encouraging designers to use a code of ethics. But the fact we're all today means that approach didn't really work.

I'll give you a specific example. There's a research paper called [INAUDIBLE] by Blake et al. 2020. In that paper, some researchers work with a ticketing website to look at the effect of hidden fees versus upfront fees. And they apparently involved several million users. It's the largest test of dark pattern that's ever been published. But the users who weren't shown the ticket fees up front ended up spending about 20% more money and were 14% more likely to complete [INAUDIBLE].

And that's a huge impact. Imagine if you ran a business and you could press a button to get your customers to spend 21% more, it's a no-brainer. Of course you'd press it. And that's the reason why dark patterns exist. They're effective, and they're [INAUDIBLE].

The reason why we're here today, of course, is that they also lead to outcomes of harm, [INAUDIBLE] spending money they didn't intend to, they share things they didn't intend to, or they get tied into agreements that they didn't know existed. To finish up, I'm going to tell you an old story. Perhaps you've heard it before. It's the old myth of the king's shilling.

So the act of touching this coin, this shilling, was taken as a binding agreement to join the English Navy in the 1700s, or so the myth goes, and you'd have no coming back. So as the story goes, they would put the coin into, say, a beer tankard without them knowing. And when it touched their lips as they took a sip, they would be [INAUDIBLE] into the Navy.

Websites and apps should not be able to [INAUDIBLE] consumers with [INAUDIBLE] deceptive and manipulative tricks. On that thought, I'm going to hand over to Arunesh Mathur, who's going to show me some different examples of different types of dark patterns that are used in the world today. Thank you very much.

MIN HEE KIM: Thanks, Harry. Arunesh, please go ahead.

ARUNESH MATHUR: Thank you, Harry, for this background on dark patterns. So since this panel is all about what dark patterns are, I'm going to focus on defining dark patterns, and I'm also going to show you some examples.

Moving past the title slide, a lot of what we know about dark patterns comes from taxonomies that researchers have carefully curated over the years. These taxonomies are terrific, and I highly recommend you to check them out. And I'm sure many of you have seen them before. These taxonomies contain many examples of dark pattern user interfaces.

Next-- but in addition to illustrating all kinds of examples, a central contribution of the literature on dark patterns is that it offers a definition of dark patterns, something that Harry just mentioned. It tries to tell us what a dark pattern user interface exactly is. In a recent article, we analyzed many of these definitions.

And we see that there's a lot of variation in how researchers define dark patterns. Some definitions highlight the characteristics of the interface, while others describe how the interface exactly is trying to influence users. But we know that even when the definitions seem to agree, they're ultimately highlighting various different aspects of dark patterns. And that makes it really hard to capture these tricky interfaces in one definition.

So how do we go about defining a dark pattern then? Next-- instead, a different perspective on dark patterns, we think, is to look at them through the lens of their attributes. These attributes describe how the dark pattern is attempting to influence users. And what we see is that there is a set of six attributes, which together describe all the dark patterns that researchers have been curating in the taxonomies.

Next-- going through these attributes, we note that there are some dark patterns that are deceptive, meaning that they induce false beliefs. And this is something that may be familiar to many of you. One example here is a countdown timer on a shopping website that you might not know is inconsequential to what is being advertised, meaning, the countdown timer either goes away or it resets when it finishes. But the offer that is being advertised is still available.

Next-- there are some dark patterns that are information-hiding, meaning they delay or they hide important information from users. A canonical example here is the hidden fees dark pattern. where important charges are only divulged to users after they've spent a lot of time selecting and finalizing the product that they want to purchase.

Next-- some dark patterns are asymmetric, meaning they make certain choices that users want to access, and also which are disadvantageous to the designer, really, really hard to access. One example here is the cookie consent dialogue, which places the option to accept cookies right in front, but the option to deny cookies behind, perhaps, several different screens.

Next-- some dark patterns are covert, meaning users likely do not understand how they're being influenced. An example is a popup may see on many websites that ask you for your email in exchange for a discount, maybe $25. But then it immediately asks you for your phone number too. And if you're a user, it may seem like you need to enter both of these pieces of information to obtain the discount. But in reality, just one is enough. And most users are very unlikely to realize that.

Next-- some dark patterns are differentially treating of users, meaning they disadvantage and treat one group of others differently from another. One example is the pay to win dark pattern that we see in many video games, where the only way you can get ahead in the game is by purchasing features in the game and not by your own skill.

Next-- and finally, some dark patterns are restrictive, meaning they eliminate certain options from the user interface altogether. For example, in this interface here, in order to sign up, you need to agree to the terms and conditions, as well as the promotional offers and marketing emails in order to proceed. You cannot agree to one but not the other.

Next-- but even within these attributes, we see two themes emerge that describe dark patterns. Next-- dark patterns that are deceptive and information-hiding, and perhaps those that are most familiar to this audience, are all about manipulating the information [INAUDIBLE] users, either by deceiving them or by delaying the necessary information to them. Next-- and on the other hand, dark patterns that are asymmetric, covert, or deferentially treating of users, and restrictive, they modify the decision space for users, and ultimately influence how users make choices.

Next-- and here we see these two themes emerge. Next-- so to answer the question, what is a dark pattern, well, we see that dark patterns ultimately modify the choice architecture for users. Next-- but there is not one single definition. It's, in fact, two themes that bind the entire dark pattern scholarship together-- next-- patterns that modify the choice set and the decision spaces, or they go about manipulating the information that is available to users. Next-- and in fact, an observation that these themes and not one definition captures dark patterns parallels Daniel Solove's now seminal observation for privacy, that privacy is not a single concept, but a family of related concerns.

Next-- so that's a very short summary of what dark patterns are, describing them, defining them. And of course, the discussion about patterns cannot take place without the discussion of the content. And you'll see that in panel 2. But if you're interested in reading more about the ideas that I just presented here, please check out our article, "What Makes a Dark Pattern Dark." And here is a short reference to it. Thank you so much.

MIN HEE KIM: Thank you, Arunesh. And we're actually going to ask you some more questions about that in the discussion portion. For our viewers out there, there is a slider bar in the top right corner of the slide. So if you would like to make them a little larger, you can adjust them yourself there. And then next, we're going to have Johanna tell us about different modalities. Johanna?

JOHANNA T. GUNAWAN: Yeah, great. Thank you so much, Arunesh, for that excellent and thorough overview of dark patterns. I'll echo the sentiment that there's an amazing body of literature to look into to learn more about these definitions.

So my colleagues and I at Northeaster are a team of computer science and policy researchers. And dark patterns represents a confluence of our various interest and privacy and consumer protections. We wanted to further the discussion by investigating dark pattern prevalence across different modalities or versions of the same web service. Next slide, please.

So as users, we instinctively know that interfaces differ across modality. How we experience a service or web service depends on how we access it. And web designers encounter this daily by working across a variety of screen sizes and device functionalities in order to deliver their product.

Our team was curious as to how dark patterns might be expressed across these different interfaces. And to test this, we ran a study of popular web services across their app, mobile browser, and desktop browser versions. Specifically, we were looking to see whether dark patterns differ in quantity, which would be the number of observed patterns per service or per dark pattern across modality, as well as qualitatively, which would be the types of dark patterns we might observe in each service.

So what did we find? Next slide, please. We ultimately found the dark pattern count is frequently higher in apps than in websites, both when you look within a service or across types of the dark pattern. And we find this especially concerning, especially for the asymmetry attribute that Arunesh outlined just minutes ago, both for asymmetry between the user and the service provider, but also asymmetry or rather an equivalence across these modalities. Next slide, please.

So this example takes a look at missing options for account deletion, which is, essentially, a dark pattern in which users aren't able to extricate themselves from a service as easily as how they entered it. And this pattern has been labeled in a number of different ways in the literature and some of the names I've provided on the slide.

So I know what you're thinking. The screenshot does show an option to delete an account. It's right there highlighted in red so it's definitely not missing. And that's true, but it's only true for this modality. This screenshot is from a hiking service's desktop website. And here, users can, in fact, delete their account. But when we visited the same Edit Profile page within the app and mobile browser modalities, we could see the options on the left to change your first name, last name, et cetera, but we didn't find this Delete Account option at all.

This finding is relevant for things like the right to delete or for opt in and opt out regimes, but generally, for providing users more meaningful control over access to their data. We were able to create accounts in all three modalities, so it's odd that we weren't similarly able to delete our accounts across all three. Next slide.

But even when controls are provided to the user, dark patterns can still interfere with user agency. The settings context example in these screenshots is for the aesthetic manipulation or interface interference dark pattern categories, which is when the format of an interface, or the way things are presented, can confuse, trap, or deceive users, and otherwise interfere with their goals.

These two screenshots are from the same booking service. The left is from the app modality and the right is from a mobile browser. An astute scholar of dark patterns may notice a number of different things going on here. But for time's sake, I'd like to focus mostly on how there are different settings available in these two modalities, even though they're for the same service and for the same group of settings. And they also have dissimilar levels of control between the modalities.

For example, in the mobile browser, users are able to control both their email and their push notification options. And to some extent, they can also use the Select All button to speed up the process and reduce the effort of managing their settings. However, in the app versions, users don't even have access to email preferences. And neither version allows users to turn off all of their notifications in bulk, which requires the user to individually toggle each content type.

Even so, the mobile browser interface provides a slightly richer experience of control, as well as more options for the end user. So in sum, both of the examples from these two slides, in apps, we find that in apps, they don't often provide equivalent user agency options as the other two modalities. So what does this kind of behavior mean for dark patterns? Next slide, please.

Well, interfaces are going to continue to differ. But these learnings provide opportunities to equalize experiences across modalities and reduce asymmetry, not only for the three modalities that our study looked at, but also for things like IoT and voice assistance and other device types in the future. Firstly, we think that user agency features, like certain settings and the ability to delete an account, should be available across all versions of a service. There's no excuse for not allowing users to leave a service in the same location they signed up for it.

Secondly, dark patterns can impact different user groups to enable outcomes. Surveys have shown that not all internet users own multiple devices, and lack of access to a device shouldn't subject users to different experiences of user agency. And lastly, on the policy and enforcement side, terms of user privacy policies that discuss compliance at a service level rather than a modality level may not provide a complete picture.

For example, if an audit only examined the desktop site in our first example, they might think that service is compliant and provides meaningful methods for users to delete their accounts, when the reality for mobile users would be quite different. All in all, a comparative look at dark patterns reveals inequal experiences across versions of a service. We can learn from these findings and work towards resolving these inequalities. Thank you.

MIN HEE KIM: Great, thank you, Johanna. I think that was helpful. And we'll follow up in our discussion portion, but I think some of the things that Johanna's mentioned here will also be touched on by our panel 3, which will be discussing the effects specifically on people of color. So that will be very interesting. Next up is Kat with a designer's perspective.

KAT ZHOU: Wonderful, thank you. My name is Kat Zhou, pronouns she/her/hers. And I want to start by thanking the moderators, Min Hee and Miry, for this opportunity, and also want to thank the previous speakers for so clearly elucidating what manipulative design practices are within our industry.

Now I'll be diving into the inner workings of tech companies and giving you an inside glimpse of how product teams and designers are incentivized to incorporate anti-pattern designs within their apps and their websites. Next slide, please. So how on Earth did we get here, where we've seen this proliferation of designing for consumers to get hooked? In the next slide, I've pasted some screenshots of various headlines, of news articles, and Medium posts written about an area of design that's been trending for a while.

And in the following slide, I've highlighted the common word that is a central theme of these headlines. Growth is what drives designers, product teams, companies, the tech industry, and the entire racial populace system. You can see the allure of growth shaping how venture capital firms make decisions in how they invest. You see the allure of growth shaping how board members at companies vote. And you see the allure of how growth shapes how product teams might infuse anti-patterns into their apps. Next slide.

So I want to emphasize the consideration, growth at what cost? Who has the power and the privilege to decide to pursue growth and who gets negatively impacted by that pursuit? And in the next slide, I want to underscore the point that, in order to survive in today's tech industry, companies must ruthlessly design for growth. The pursuit of growth is so baked into our industry that we not only regularly lobby against policies that regulate our growth, but we have also crafted processes of working that emphasize and prioritize growth.

In the next slide, you'll see a typical product cycle within teams. Note this is a very zoomed-out and generalized cycle. So typically, each company or organization might have their own variation of this cycle, likely with a lot more detail fleshed out. So this is just a very scoped-out, bird's eye view. But I think it gives a good representation of how things generally work.

So first, leadership in the company and within more specific teams figure out their quarterly objectives and key results, or OKRs. And afterwards, throughout the quarter, product teams consisting of designers and engineers, as well as researchers and product managers, will work to tweak the product accordingly so that it achieves said OKRs. And do this again and again, rinse and repeat. Next slide, please.

So to further clarify what OKRs are, they're simply measurable metrics that are used to shape what teams and companies do. And oftentimes, they are related to the bottom line of the company. So an example might be percentage increase in subscriptions to a service or a percentage increase in clicks on a particular call to action button on a landing page of a website. Next slide, please.

So not only are teams incentivized to achieve those OKRs to the best of their ability, but oftentimes, individual employees are as well. So typically, when employees go through performance reviews with their managers, they are expected to demonstrate how they've created so-called impact. Namely, they're expected to show how the work they've done has helped the larger team achieve their goals and their OKRs.

And promotions and raises are often linked directly to these conversations. So designers and others have a very strong incentive to craft experiences that fulfill those OKRs. And in the next slide, I've included some screenshots of a variety of Google searches for adjacent tactics to anti-pattern design. And these tactics, from gamification to nudging to A/B testing, have all been normalized within our industry and all involve some kind of manipulation or obfuscation for our users. And as you can see in these Google results, they're all generally positive and written in a how-to or "about such and such" format. This is from the first Google page for each search.

And in the next slide, I've overlaid the caption on top, a quick lay of the land. So as designers, we learn how to employ these techniques, as well as other anti-patterns, to accomplish our designated OKRs. And the extent to which how much all of this is normalized is pretty wild.

We learn these tactics in universities. We learn these tactics on the job. And that normalization makes speaking up about these anti-pattern designs and these tactics very, very tricky. So that's why the organizing and unionizing that's happening in tech today is so important, because there's strength in numbers.

And finally, on the last slide, I want to end my segment with this quote from one of my favorite authors, Bell Hooks. She notes that being oppressed means the absence of choices. And if you take anything away from today's workshop, I hope it is that, as we eventually create stricter regulation around manipulative design practices in tech and beyond, we have to prioritize user autonomy, especially amongst our most marginalized and vulnerable communities. And we need to move forward toward the future where deceptive and paternalistic practices are not normalized for the sake of a bottom line. Thank you.

MIN HEE KIM: Great, thank you so much, Kat. And we'll turn to Katharaina her perspectives on the underlying drivers of dark patterns.

KATHARINA KOPP: Well, thank you, and also thanks to the commission for holding this important workshop today. My name is Katharina Kopp, and my pronouns are she/her/hers. I'd like to continue on Kat's theme looking at business practices and broaden the scope of the conversation to suggest developments, which we must take into account in order to develop the right public policy solutions. I like to suggest that manipulative and deceptive patterns, or dark patterns, are part of a larger system of deceptive and manipulative practices, which are driven by the logic of the unregulated commercial digital surveillance economy.

The entire digital user interface-- that is advertising, content, or what we refer to here at this event as design-- so this entire user interface is used to optimize for business outcomes at the expense of individual and group interests. AI systems and data play a critical role in the personalization of information and experiences online and for the optimization of business outcomes. We know this from targeted advertising and social media, for example.

So it is not much of a leap to suggest that the entire online experience, including design, is part of this effort by companies to get the most out of their customers. AI-driven and automated design may not be fully implemented yet, but much evidence suggests that we are on the way to fully personalized manipulative design that exploits users' vulnerabilities at scale.

So dark patterns will increasingly look like other algorithmically-driven marketing practices, I suggest. That is, they will be hidden and yet pervasive. They will be data-driven and personalized. There's a company online called Personyze. It promises that it can build personalized learning pages and websites based on 70-plus. That's 7-0 plus user attributes. And that's probably only the tip of the iceberg. I also suggest that these patterns will use machine learning and that they are AI-driven, which allows for automated testing with no human intervention, allowing for the selection of the design that serves corporate interests the most.

I think also important for public policy is to consider-- and Kat and others have touched on that-- that dark patterns are likely to have differential impact on users, meaning that the harms will be distributed unevenly and that we are likely to see disparate impact, just as we are observing this for personalized targeted advertising. We know from a study that we will hear about later that low socioeconomic users are more harmed. And the same is true for people of color, true for children, for sure, the elderly, people with disabilities, and others. And we need to learn much more about these impacts. Next slide, please.

So dark patterns are here and are likely to become more sophisticated and pervasive because drivers of dark patterns are unrestrained. The main driver of dark patterns is the business model. I see three main objectives for which businesses design user interfaces to maximize outcomes. They design to maximize sales and revenue, resulting in financial harms. They design optimal interfaces for engagement and attention, which may lead to harms of addiction. They design for data use and collection maximization, which results in privacy harms.

Technology, of course, is also a big player in this. Technology can be deployed very differently online than offline. The ease of deployment, the speed, the scale, the precision, the control of variables and connected devices enable the optimization of this technology in the interest of companies. Of course, ubiquitous data collection also drives an increase in deployment of dark patterns and an increase in deception and manipulation, as known vulnerabilities about users can be exploited. Automated machine learning and AI, in particular, can produce optimally personalized design patterns at scale.

So lack of legislation and absence of a comprehensive FTC response to dark patterns has allowed dark patterns to flourish. I'd like to suggest then, in conclusion, that deceptive patterns are increasingly algorithmically-driven, personalized, automated, and likely to produce deception by default, as Professor Willis, who will speak at the last panel, has recently written so eloquently about in her paper. So we have to legislate and regulate manipulative patterns now. But as the Chair Commissioner Slaughter said, the FTC can enforce against deceptive practices without further delay. They have enough tools at hand to do this and address the underlying issues. So I'd like to thank you.

MIN HEE KIM: Thank you so much, Katharina. I think that was really helpful. And yes, I think panel 2 will get more into some of those line-drawing questions. But I think there will be enough here to touch on for our discussion portion.

And sort of in terms of FTC's actions, I just want to mention, we've been pursuing these types of cases, such as Commerce Planet, with its hidden fees, where they hid it behind links and fine print, kind of using visual interference and bait and switch and some other examples that will come up during the workshop. So building on the concepts that everyone's outlined, we want to dig a little bit more into the definition of dark patterns, their attributes, and prevalence in different digital spaces, and then the factors and incentives that give rise to these designs. So I'll hand it over to Miry to kick off.

MIRY KIM: Great. So let's turn to the topic of, what is a dark pattern? In getting to the definition of dark patterns, I'd like to begin the conversation with the actual term dark pattern. So does the actual term dark pattern adequately convey what these design patterns are, or should we be using another term?

So I want to turn to Harry. In your presentation, you mentioned that dark patterns are sometimes referred to anti-patterns, which you pointed out is a mistake. Do you still think that dark pattern adequately convey what these design patterns are?

HARRY BRIGNULL: Well, I would say an anti-pattern being a mistake is something that isn't done intentionally, whereas dark patterns, or whatever we end up calling them, is something that has been done intentionally in full knowledge of trying to [INAUDIBLE] or manipulate users. So I think we probably do need a term that's separate from anti-patter. But the question is, is dark patterns the right one? It serves certain functions. It's been useful, for example, on Twitter, because it's a unique series of characters, using the hashtag.

It's [INAUDIBLE] been quite good for marketing and spreading the word of it. But from a sort of legal perspective, I think it's a bit vague and sloppy. And you're probably going to be wanting to use something a great deal more specific when writing laws and what have you.

MIRY KIM: Great. And Arunesh, you did a lot of work, and you spoke about it in your presentation, about defining dark patterns. Have you thought about some other terms for these types of designs?

ARUNESH MATHUR: Thank you. No, that's a great question. I think when we started looking into many of the definitions of dark patterns, what stood out to us was that people were defining it in a variety of different ways, researchers. They were to dark patterns course, said they are some that are deceptive, malicious, misleading, obnoxious, steering, that dark patterns undermine autonomy. They subvert user intent. They sort of influence users without their consent and their knowledge.

And we saw a variety of these terms being used. They're all equally valid, I think. But to think of them in a more succinct manner, I think it was interesting for us to look at the specific attributes. Because these attributes are ultimately those that describe what the dark patterns are ultimately trying to do.

So we can call them whatever we want. But I think at some point, we need to have a discussion about what these interfaces are and what they look like. And I think the attributes really help us, in my opinion, understand and capture all of these different interfaces together.

MIRY KIM: Yeah, Arunesh, you make a good point. In fact, in the late 2000s, Greg Conti used the term evil interfaces and malicious interfaces to describe the designs that attempt to trick, coerce, or manipulate users into taking some undesired action. Kat, do you have a perspective from a product designer's perspective as to what you think-- how you would identify these design patterns today?

KAT ZHOU: Yeah, yeah. So I think that it's incredibly important that we've pinpointed that there is this malicious or manipulative design practice that is out there. I think the term dark pattern itself is a bit of a-- it's a metaphor, in some ways. And I think we realize that, in Western society, we've had this kind of dark-light dualism that we've used to frame our moral outlook, associating dark with bad and light with good.

And for me, I try to move away from that, and placing that metaphor with a more explicit phrase, perhaps hostile design or just manipulative UX. I like what Harry said about anti-patterns being maybe unintentional in that sense. But being more explicit about it, I think, is a step in the right direction.

MIRY KIM: And Khatarina, do you have a perspective regarding the term from a consumer's perspective?

KATHARINA KOPP: Yeah, I mean, I agree with Kat. I think it's a little vague, the term. So most consumers, in general, are actually not aware that these patterns are going on. So the more we can do to use terms like manipulative design, I think, the more helpful it is to start educating consumers about it, but really, much more actually for public policy makers. Even they, I think, are struggling with that term. So I would prefer manipulative designs as a way in. Because they also so hidden and so deceptive, happening behind the scenes, that that's an appropriate term I think.

MIRY KIM: Great, great. And this is a good transition into discussing the actual specific characteristics of dark patterns. So Arunesh, when you are giving your presentation, light bulbs are going off in my head. Once someone points out the tricks to you, they seem so obvious. There's almost a, how does everyone not see it reaction, even though I've never noticed it before. So the question I have for you, Arunesh, is, how challenging was it to initially identify and categorize these attributes?

ARUNESH MATHUR: Thank you. That's a great question. I think the we came about defining many of these attributes was to gather as many examples as kind of a variety of these interfaces, and then maybe try to look for commonalities across these various designs. And I think that ultimately helped us define many of these attributes together. And you're right. There are many user interfaces that you might, in retrospect, say, well, as a user, how did you not see that, how did you fall for that?

But there are many, many interfaces that can be very tricky and really hard for users and, in fact, even for researchers like us to reverse engineer and to determine what's exactly going on. And I think that complicates it. But in terms of defining the attributes, I think, once you know what a dark patter is doing and what the interface looks like, it's-- many of these attributes sort of pretty-- in a very straightforward manner flow out of the various examples.

MIRY KIM: And Johanna, if you could chime in, because you were looking not just at the desktop browsers, but at mobile browsers and mobile applications as well. Could you talk or address the process and challenges that you had in identifying dark patterns across modalities?

JOHANNA T. GUNAWAN: Yeah, no, that's an excellent question. And I think, kind of similar to what Arunesh said, it can be really hard to label these when there are some that are so subtle and some that are more obvious. But when you're looking across modalities, too, there's a lot of labor-intensive work that has to be done, particularly because apps can also be very difficult to kind of look at.

There's a previous study on dark patterns in apps, which is a fantastic paper, but they also had to manually investigate these services by spending several minutes interacting with them, just to see the interaction flow for the visual types of dark patterns. So complexities in design, interaction flow, modality differences, and other avoidances can make it really hard to automate this analysis, even within one modality, but even more so across them. And that makes it doubly difficult to try to scale up operations and get more samples.

Each modality has its own unique opportunities and challenges for automation, so building a controlled model that fairly compares these modalities would have to take all of these moving factors and subtleties into account. So when you consider this kind of extension to other modalities beyond the ones we looked, by IoT devices, the additional operating systems, and other digital platforms, you're going to continue to find that scaling up the investigation is going to continue to become difficult, but still very important work.

MIRY KIM: Great, thank you Johanna. And to Johanna's point about having found variations of our patterns on the website versus application, it begs the question, how does that happen? Wouldn't a company want to be consistent across platforms? Wouldn't a company want to maintain a certain look and feel through different mediums? Harry, as a user specialist, design specialist, I'd like to address this question to you. Do you have any theories as to why or how that could happen?

HARRY BRIGNULL: Yeah. So I mean, mobile apps often have different business objectives. So the revenue might come in via app payments and subscriptions, which work differently to the web. They often have actual gatekeepers, so there may be rules that exist in the app stores that don't exist on the web. And of course, they might be run by different teams in different countries. So achieving consistency in a big organization is actually quite difficult, even when you're trying. So there are a lot of reasons for it to be different. It's kind of easier for it to be different than it is for them to be the same, in all honesty.

MIRY KIM: And Kat, do you have a perspective from a product designer's perspective as to how this might be the case?

KAT ZHOU: Yeah, sure. [INAUDIBLE] a plus-one to what Harry just said. Oftentimes, I think they are-- the differences are unintended. In design, we try to have design systems that create a library of consistency for different modalities. And it can be tricky to adhere to them, especially when you're spread out in various teams, you're very siloed.

So more often than not, I'm willing to bet that inconsistencies are not quite intentional. But of course, as Johanna was saying, sometimes it depends on the modalities. And sometimes we've come to realize that users react and use mobile apps very differently than how they might interact with websites. So there's probably some rationale, as well, there too.

MIRY KIM: And Johanna, do you have a sense why dark patterns are more prevalent in apps versus websites?

JOHANNA T. GUNAWAN: I think the things Harry and Kat just mentioned are kind of the things that we-- that our data can sort of support. And we don't also think that it's going to always be intentional. But in terms of how they differ, there are-- if you think about what an app can do, it can be with the person 24/7. It can go anywhere. It can be always on.

You can't quite bring a desktop machine with you wherever you go and have that kind of same level exposure. So I think that's an area of future work that we're interested in looking at or seeing from other researchers as well. But as far as why that might be the case, those are some initial thoughts.

MIRY KIM: Great, thank you. And I'll turn it over to Min Hee Kim for the next topic.

MIN HEE KIM: Thanks, Miry. Yep. And yeah, and I think that was all a great overview of what these design patterns might look like, since I think the consensus we've reached here is that there isn't a one-size-fits-all term or even a definition for them. And that context really is important when we're evaluating these designs, not to mention how challenging it is, depending on, you're looking at it on desktop versus mobile, and so on and so forth.

So for now, I think we'd like to talk about the prevalence of dark patterns, so kind of carrying on that conversation of how we're seeing them in not just websites, but in these different arenas. So maybe, Katharina, if I could turn to you, you spoke about pervasive hidden practices in your presentation. Could you give us some examples of the practices you've seen and whether you think these design patterns will become even more prevalent?

KATHARINA KOPP: So I don't have much information about how, in different industries, they are distributed differently. I think what we need to look out for is how technology and AI-driven technologies are going to make this more automated and more easy to deploy. And so that's what I would look out for, that larger companies or most sophisticated companies will be able to deploy that more easily. But over time, I would expect that to become more the industry standard.

MIN HEE KIM: Gotcha. And I think you've mentioned a company called Personyze, where they were using up to, or offering up to 70 personal attributes for others to use. Who do you think is the audience or who do you think are the customers of companies like Personyze? Is it just the big techs? Do you think it's also the small companies?

KATHARINA KOPP: It's the companies who have the resources to use that and have the sophistication. At the beginning, as I said, I think as the marketplace becomes more competitive, and it becomes standard. We've seen this in other areas, in advertising. It becomes just the dominant way to conduct business online. So I think we are probably at the beginning of this. But I expect this to be, just like in other areas, just the pressures of the marketplace that everybody has to convert to that, if it's unregulated.

MIN HEE KIM: Gotcha, yeah, so an area to watch out for and make sure we're all aware moving forward. So speaking of moving forward, what's the likely forecast to be on these dark patterns? For example, Arunesh, I think you mentioned you've studied shopping websites, as well as endorsements on social media. Do you think there are certain categories of digital interfaces that tend to have more dark patterns than not?

ARUNESH MATHUR: Thank you, that's a great question. I think what we're seeing are differences across the web and mobile. As Johanna mentioned, there are differences there. But I think, oftentimes, dark patterns also show up and become very prevalent in response to regulation. So we've seen how, in response to the GDPR, there have just been a sort of a barrage of cookie walls that have been-- that have shown up and are really sort of pushing users to select the ads and not reject cookies.

So I think that there are going to be various different points at which we see many of these dark patterns appear, and maybe even go away in response to regulation. But of course, there are going to be, also, differences across web and mobile, as was mentioned here, also differences in prevalence that might vary as a result of certain kinds of technologies that get developed. Like Katharina mentioned, the user [INAUDIBLE] technologies to develop many of these interfaces might be something to watch out for.

And something our own research also shows us is how many third parties and then how the web is architected in a very specific way can lead to certain kinds of patterns becoming very, very prevalent. Because it's now-- it becomes very easy for websites and apps to include these third parties. And voila, you have an interface that lets you get what you want to get from users.

MIN HEE KIM: Thank you. Yeah, and maybe following on, if I could turn to Johanna, what do you think in terms of-- as you were studying the different platforms and the mediums, do you think dark patterns will sort of continue their trajectory on the internet in general and spread onto these other platforms that you've studied?

JOHANNA T. GUNAWAN: Yeah, I think that dark pattern prevalence will probably-- well, I suspect that it's going to follow where more of the users are, where more of the exposure is. And we're kind of seeing, in terms of usage, device usage, that it trends to mobile, it trends towards apps. And it's not different from what we're noticing as well.

So maybe if it becomes that IoT becomes much more widespread and much more common, to the extent that we have smartphone usage, then we might actually see dark patterns there, too, just because that's where more people are designing, more people are developing, and the way that these efforts are. If you think about the OKR dimension, too, and the motivations that Kat outlined, if those motivations are going in a certain direction, I suspect that dark patterns will follow.

MIN HEE KIM: Yeah, and speaking of Kat-- thank you, Johanna-- maybe I can get your designer's perspective. So kind of to Katharina's point of, this might be sort of like the beginning, or that we'll sort of see it proliferate because people are seeing it working, what do you think from a designer's perspective? Is this how generally design sort of develops and evolves?

KAT ZHOU: Yeah, I want to echo what other folks have been saying. I think they're not going anywhere anytime soon unless we regulate it and we take a more comprehensive approach towards regulating them, as well as the motivations behind them. I also think that, as new modalities pop up, whether that's beyond our mobile devices or our laptops, but also within VR and AR, I think any chance that companies get to really market to their users and to get more money from them and more data from them, they will take that, so yeah.

MIN HEE KIM: And maybe as a follow-on to that was to where companies are taking those actions. Is it possible for-- where is the role of the typical product teams? Is it possible for employees to push back against these as they start to recognize some of the different-- these elements?

KAT ZHOU: Yeah, it's very much possible. It's not easy, though. And we touched a little bit before on how teams are so siloed within companies, especially in the larger companies. So oftentimes, what happens is, you find out about something that's been shipped out to the public that was made by a team in five time zones over, that you were like, I didn't agree with this, but I guess my company has now released this feature out to the public. And there's ways to speak up. You can speak up in town halls. You can write petitions. You can organize within your companies, which is a lot easier said than done. And a lot of tech companies make it quite hard to organize. And we've seen some retaliation from companies just in the past few months. So it's definitely possible, but it takes a lot of guts.

MIN HEE KIM: Yeah, thank you. And I think, Harry, I think you mentioned in your presentation that, initially, you kind of thought designers could be the ones, sort of the gatekeepers, not participating in this. But that, given the cost, how cheap it is and how easy it is for companies to implement these, that putting the onus on the designers obviously didn't work, because here we are. Do you think there is a counterbalance that will impact companies' decisions? Is it possible that so many consumers will get annoyed that they will have a different reaction? What do you think?

HARRY BRIGNULL: I think the opportunity has been on the table for 10 years, more than that. And that hasn't been the case. That hasn't happened. Consumers haven't kicked back. Companies haven't self-regulated. They haven't made it easy enough for designers to speak up either. So another approach is needed. We're here to talk about regulation, I suppose.

MIN HEE KIM: Yeah, and I think our panel 5-- good point on the regulation and what we can do in terms of enforcement. I think panel 5 will dig a little deeper into that. Although, as we've started talking about, the lack of shared vocabulary and, I think, definitions haven't prevented the FTC from pursuing cases that have been using these types of digital designs, especially when it comes to deceptive designs, where I think Kat and Katharina both mentioned, where it gets into explicit definitions or explicit acts.

So like in Match, where they had fake love interest ads that sort of lured people in, and then it was very hard to cancel those subscriptions. So we'll look forward to hearing a little bit more from panel 5. But maybe for now, we'll touch on some of the questions that we received from the audience.

So are dark patterns being implemented intentionally? And if so, doesn't that just make them plain fraud? And I think we kind of talked about the intentionality versus results. I think Kat spoke to that. But any thoughts in terms of, are companies sitting there really thinking, let me use this dark pattern to draw consumers in or to keep them once I have them?

HARRY BRIGNULL: Sorry, is that one addressed to me?

MIN HEE KIM: It's open to anyone. Harry, would you like to start?

HARRY BRIGNULL: Yes. Are they being implemented intentionally? Yes, they are being intentional. I think that's pretty clear. They take a lot of work to build. Does that make them fraud? I guess that's a kind of legal question, really. Some dark patterns are very easy to identify. It's earnest, that you can look at them and go, that's a specific case because X. And others are quite complicated, and they take a lot of analysis to understand. It's like a complicated little machine. You have to really to take apart [INAUDIBLE].

So it's quite hard-- yeah, there will be some which will be quite easy to identify and define [INAUDIBLE] thing. I'm not quite sure the legal definition for and whether or not [INAUDIBLE]. But there will be some where you'll be able to literally point out and go, that's bad, and identify. There are others that are much more on the cusp of legality and [INAUDIBLE] so complex it'll take a while to unpick and argue about. And they're the difficult ones.

MIN HEE KIM: Thank you. And Katharina, did you have some input?

KATHARINA KOPP: Sure, yeah. I think that the challenge with intentionality is that, the more we drive towards automated decision-making, where you really don't need human intervention and where the testing can happen, as I said, automatically, to then suggest this was intentionally is, I think, much harder to prove. The intention is to optimize business outcomes. That's probably the intention. But these AI-driven design solutions will be just automatically on their own develop these dark patterns, I would suggest.

MIN HEE KIM: Yeah, thank you. And I think that probably gives a little bit more input for panel 5 to dig deeper into all of those issues. So we're getting near the end of our time. And if I could just turn back to all of our panelists. Once more, tell us which aspects of these design patterns you find to be the most concerning? And we'll start with Katharina.

KATHARINA KOPP: Yes, thank you. I think really important is to understand that these patterns are largely hidden, especially when algorithms are involved. And the FTC should use its power to investigate how companies are using this. And FTC is looking right now-- has a six-piece study looking at user algorithms in advertising. But it could also use that to look at how algorithms are used to design.

And I think we should focus on practices that are most exploitative and that use the biggest differential impact, so that we can minimize the disparate impact of these practices. That should be our goal, if we are concerned about equity, which I am, and I think many of us are. So focus on practices that affect the most vulnerable, such as children and teens and other vulnerable populations. I think that should be our priority.

MIN HEE KIM: Great, thank you. And those are upcoming in panel 3 and panel 4. Kat, for your closing thoughts?

KAT ZHOU: Yeah, I want to echo what Katharina just said. Also, dark patterns that are extractive when it comes to money and [INAUDIBLE] data and doing it in a confusing way for users, I think, those are probably one of the first things we want to regulate against. And they're pretty blatant in how they are extractive, so that would be my suggestion.

MIN HEE KIM: Great, thank you so much. And Johanna.

JOHANNA T. GUNAWAN: Yeah, something similar, dark patterns that disproportionately impact different user groups, and especially so, I think, when it comes to compliance, when you find dark patterns that circumvent terms of using privacy policies and kind of get around some of the mechanisms that we do have currently as a regulatory perspective.

MIN HEE KIM: Fantastic. Arunesh?

ARUNESH MATHUR: Thank you. Yeah, I'm going to echo a lot of thoughts here. I think context matters. Where exactly a specific dark pattern is being used, who is it trying to affect, and what is the magnitude of the harm to that user or to that community, I think is very important. So to answer the question, I think, that context is very, very important.

We can't just say that, hey, here's a deceptive pattern, and we should only care about deceptive patterns, perhaps. But what we really need is a discussion around the specific context in which that is being used. And I would echo everyone's thoughts here, especially to lower income individuals, to [INAUDIBLE] populations, looking into those kinds of dark patterns is very important, those that influence these users.

MIN HEE KIM: Great, thank you. And Harry.

HARRY BRIGNULL: I can speak to sort of what's common, to talk in a slightly different outcome. I think it's common to see privacy related dark patterns on social media. News websites, for some reason, seem to frequently make it very difficult to cancel renewed subscriptions. Domain name registrars often use tricks to get you to buy things you didn't intend to through the use of confusing terminology. And I also suspect that, as air travel starts becoming more popular now, we'll need to watch out for dark patterns in the checkout of low-cost airlines, tricking you into buying things like insurance, which used to be a popular thing and I think we'll see a resurgence of.

MIRY KIM: Great, thank you, Harry. And thank you. We've discussed various topics today, such as what are dark patterns, what makes them dark, why they seem to be everywhere, and what that might mean for consumers. So the next four panels will be delving into these concepts and much, much more. So stay tuned.

We want to thank all of our panelists for a great discussion and helping us set the stage for the rest of the day. A special thank you to those who are joining us from different time zones. Now we'll turn to Dr. Strahilevitz from the University of Chicago Law School for a presentation of findings from his paper with Jamie Luguri, "Shining a Light on Dark Patterns." Welcome, Dr. Strahilevitz.

LIOR J. STRAHILEVITZ: Hi, I'm Lior Strahilevitz, a law professor at the University of Chicago. And it's a real pleasure to be able to present my research that was co-authored with Jamie Luguri, who's a JD PhD with a PhD in experimental social psychology. And Jamie's currently clerking on the United States Court of Appeals for the Seventh Circuit right here in Chicago. So let's go to the first slide.

And one of the things I'll say at the outset is, if you're watching on a desktop, there is, as was mentioned in the first panel, a really handy slider in the upper right-hand corner. And I'll show you a number of slides. You might want to shrink me and grow the PowerPoint slides so that you can take a look at some of the data and some of the text that I'll be showing you at a reasonable size.

Well, so starting with dark pattern empirics, good scholars stand on the shoulders of giants. And we're really fortunate to have had great work done by people like Harry Brignull, who you've just heard from, Colin Gray, to build a typology of dark patterns. And then recent work by Arunesh Mathur, who you also just heard from, and his team of co-authors really demonstrated the prevalence of dark patterns in the United States. And there's similar work by [INAUDIBLE] and his co-authors demonstrating dark patterns prevalence in Europe, despite a GDPR.

And there's also really important work, including work by Alessandro Acquisti and his co-authors, that talks about uses of techniques like dark patterns to prompt consumers to share personal information about themselves that they might not otherwise be interested in sharing. But what Jamie and I were interested in doing in this project is really trying to test experimentally whether dark patterns would be effective at prompting consumers to purchase goods and services that they otherwise would not be inclined to purchase. And so that's really what the study that I'll be talking about for the next 20 minutes deals with. And we've got a lot of very interesting results that shed light on that question.

Now, I'll say at the outset, I'm going to be talking about a series of very large-scale experimental studies. One involved just under 2,000 subjects, the other just under 4,000 subjects. But what's really critical to understand is, this is not like some research done on Amazon Mechanical Turk. This is not a convenient sample of undergrads.

This is-- the studies were basically done on a population that looks like the US adult population. So we were very careful to get a census-weighted sample that reflects all the heterogeneity with respect to race, gender, region of the country, education, et cetera. And so when I show you these studies, this is really how American adult consumers generally respond to dark patterns.

So let's move to the next slide, study 1, our experimental setup. All right, so I'm going to tell you about two studies. In the first, we had nearly 2,000 respondents. And we began by asking them a series of demographic questions to collect personal information about them, and then a series of questions about how much they cared. So I want to make sure one study 1, experimental setup.

So let's get caught up in terms of that slide. So we wanted to make sure-- first, we asked consumers a whole bunch of questions about how much they care about their privacy. And then after they spent about 10 minutes doing that, we showed them a screen indicating that we were calculating their privacy propensity score. In actuality, we weren't calculating any such score. But we told everybody that, based on their responses to the earlier questions, we had identified them as someone who cared a great deal about their privacy.

And we said we had good news. We had partnered with the nation's largest provider of data protection and identity theft protection services. And based on the information that they'd already given us, we'd gone ahead and signed them up for a data protection plan. They were told that this plan would be free for the first six months, but that after that period of time, they'd be charged a monthly fee. And of course, they could cancel this subscription at any time. Let's go to the next slide.

Now, the way we ran our study, study 1, we randomly assign people to one of three experimental conditions. There was a control group that wasn't going to be exposed to dark patterns. There was a group that was going to be exposed to some mild dark patterns, just a few. And then there was a group that potentially was going to be exposed to five or six different dark patterns. And this is the aggressive dark pattern condition. So let's go to the next slide.

This is what the control group saw. We basically had a box that popped up. And we said, based on what you've told us, we know who you are, and we've gone ahead and signed you up for this identity protection program. You can choose Accept or Decline. And that's it, straightforward, not a dark pattern. We think of this as a neutral choice architecture.

And once consumers click through this screen, we sent them to a number of other questions, where-- and I'll show you those final questions in just a little bit. OK, so that was the control group. That establishes our baseline. Let's go to the next slide, and we'll see what the mild dark patterns conditions saw. First, they we're going to see a screen that didn't quite present it as cleanly, the choice, as Accept or Decline. But rather, we're already trying to bias them in the direction of accepting through some dark patterns.

Notice that Accept and Continue is in a bright red. So this should be the mild dark patterns conditions slide. I understand there's some issues keeping up with the slides. So hopefully we're on the mild to dark patterns condition slide. So you'll see that Accept and Continue is in red. And it's also labeled, the recommended option. And then you'll see that the choice is not Decline, but rather Other Options.

So let's go to the next slide. And we'll see that, if they selected Other Options, then they were going to be taken to an additional screen, one that gave them a choice between, I do not want to protect my data or credit history, and, after reviewing my options, I would like to protect my privacy and receive data protection and credit history monitoring. So the idea here is-- this is what Harry Brignull called confirm shaming. If someone wants to reject the data protection plan, they're going to have to make a statement that they don't necessarily agree with, which is, I do not want to protect my data or the history.

Well, nobody wants to say that. That sounds like a dangerous admission to make, even to an electronic interface. And then, if they selected, I do not want to protect my data or credit history, there was a final screen that gave them a range of options. Tell us why you declined. This is a common form of dark pattern. But actually, that slide-- or that box converted almost no one. So among those who accepted, in the mild dark patterns condition, 98% of them did so in response to one of these two boxes that are visually displayed on the slide.

OK, let's advance to the next slide. So this slide says, aggressive dark patterns condition, first two screens identical. All right, so what that tells you is, the people who were randomly assigned to the aggressive dark patterns condition saw the first couple of screens that looked exactly like the screens I just showed you for the mild dark patterns condition. And it was only those people who continued to say no after running through all the mild dark patterns screens that we're going to be exposed to some further dark patterns.

Not surprisingly, given that, at least at this stage in the experiment, the mild dark patterns condition and the aggressive dark patterns condition were exposed to the same stimulus, their acceptance rates across these first two screens were nearly identical. So let's go to the next slide, and I'll show you what the aggressive dark patterns condition looks like for those people who initially declined the service when exposed to the milder dark patterns.

So if people were in the aggressive dark patterns condition, then they were going to see this screen. We should be on the one labeled screens 3 to 5. And that screen provided them with additional information about the prevalence of identity theft and how inconvenient it could be if they were victimized by identity theft. So there were up to three screens that people would see in this condition. And we didn't let them advance to the next screen until a 10-second countdown timer had elapsed.

So this is a classic obstruction dark pattern. We're making it a little bit more painful. We're sucking up a little bit of their free time if they want to say no, while if they want to say yes and accept the data protection plan, we're sending them along on their way. And people who three times selected not accept data protection plan and continue, but rather I would like to read more information, finally were shown to this last dark pattern screen.

Let's advance to the next slide, which is aggressive dark pattern screen 6. And here, we employed another common dark pattern technique, which is an intentionally confusing prompt. So the question here is, are you sure you want to decline this free identity theft protection? And if people said, no, cancel, well, then they accepted the dark pattern. Even though the word cancel is in there and that might confuse a consumer, no, cancel, means cancel the earlier decision to decline. And so if they wished to reject the data protection plan, then they were to select Yes.

Now, this is, logically, an effort to extract consent. So they really should say, yes, I do want to decline. But consumers are going to see the word cancel there. They may be pressed for time, and so that might confuse some of them into accepting a program that they'd prefer to decline. Let's go to the next slide.

So final screens across all conditions-- this is what everyone saw-- control group, mild dark patterns group, aggressive dark patterns group. We asked them a series of questions, including, can you assess your mood on a scale of 1 to 7? Can you tell us how free you felt to decline the service? Can you let us know whether you'd be willing to do followup research with the same researchers-- that's us-- when we run our next experiment? And then we also had a freeform open-ended box that said, do you have any comments or questions for the researcher.

OK, so after the experimental subjects went through this material, we debriefed them fully and said, we haven't signed you up for a data protection plan, there is no data protection plan, and this is what we're studying. So let's go to the next slide. And this slide is a table that shows whether the dark patterns were effective or not. How effective were the dark patterns? They were quite effective.

So let me just explain why there's three columns here rather than two. In the obstruction dark pattern condition, where we had that 10-second countdown timer and we made people read information about identity theft, we noticed that a fair number of our subjects dropped out of the experiment at that point. And in so doing, they forfeited the compensation that they otherwise would have been entitled to.

So there's an interpretive question as to whether you want to treat those people who Xed out of their browser as having declined the service, or do you want to exclude them from the denominator in the study? And depending on what choice you want to make there, you can interpret the acceptance rate in the aggressive dark pattern condition as either 37% or 42%. I'm inclined to think that the 37% is the most meaningful number.

But let's compare that to the control group. When people were exposed to no dark patterns, only 11% of our subjects wanted to sign up for this data protection plan. When they were exposed to mild dark patterns, that figure more than doubled to 25%. And then when they were exposed potentially to multiple dark patterns in the aggressive condition, then the acceptance rate grew from 37% to 42%. So at least it doubles in the mild dark pattern condition, and at the very least, triples in the aggressive dark patterns condition. In other words, as Harry Brignull said on the first panel, we're seeing dark patterns proliferate because they're extremely effective.

Let's go to the next slide. Which dark patterns were most effective? I'll say more about that in a second in the context of experiment number 2. But at least for starters, what we can tell you about this experiment is it was those first two screens that did the bulk of the work, as far as dark patterns are concerned.

It was that choice not between Accept and Decline, but rather Accept and Continue, Recommended, or Other Options that generated upwards of 65% to 75% of the acceptance rates-- acceptances. And then it was that second screen that generated the next greatest percentage of the acceptances. So in other words, when people are successfully manipulated by dark patterns, they tended to be manipulated by one of those first two screens that I showed you.

And we'll go to the next slide, the aggressive dark patterns condition. Remember that they additionally saw several obstruction dark patterns screens and intentionally confusing No, Cancel, or Yes prompts. And on those screens, at least, they only accounted for about 13% and 11% of the acceptances in the aggressive dark patterns condition.

Now let's go to the next slide, which talks about susceptibility to dark patterns. So we analyzed a host of demographic factors because we were really interested in trying to figure out, A, whether dark patterns work-- and so far, I've shown you some data to indicate that they do-- and second, if they do work, do they work in a differential way.

And going into the study, our hypothesis was that less educated American consumers would be easier to manipulate via dark patterns. And it turns out that our study and the results validate that hypothesis. They confirm it. So it's not only the case that the results sorted by education are statistically significant, but they're also really giant in terms of our magnitude.

So if we think about what happened in the mild dark patterns condition-- remember, that's just a couple of screens-- among highly educated Americans, people who went to college, the acceptance rate was 21.2%. Among Americans who didn't go to college, the acceptance rate in the mild dark pattern condition was 34.1%, so really a very significant gap that correlates with education. Less educated Americans were significantly more vulnerable to dark patterns.

We saw something similar in the aggressive dark patterns condition, though the difference between highly and less educated people was not as pronounced there. And once we controlled for income and other demographic factors, the educational differences in the aggressive dark patterns condition were no longer significant, but they remained highly significant among the less educated versus the more educated in the mild dark patterns condition. And I'll show you more data from study 2 about educational differences in just a minute.

So let's advance to the next slide. This slide provides some data on respondents' moods, sentiments, and revealed preferences. We want to figure out whether firms that deploy dark patterns maybe gain additional acceptances, but also erode goodwill that they built up with customers. And so that's what these questions were about.

And what's really striking about our findings here is that there's no statistically significant differences between the moods of those who are exposed to mild dark patterns and the moods of people who were in the control group. They look very close to identical in terms of what their responses were. People in the mild dark patterns condition also weren't any more likely to express anger in that freeform box we had at the end of the survey. And very few of them dropped out of the survey by clicking out of their browser.

But look at what happens when we go to the aggressive dark pattern condition. This did generate a backlash. People in the aggressive dark patterns condition were substantially angrier. They were a lot more likely to express anger. They had some choice words for us, in some cases, about what they had been subjected to. More than 10% exited the survey and forfeited their compensation. And it's also the case that people exposed to dark patterns were much less likely to express willingness to do further research with us down the road.

Now, interestingly-- let's go to the next slide-- these results were concentrated among people who said no. So people who said yes to the data protection plan, even if they were exposed to our aggressive dark patterns condition, tended not to express anger or have their moods adversely affected. It was the people whose time we had wasted or people who noticed that we were trying to manipulate them that were more likely to experience a backlash effect.

OK, so I think the last thing I'll say is, we would have loved to have been able to run an experiment like this on Amazon or on eBay. We were not an entity that had a prior existing relationship with our experimental subjects. We didn't actually have their credit card information.

And so I think there are interesting questions about external validity to be asked, though our impression is that a firm that already has an existing relationship, has built up some trust with consumers through their previous interactions, might be even more successful at manipulating them via dark patterns than we were in sort of our first interaction with these thousands of American consumers.

Let's advance to the next slide, which talks about study 2 format. So we were really intrigued by our results in study 1 and decided to basically double the size of our experiment in study 2 in order to address some of the shortcomings in the first study. So you'll remember in the first study, everyone is subjected to dark patterns in the same order. And that does create some questions about which dark patterns are most effective. And study 2 really has a design that's targeted at identifying the different strengths of dark patterns compared to one another. We didn't really try out aggressive dark patterns in study 2. Everyone was exposed to either a control group or mild dark patterns, either one or two different dark patterns. And we were also able to test some dark pattern techniques that we didn't test in study 1.

Finally, we were really intrigued by some of our findings about pricing of the data protection plan. And so we did change the pricing scheme. Now we randomly varied whether people were going to be charged $8.99 a month or $38.99 a month for the data protection plan. $38.99 a month is a terrible deal. That's substantially more expensive than the most expensive identity theft protection programs aimed at consumers that we could find online.

And we also shortened to the free trial period to one month, wondering whether this dramatic price difference would affect acceptance rates for our data protection plan. Let's advance to the next slide, which shows you, basically, the experimental design in a four form, five content conditions. So to conceptualize this, consumers were randomly assigned to one of these boxes on the matrix, where they might see just control group, and with respect to the content and the form of the we were making them. Or they might see a mix of dark patterns that were focused either on the way in which information was presented or the content of that information. And to make this a little more concrete, I'll show you what the form and content dark patterns looked like.

Let's go to the next slide. So this shows you content conditions, study 2. And the different dark pattern manipulations that we tried out were as follows. There's a hidden information condition, one that makes the positive aspects of the data protection plan prominent visually, puts them in an attractive large font, and puts the less attractive aspects-- the fact that after a month you're going to be charged-- makes them less visually prominent by putting them lower down on the screen and making them-- having them appear in a gray font rather than a black font that might have been a little bit harder to see. This is your classic fine print dark pattern.

Another kind of dark pattern we employed is social proof. We create a bandwagon effect. We tell them that a certain number of people in the last three weeks have signed up for this plan. Would you like to join them? Arunesh Mathur has really interesting work talking about the prevalence of these kinds of dark patterns and also showing that sometimes firms make up these numbers in order to generate these kind of bandwagon effects.

We tried out a scarcity message. Congratulations, you're entitled to this data protection plan. But if you don't act within the next 60 seconds, you're not going to be able to enjoy this offer. And finally, we tried confirm shaming, similar to what we showed you in study 1. In order to decline the program, we've simply made people say, I don't care about protecting my data or credit history.

OK, good, let's advance to the next slide, and I'll show you what the form conditions looked like. So the form conditions, well, the control group just saw a choice between Accept and Decline. That's not a dark pattern. But then we tried out three different kinds of dark patterns. We made Accept the Program preselected, so selected by default.

In another dark pattern we made acceptance of the data protection plan the recommended option. And then in our final dark pattern that was based on content, we created some additional hurdles that they had to jump through. It was a choice between Accept and Continue or Other Options. This is a kind of obstruction dark pattern that's going to be easier to say, yes, sign me up, than it was going to be to say, no, I'd prefer to decline this.

Let's advance to the next slide, and I'll show you one final dark pattern. This is the trick question that half the sample randomly selected was exposed to. This is a confusing prompt that asked, would you prefer not to decline this free data protection and credit history monitoring. OK, prefer not to decline-- all right, so there's a double negative embedded in there. And we know that double negatives tend to confuse consumers. So if they did wish to reject the program, they had to select, no, I would not prefer not to decline free credit history and data protection monitoring.

OK, so let's advance to the next slide, and I'll start to show you some results. OK, so when results are highlighted in yellow, that means that the variation between this group of consumers and the control group were statistically significant. And so when we look at the content conditions, this is actually kind of interesting. The scarcity dark pattern, the one that said you've got to act in the next five minutes-- or sorry, you've got to act in the next 60 seconds or the deal disappears, that actually didn't move people.

It didn't significantly increase acceptance rates. But the other dark patterns-- confirm shaming, social proof, and hidden information-- very substantially boosted acceptance rates. And look at what hidden information all by itself does. It more than doubles the acceptance rate, 14.8% in the control, 30.1% in the hidden information dark pattern, one dark pattern doing that all by itself.

Let's go to the next slide. What about the form condition? Well, here again, making something the recommended option wasn't statistically significant with respect to the control group. But making acceptance of the data protection plan the default option or using obstruction techniques substantially boosted the acceptance rates, here from 16.7% up to 20.1% for the default, and 23.6% for obstruction. And what's really intriguing about our study is, if we go to the next slide, study 2, acceptance rates by condition, we can see how these different dark patterns work together.

And let's look at how potent dark patterns can be by examining the difference between the number in the upper left quadrant and in the bottom right quadrant. So when exposed to the pure control, 13.2% of our subjects in this experiment wanted to accept the data protection plan. But if we just combine an obstruction dark pattern with a hidden information dark pattern on one screen, then that's going to boost the acceptance rate all the way up to 34.5%. So it's a striking, striking increase in the acceptance rates. And you can see that other potent combinations of dark patterns, like combining obstruction with social proof or making something the recommended option with hidden information, also substantially spiked the acceptance rates.

Let's go to the next slide. Another very powerful dark pattern was the trick question. Would you prefer not to decline? Here, the acceptance rate, when people were exposed to this dark pattern, exactly doubled from 16.7% to 33.4%. And what's striking about this dark pattern is, as part of the debrief, we did ask consumers whether they had accepted or declined the data protection package.

And about half the people who accepted in response to this confusing prompt thought that they had declined. In other words, not only were consumers manipulated into selecting acceptance, but they didn't understand that they had been manipulated and signed up for something that they didn't actually want to sign up for. These results, of course, were highly significant. And the more time people spent on this confusing screen, the less likely they were to be willing to do research with us going forward.

OK, so let's go to the next slide. This is about the cost, randomly bearing the cost of the service. It turned out it didn't matter. So when we randomly assign people to be charged either $38.99 a month or $8.99 a month, that huge difference didn't produce any statistically significant increase in the acceptance rate. Well, how can that be?

Well, in some of our questions that we asked folks, they told us that they were optimistic that they would cancel the subscription plan in the first month before they were ever charged. And it may well be that consumers were highly optimistic, perhaps unrealistically optimistic about how likely they were to cancel before those credit card charges would kick in. And this idea that price doesn't matter but dark patterns do replicated a similar finding that we saw in study 1.

Let's advance to the next slide. Replicating study 1 again, it turns out that, other than the confusing prompt, these dark patterns didn't worsen the moods of consumers. There's no backlash for companies that employ these kinds of techniques if our results are externally valid.

So techniques like social proof, confirm shaming, didn't worsen the moods, hidden information actually improved the moods of our subjects, most likely because they thought they were getting something free and didn't realize that, after a month, those charges would kick in. Other words, this second study confirms the earlier results that employing dark pattern seems to be all upside for firms. As long as they don't push things too far and really annoy their consumers, they can employ just a couple of dark patterns and get away with it without alienating their consumers.

Let's go to the next slide. OK, now we're seeing education levels. And once again, in study 2, we do confirm that dark patterns are much more effective with respect to lower educated Americans than they were with respect to more highly educated Americans. So the jump among the less educated is much more pronounced in response to the dark pattern stimulus than it is for the highly educated.

Now, we did run this second experiment during the start of the lockdown in the United States. People with lower education were more likely to be out of work as a result of the lockdown. Their jobs didn't lend themselves to social distancing to the same extent that college graduates' jobs did. So we do see in the control group a lower acceptance rate. But notice that the dark pattern manipulation boosts the acceptance rates of the least educated Americans so that it more closely resembles the acceptance rates of their more highly educated and more affluent counterparts.

Let's go to the next slide. So as we think about normative takeaways, really, from our perspective, the data indicates that it's these mild dark patterns that are most insidious. They significantly increase acceptance rates for dubious service that we were offering, without substantially generating any consumer backlash. The less educated are more vulnerable to dark patterns. And when we think about sort of standard economic models of consumer behavior, consumers are supposed to be really responsive to price increases when they decide whether to buy something. And yet what we find is that they're so much more responsive to dark patterns than they are to reductions in price, even if those reductions are on the order of $30 a month.

Now, we don't claim to be the first social scientists who've ever used a randomization to run these kinds of studies to figure out the effectiveness of dark patterns. In fact, I suspect that a lot of social scientists working in-house for e-commerce companies have been running studies exactly like the ones that Jamie and I ran for years. We're just the first to publish our results and to share this data with the world. But we think it's precisely because social scientists working in-house for tech companies have done studies like ours on their own and seen how effective these dark patterns are, we think that explains the proliferation of dark patterns on various electronic platforms.

And as we go to the next slide-- and this will really be our final visual-- there's a lot more to say about dark patterns. And in our paper that just came out in the *Journal of Legal Analysis,* Jamie and I provide a lot more data, as well as a lot of legal analysis about the legality of dark patterns. So if you're interested in learning more and seeing the full results, you can simply go to bitly, and then backslash, darkpatternspaper. That's all in lowercase. Or you can just Google, Shining a Light On Dark Patterns, and that'll take you either to the SSRN website or the *Journal of Legal Analysis* website, where you can download the full paper for free. So thank you so much for listening. And if we have time, I think we might be able to take some questions.

MODERATOR: Thank you, Dr. Strahilevitz. And actually, no questions at the moment. But with that, we're on break. And we'll return at 12:30 PM Eastern to discuss how dark patterns affect consumers.

LIOR J. STRAHILEVITZ: Great, thank you.

EVAN ROSE: Hi, and welcome back from lunch-- or brunch, I suppose, for my fellow West Coast viewers-- to the FTC's Dark. patterns workshop. My name is Evan Rose, and I'm an attorney in the FTC's San Francisco office. And my co-moderator is Andi Arias, from [INAUDIBLE] the division of Privacy and Identity Protection. Before we get started, let me remind everyone of our standard disclosure. On behalf of myself, my co-moderator, and the panelists, I would like to note that the views express today are our own and do not necessarily reflect the views of the commission or any one particular organization or company.

Now, if we have time, we will try to incorporate questions we received from viewers. Please send those questions to darkpatterns@ftc.gov. And now I will turn it over to Andi to introduce our panelists.

ANDI ARIAS: Good morning or good afternoon. I'm going to briefly introduce each panelist, but you can look at their more fulsome and impressive biographies on our website. So first up is Professor Ryan Calo, a professor of Law at the University of Washington. He co-founded the university's Tech Policy Lab and Center for an Informed Public. And he is the author of a 2014 article entitled "Digital Market Manipulation and Other Work on Dark Patterns."

Next is Dr. Jennifer King, who is the Privacy and Data Policy Fellow at the Stanford University Institute for Human Centered Artificial Intelligence. She is a recognized expert and scholar in information privacy. Next is Professor Jonathan Mayer, who is an assistant professor at Princeton University, where he holds appointments in the department of Computer Science and the School of Public and International Affairs. He has previously served as technology counsel to United States Senator Kamala D. Harris, the Federal Communications Commission Enforcement Bureau, and the California Department of Justice.

Finally, but certainly not least, it's Finn Myrstad, who is the director of Digital Policy at the Norwegian Community Consumer Council, focusing on national and international issues related to privacy, cybersecurity, net neutrality, copyright, telecommunication, and more. Whew! He is also the EU co-chair of the Transatlantic Consumer Dialogue's Digital Committee. Thank you all for being here today. Now I'll turn it over to Evan, who will introduce our panel topic.

EVAN ROSE: Thanks, Andi. So we know that many of the design techniques underlying dark patterns can also be used to promote ends that most would agree are good for individuals and society. For example, research shows that making organ donation the default option when people apply for driver's licenses increases the number of people agreeing to donate their organs.

We also know that many of these design techniques are not new. Many have long existed in the brick and mortar context. For example, grocery stores deliberately place certain products, like candy bars, in the checkout line to encourage impulse purchases. This panel then will take up the question that Professor Mayer and his colleagues at Princeton pose in the title of their recent paper, "What Makes a Dark Pattern Dark?" That is, what are the properties and attributes of online user interfaces that might lead us to label dark patterns? Where do we draw those lines? I think I'll turn it back over to Andi to get us started.

ANDI ARIAS: Great, thank you. So Professor Mayer, you answer the question posed by Evan, in part, by urging dark pattern researchers to apply four normative lenses to their work. Can you give us a brief overview of those lenses and how they might bear on legal conceptions of harm?

JONATHAN MAYER: Sure. Andi and Evan, thank you for the question. And thank you to all the commission staff who have worked so hard to make today's important workshop possible. And the aim of our recent article was to provide a conceptual framework for the developing academic literature and policy discourse on dark patterns. And we sought to answer exactly the question that Evan and Andi posed. What properties make a user interface so problematic that we should describe it as a dark pattern?

And the reason we focused on that question was because we perceived an inflection point for dark patterns, where academic research could translate into and facilitate policy action. The dark patterns literature has predominantly been descriptive so far, identifying specific user interface designs that researchers find problematic. And that's invaluable work and, in no small part, why we're here today.

In reviewing the literature, though, we found that projects typically were not in dialogue about why the user interfaces were problematic. Answering that question is essential because it's the basis for possible interventions. As my colleague Arunesh Mathur explained on the first panel, papers have had varying definitions of dark patterns. They have typically relied on colloquial and interchangeable uses of terms like trick or manipulate or deceive.

And our first step in the project was to review the descriptive literature and organize specific attributes of dark patterns into a coherent taxonomy, which we grounded in choice architecture. Our next step was to explore how normative perspectives could provide a foundation for dark patterns research and policymaking. And we suggested four possible lenses for evaluating whether a user interface is a dark pattern-- an individual welfare, a collective welfare, regulatory objectives, and individual autonomy.

Since those ideas are very much abstract, I'll illustrate the four lenses with an example. So suppose that an online service offers a privacy choice to its users, and that choice is buried in a settings and requires clicking through multiple confirmations. And I appreciate that this example is entirely unrealistic and fanciful and no online service would ever do this, but please bear with me.

So one way to evaluate the user interface is whether it diminishes individual welfare. So for instance, we might think that when an online service's business practices are consistent with the consumer's privacy preferences, that's a form of individual welfare. There are, of course, other ways to think about privacy. Welfare is just one of many of those. What may make a buried privacy choice a dark pattern then is that it results in a significant deviation from user privacy preferences.

Another way to think about this user interface is whether it diminishes a form of collective welfare. We might be concerned, for example, that the user interface undermines competition by enabling an incumbent online service to extract valuable consumer data and entrench their market dominance. And so may that's why it's a dark pattern. There are, of course, other ways to think about competition. Consumer welfare-- or I'm sorry, collective welfare is just one of many ways, but just an example of how you might apply that lens.

Another way to think about the user interface is whether it frustrates regulatory objectives. We might be concerned that the user interface undermines the California Consumer Privacy Act or the European Union's General Data Privacy Regulation-- I'm sorry, General Data Protection Regulation. And that's why it's a dark pattern. And the fourth normative lens that we offer is individual autonomy. So we might find the user interface problematic because it denies consumers authorship of their own decisions.

And that would be because they're unaware of the choices they have or what those choices mean. And so maybe what makes the buried consumer privacy choice a dark pattern is that it's no choice at all. So those are the four lenses that we offer. They are by no means exhaustive and by no means exclusive. But our goal in the paper is to highlight that, by starting with these types of normative frameworks, we can conduct research that better translates to regulatory rules, standards, and eventually action.

ANDI ARIAS: Thank you. That's really helpful. And of course, as a privacy attorney, I'm not going to-- I have to ask this question. I'm going to direct it to Dr. King and Mr. Myrstad. But obviously, Professor Mayer and Professor Calo, feel free to jump in. Can you tell us a little bit more about how dark patterns are used to undermine consumers' consent to the collection and sale of their data? And I recognize Professor Mayer gave us a little bit of a flavor with his hypothetical, but I'd love to hear more from you.

JENNIFER KING: Sure, I'll go first. First, thanks, Andi, and thanks to the commission for putting up this event and having us here. So one of the primary outcomes of dark patterns in this area is to coerce or manipulate individuals into consenting to the collection or disclosure of their personal information, either at all or more than they would actually prefer to disclose. And so there are a few ways of doing this.

You can do it through defaults, I think Jonathan mentioned that as well, especially defaults that are set up to kind of maximally share your information. You can make requests for data difficult to avoid, even if they aren't necessary. I think we see this a lot with things like cell phone numbers, for example. There's often not a really valid reason why a service might need your cell phone number.

But cell phone numbers are kind of the new social security number in some ways because we don't change them. And so companies are often eager to get those because it's another way they can identify you. Having to opt out versus opt in-- this is a key tactic that we see with Apple and the new iOS update in the positive sense, in terms of it now switches the burden on the apps to have to ask you whether you want to be tracked across iOS or across mobile platforms rather than being opted into that by default. And so I think this is going to be a very interesting live experiment to see how well this works.

But one area that I'm concerned with, in general, is that of online consent. And so those existing mechanisms by which we ask for consent, I think they're pretty much a complete failure at this point. I don't think a lot of people argue with me. And what I really want to challenge is the design community to reimagine consent completely from a human-centered perspective, and not just one that translates contract law as we know it into an online interface.

So I'm of the opinion that the present mechanism of hitting I Accept with no attempt to actually inform you in a user-friendly way of what you're consenting to is potentially inherently manipulative. And I'd really like to see solutions that go further than just giving us kind of new looks on existing interfaces, such as, one of the things we look at is whether the Accept button is highlighted in advance. We can argue that's a dark pattern. I think we should just potentially consider even going further than just making small tweaks like that.

FINN LUTZOW-HOLM MYRSTAD: Yeah, and if you-- I guess the question was also centered to me. So I want to build on that, and I think Jen makes a really important point. And I can give you a few examples of how everyday consumers are encounter to this. If you take a step back, everyone has a digital device, a phone and a laptop or whatnot. And our lives are seriously very busy.

We have things to do. We have to make food. We have to work. We have to study. We have to talk to people. So our lives are super busy. And using our digital devices and using digital services is often something we do-- now, obviously, with COVID we use them more than before. But it's something we do to do something else.

We don't go online just to be using an app. We do it to talk with someone or to complete a transaction. And you just want to get to the point as soon as possible. And already there, right, we are often at the point where the asymmetry of information is huge. The power imbalances are massive because companies know so much about us. They have the ability to adapt the dark pattern to your state of mind, for example.

The timing of the dark pattern is often not by random. So it can happen when you are very busy, when you're on your way home, you're on a bus, whatnot. So my point here is that consumers, we are all vulnerable because of the pervasive nature of dark patterns. We're all vulnerable because dark patterns can be timed to our weakest point. And obviously, has already been pointed out here today, dark patterns can impact certain groups more than others. And I think the previous presentations regarding vulnerability have been highlighting that quite a lot.

So just to give an example to what Jen King said, we've all been in a situation where we've received a popup. We downloaded an app, we want to use the app, and then a popups comes. Do you agree to the terms and conditions? And already there, most consumers, we just click Yes because they want to use the service. And most companies know this. And even if you try to read the terms and conditions, it will take forever.

We actually did a stunt a few years ago where we printed the terms and conditions on an average telephone. It was more than 900 pages. And it took us 31 hours and 49 minutes and 11 seconds just to read the terms and conditions. And it goes without saying that any normal human being will fall asleep or need to go to eat, and a phone will run out of battery long before you've ever completed that. And that does not even include comprehension.

So that's like a dark pattern design or dark pattern by default that companies have introduced. And that's just the most basic one. And I think there's been a great description of patterns. And I'll give you some more examples afterwards, some types of work that me and myself and other consumer groups have done to sort of try to deal with dark patterns. Because they're really everywhere, and we need to do something about them.

EVAN ROSE: Great, well, I can pick it up from there and take this discussion in a slightly different direction. I do want to think about what we've been talking about in connection to techniques that have been used in the offline world basically forever. Professor Calo, in your writing, you've examined the differences between digital manipulation and manipulation in the brick and mortar environment. And you talk to us about what you see as the salient differences? Do digital dark patterns pose different harm to consumers, or is the difference primarily one of scale and pervasiveness?

RYAN CALO: Thank you, Evan. Thank you, Andi, and to the commission for putting this together, one of my favorite topics to talk about. So the idea that you can manipulate an environment to channel behavior has a long pedigree. We think about, for example, an earlier-- an example from the 1920s of the bridge that Robert Moses allegedly made have a low clearance so that only wealthy people could get to the beach, because public transportation was difficult under a low bridge.

This is an example that was brought up in the '70s by his biographer Robert Caro, and later by people like Langdon Winner and Larry Lessig. And so we know you can shape people's behavior by shaping the environment. Another great example that's less often talked about is the hotel keys in Europe, why they're so unbelievably bulky and heavy. Bruno Latour has a great discussion of this. It's because clerks in hotels had got so sick of people losing the keys that they made them heavy so that they had an incentive to leave them at the front desk. So just as you can get people to-- you can foreclose possibilities in service of essentially racism, you can nudge people's behavior to stop them from losing something.

Obviously, firms, companies are aware that you can use these techniques in order to profit. And so that's why you see, for example, Evan alluded to grocery store items. One of my favorite examples is the idea that a sugary cereal would be about eye level for a toddler so that they'll nag you about it, or the idea in a casino of hiding the clocks so that people who are gambling are not aware of what time it is, and they don't check themselves.

Well, when you move from a brick and mortar environment to a digital environment, there's more aspects of the environment you can manipulate. You can jettison physics at one level. Not entirely-- computer scientists, don't get mad at me. But you can manipulate more aspects of the environment. But I think, as importantly, you can also collect and leverage information about consumers.

So let me just use one concrete example of this and follow it through. So everywhere you go, things cost $9.99. Why? Because it feels further away than $10 than it really is. It's done on purpose to manipulate you into believing that you're paying less for something.

But imagine that you walk into a grocery store and you're greeted by somebody, and they tell you, we're going to do, a survey, and we'll give you something for it and whatever else. You fill out the survey, and it's all about how much you'd be willing to pay for this or that kind of product. And by the time you get to that kind of product in the aisle, all of a sudden, it's changed based on what you insert in the survey.

It's changed to your reservation price, the most you might be willing to pay. It's not $9.99. It's literally what you said you would pay. And if that sounds sort of strange and dystopian, then consider that in-- decades ago, even before I wrote my paper seven years ago, Amazon was doing this thing where it would charge you one price when you visit the website for a while, and then afterwards would figure out that you were a repeat customer and begin to inch up your price, knowing that they had you.

Or another example is that Uber, obviously the ride-sharing company-- I have a second paper a few years later with Alex Rosenblat about Uber. Uber investigated whether perhaps you might be willing to pay more for surge pricing if your battery were very, very low, because you would fear that you wouldn't be able to get home. Now, they didn't deploy this.

But what they did deploy was another interesting example, where during surge pricing, the idea being that when there's a lot of demand, the price is going to go up, initially, Uber would make things be twice as much or three times as much. But what they figured out was that consumers would sort of look at that and say, gosh, this seems really artificial to double the price. You know what I mean? And so they would introduce a false specificity. And they would say it's actually 2.2. Why 2.2? Because consumers would feel like it was more specific, and so they leveraged that effect.

And so what you really-- what the digital environment allows is, it's the mediation of the consumer that permits a lot more of this kind of a pattern that we're talking about, we now call dark patterns, and which I call digital market manipulation, and other people call other things. So it does really truly seem like it's different in kind in the end. It feels like it's a qualitative difference in addition to quantitative.

EVAN ROSE: Thank you. Mr. Myrstad, can I ask you to kind of weigh in on this topic generally?

FINN LUTZOW-HOLM MYRSTAD: Oh, I think that's a good point by Ryan here. We see it-- I mean, I think the issue of vulnerability, again, I want to go back to that because it's so key to understanding dark patterns in a digital context, that when a company knows when I'm hungry and they can target me with a content popup when they know I'm annoyed and just want to get rid of things, it is so much easier to do it in a digital world.

The experience we have from a brick and mortar stores informed how we're doing it online. But now it's taken up to another scale and can also do this in real time and really see how it works and how different consumer groups react to different types of dark patterns. So the example that Ryan had from the store, now, we all have different stores that we go into, adapted to our interest.

That data is collected from all the different apps that we're using on all the websites we're using, collected by data brokers who analyze our moods, compare us to people that look like us, so that they can much more accurately predict our future behavior. And this makes the discussion that we're having today much more serious, because you could, in an ideal world, at least train people to detect dark patterns in the physical world. If you read about it, you'd sort of see them, and they'd be kind of repetitive. But if the dark patterns always change, and the timing of them will always differ, you're not really looking out for them.

And I can give you one example from a gaming app for children, where let's say the green button is the button they click to advance from one level to the next level. And then suddenly, that button is suddenly a Buy button. Most children will have been taught, because they've been clicking, clicking, clicking, clicking, and suddenly that's a Buy button. And we have lots of complaints about that. And there are lots of cases about that. But the practice still continues.

And the scale of this is massive. And I think one of the reasons why we are not more aware of the dark patterns and it's hard to do research-- and Jonathan Mayer and Arunesh and all the people at this distinguished panel have done amazing research in this field-- but it's still so hard to detect them because they are individual and they happen in moments. And I think a lot of consumers feel shame regarding dark patterns. They feel like, oh, I should have known better. I was stupid to fall for that trap.

And therefore, it is completely underreported. And we have real problems documenting it. So this is why I think it's great that we have this workshop today. And I think it's really important at the FTC, but also enforcement agency in Europe and around the world. I'm based in Europe, even though I'm chair of the Transatlantic Consumer Dialogue. Really, we need to have enforcement agencies looking into this. I'm really happy we're doing this workshop today to, yeah, shine some light on this issue.

EVAN ROSE: Thank you very much. Could I actually go back to Professor Calo to address the sort of feedback loop back to the physical space. And also, if you don't mind, just maybe-- we're always short on time-- but talk about what degree marketers-- to which degree marketers can actually deliver on this level of customization and personalization of online spaces that we're talking about, if it's actually a realistic threat in the near future that companies will be able to make this ultra personalization exploitation of individual biases and personality traits happen.

RYAN CALO: Yeah, so just building on what Finn was saying-- and thanks for the great question-- so first of all, when I was doing this research for my paper, "Digital Market Manipulation," one of the areas of study that I came across, which precedes, of course, a lot of this work, is this idea of persuasion profiling. And the idea of persuasion profiling is a whole literature around, how do you find out what persuades a person, and then leverage that in real time?

So one of the ways you evidence that this is actually happening is that there is a robust literature on this. For example, finding out that Andi cares a lot about-- doesn't care about-- she doesn't care about what's popular, but she really cares about how much is left of something. So when Andi goes to a website, what she sees is "while supplies last," because they figured out her persuasion profile. Versus I don't care about whether supplies last. I just care about what's popular. That's all I care about.

I'm being that way. And so when I see the very same thing, I see a message that says "our most popular item." And it's tailored to Andi and I on that basis. The very people doing that research subsequently went into companies, went to work for companies whose business model is digital advertising. And so while we can't-- so we should strongly suspect that the reason that they were taken into these companies was because of this kinds of research.

Now, one of the phenomenon I found so interesting-- and I'll try to be short about this-- is, you might think of it as being, hey, what you could do in the brick and mortar world you can do much more online. And so it's the sort of idea that once you go online it's more dangerous. But one of the most interesting facets of this is, as you do-- as it becomes obvious that you can do more of these kinds of tricks online, so do you begin to imagine that you could also do them in a brick and mortar universe.

And so one example of that is-- one of my favorites-- is the idea that, in privacy, we worry that ISPs, internet service providers, will engage in packet sniffing. It's actually something that both the FCC and the FTC have looked at. Packet sniffing is the idea that the ISP, being an intermediary, would be able to look at the packets of internet information as it goes across the network and use it to advertise to you or to detect that activity.

So in around 2005, the Motion Picture Association of America just looked at that and go, wow, that's really interesting that there's this thing called packet sniffing. So they literally trained dogs to sniff actual packages in order to detect the telltale sign of a newly minted CD in an effort to detect when people were engaged in piracy. And I can give you many more examples. And as our real world, as the physical world becomes more connected, as it becomes more fused with the digital, I think we should fully expect these techniques to begin to percolate into our everyday brick and mortar lived experience, for example, through Amazon's Echo or any number of other sort of hybrids of the digital and physical space.

EVAN ROSE: Thank you, Professor Calo. Dr. King, can you weigh in here, please?

JENNIFER KING: Just very briefly, I wanted to say, for decades, information brokers have sold what are called so-called suckers lists, for example. So this is well-trod territory, in the sense that, even if the types of inferences that information brokers have been made in the past aren't precisely as good as what you might be able to do today online, we've known for a long time that we can target people based on their perceived vulnerabilities. And so sweepstakes are one of those common ways that we've seen over the last-- well, really, a very long time, decades. Thanks.

EVAN ROSE: Thank you. Dr. Mayer, can you speak to the role of experimentation as it relates to this question?

JONATHAN MAYER: Yeah, I just want to highlight a dimension of dark patterns that Ryan touched on his comments that is a bit different between the retail context and the online context. And that is the ease of experimentation. Ryan gave the example of putting serial in a certain place in the store. If you want to try that design in a store, you have to actually have people go move the cereal. If you want to try that design online, it's, change one line of code. And then you can run many more experiments. You can run them at much larger scale and identify designs that shape user behavior in a way that wasn't really possible offline.

Just to give a very clear example of that from my time at the Federal Communications Commission, there was an instance where there was a back and forth about a certain form of consumer consent. And what the company had been doing, and was publicly reported, was they'd literally been trying different shades of gray, of gray on gray text to see how that changed user behavior with the consent notice. You couldn't do that in the real world. Online you can just try different shades of gray.

The last point I want to make about experimentation is that it can result in emergent dark patterns properties. And what I mean is, you might have a system that tries to optimize around consumer engagement in some way or some business metric, that when comparing a bunch of different designs chooses a dark pattern and deploys a dark pattern, even though there was no person actually sitting there saying, let's pick this design because we think it will change user behavior in a certain way. Just by process of experimentation, you can wind up automatically picking designs that shape user behavior in ways that we might find problematic, absent intent.

EVAN ROSE: Thank you so much. Let me hand it back to Andi to move on to our next question.

ANDI ARIAS: Yeah, so I think it's great to segue hear from the experimentation to actually maybe some concrete examples and the effects of those examples. So Mr. Myrstad, your organization, the Norwegian Consumer Council, has been at the vanguard sounding alarm about the harms of dark patterns. Can you give us some examples of some of the harms produced by dark patterns? And more importantly, can you discuss how able or unable are consumers, in your view, to avoid them?

FINN LUTZOW-HOLM MYRSTAD: Yeah, thank you for the question. Yeah, I think the harms can range from anything from financial harm to privacy harm to even harms to your personal integrity, freedom of thought, freedom of choice. So we've looked into this issue for several years now. We've written several reports on documenting dark patterns, with Facebook, Google, Amazon, Microsoft.

And this has led to several legal complaints in Europe and in the US, actually. We work together with consumer groups all over Europe and all over the US. And I think this sort of cooperation is really important. But it also highlights that this is an issue that a lot of consumers are really worried about and probably feel quite helpless about.

To give you a few examples, I can start with sort of a very typical COVID example, which I think a lot of consumers have been faced with. When flights were canceled last year all across the world, airlines were faced with the demands for refunds. And the way they designed their website really did not encourage users to demand that refund. What they did instead was they designed their website in a way that-- the big blue button, the one we're trained to click on, that was, a Get a Voucher button or rebook your ticket for a later date.

And both of those options are actually giving the airline an interest free loan with no security. Basically, the airlines could go bankrupt any time and you will never get your money back. If you wanted to get your money back, a refund, a lot of the airlines would have an impossible-to-see button in very small writing, far down below in the website, and you almost have no ability to click it. So that's just one example.

This year, we launched a report on Amazon where we analyzed the click flow of Amazon in Amazon Prime. And what we discovered was they were putting so many hurdles in front of the consumer, including complicated navigation menu, skewed wording, confusing choices, and repeated nudging, confirm shaming, which we've been discussing here today, are really effective sort of measures to manipulate consumers. And this led to us filing a complaint in Europe in several countries. And also seven organizations here in the US sent letters to the FTC asking you to investigate. And I think this is really important.

One other issue we also looked at was Google a few years ago. Google has 85% of the market share when it comes to mobile operating systems. And this is an important point because that gives them immense market power. And if you buy a Sony phone or any other phone that's really not the iPhone, you have to have the Android operating system. And so what happens is, a consumer comes home, they open the phone, they open the packages, they start using the-- they want to use the phone.

What happens then is that you get a menu from Google that you have to log on with Google to get into the phone. So what we saw was the click flow and the way they portrayed the choices was in such a manner that you would turn on location tracking. And location data is extremely valuable and can reveal so many things about you, such as where you live, where you work, where you spend your nights. It can reveal your political affiliation, your religious affiliation, your sexual orientation, and so on.

So then again, we also filed a complaint in Europe against Google for breaching the GDPR. Sadly, that complaint is still pending in the Irish Data Protection Authority today and shows some of the weaknesses with the way we enforce dark patterns today, which I'm sure we'll get back to talking to towards the end of the panel a bit. So you see them everywhere. And they affect, really, all kinds of aspects of our life online.

ANDI ARIAS: That's great. So obviously, you've seen some severe effects from these dark patterns on consumers. But earlier today, we had-- there was an article published by the NAI, in which the NAI stated-- and this is the NAI, not the FTC-- that, while some users actively prefer to know more about what type of data is collected from them and control whether and how data is used, others have less interest in managing their data sharing practices. As such, they claim that regulators assessing potential policies around dark patterns should refrain from dictating specific practices or erecting limitations on the ability of businesses to effectively communicate with their users. And more importantly, the NAI states that it can be difficult to draw a clear line between practices that harm consumers and those that simply inform users that the free digital content they're consuming is supported by data-driven advertising, which is consistent, they claim, with the preferences of many people. So with that kind of devil's advocate question there, Dr. King, any thoughts on how difficult it is to identify true harm and which harms we should prioritize?

JENNIFER KING: A little-- I'll take a couple of those pieces and then see if my co-panelists have anything they want to add. So economic harms, I think, are always easier in the dark space to identify them privacy harms, for example. And really, on the privacy question, I would point everyone to Danielle Citron's and Dan Solove's recent article on privacy harms, because I think it does a great job of kind of unpacking a lot of the issues here and why proving privacy harms can be so challenging, especially in a framework where we're looking for financial compensation or financial value. And a lot of the privacy harms we experience don't necessarily have a monetary value tied to them, especially with free services.

So a couple of other things quickly regarding that NAI response, I just want to make the point that, in terms of informing consumers, from my background-- and again, I'm not a lawyer. My focus is information science, human computer interaction. When I look at notice and consent, what I see is a framework that's never, from the start, as far as I've ever seen-- and I used to work in industry before, switching to academia-- we've never seen design around notice and consent. We've never actually thought about these mechanisms from a human-centered computing standpoint. From my observation it's always been kind of an application of contract law and taking that kind of physical, literal paper contract that we had in the real world and plopping it online. And we really haven't done much better since then.

And so I actually just published an article about this this morning, coincidentally, based on work I've done in this space with the World Economic Forum. And so that's up today. And so I guess one of the challenges I see here, in terms of informing consumers, is that we've never really tried to think outside of our existing constraints. A lot of the work in this space has really just kind of taken the notice and consent framework for granted and just tried to kind of chip around the edges and make that better.

And my standpoint is, we've never actually taken even a further step back to say, what would we do if we started from scratch and really tried to design this in a way that actually spoke to how people process information and understand things? And of course, this kind of challenge with having this one-size-fits-all system, when we have people coming from so many different places when they use online services, from anywhere between older people who are more vulnerable, who may not understand how the internet works, to children. And so there's a really broad range here. And so I would really kind of challenge-- put the challenge out there that we need to rethink these processes entirely and not just focus on kind of twiddling around at the edges.

ANDI ARIAS: Great. I think-- Professor Calo, did you have some thoughts on this? And then I have a few questions from the audience and even from one of our commissioners.

RYAN CALO: I want to get to that. So I think, first of all, sometimes the privacy harm case is pretty straightforward. Sometimes it's more about how people feel like they're being taken advantage of, and they feel icky, and they subjectively feel [INAUDIBLE]. But remember my example about the grocery store where you walk in and someone gives you a survey? And purportedly, the survey is about figuring out whether they should carry this product. But then they go ahead and they use it against you later by charging you a higher price.

There are any number of examples where you believe you're giving up information, or you don't know you're giving up information, and that information is later used to exploit. And that's a pretty obvious-- in a sense, pretty obviously privacy harm. The funny thing about industry is, at one level, digital advertisers would like you, the client-- that is to say, the advertiser-- to believe that they are capable of exquisitely manipulating everybody at all moments. You see what I'm saying?

Their business model is, we can manipulate people in your service, exquisitely, exquisitely. And I think sometimes those kinds of claims are dramatically overstated. I think, since the days of hidden persuaders and Vance Packard, it's been overstated. But they want you to believe that. But at the same time, any time we're going to intervene from a regulatory perspective, all of a sudden, oh, no, there's different kinds of consumers out there, and they're pretty savvy, and the like.

I think just our baseline assumption should be that the digital ecosystem, digital advertising is not the first multibillion dollar activity in human history that has no significant externalities. And as such, there needs to be much more assertive regulation, as we are beginning to see in the United States and have seen in Europe.

ANDI ARIAS: So I'm going to combine, I think, two questions together in the interest of time. So we're getting questions about the use of dark patterns not just online, but outside of the online world, whether it's in IoT products, like smart TVs, or, as one of our commissioners asks, about the use of Wi-Fi to track consumers throughout their stores. That's a different type of IoT-- how long they linger in front of certain products, et cetera. So can you-- can one of you, whoever wants to jump in, can you provide more specifics about these tactics and how they might lend themselves to the use of dark patterns in the brick and mortar world, or even in IoT products like your smart TV?

JONATHAN MAYER: Sure. Maybe I'll just take a first stab at this. So with respect to smart TVs and Internet of Things devices, very clearly, the concerning business practices that first were on the web and then were on mobile devices are now coming to smart TVs and IoT. There's this starting to be our research in this area on privacy practices, in particular, some of which involve dark patterns, for those kinds of devices.

And the results are about what you would expect. The things that have been concerning for these other technologies are now concerning to these technologies. With respect to the shrinking difference between the retail space and what's possible online, it is certainly the case that there is more data available in retail than ever before. That, in turn, enables a form of experimentation that we haven't seen before or a form of analysis that we haven't seen before.

So it would certainly come as no surprise to start seeing more physical space designs that involve individual data, in the same way that we are already seeing in the retail space more and more concerning practices with respect to the collection and use of consumer data as a whole. So yeah, I think that difference is definitely going to continue to get smaller. And that feeds into problems in the retail space that maybe did not historically exist there and really originated in the online space, as Ryan talked about earlier.

ANDI ARIAS: Great, I'm going to turn it over to Evan. I think we have one more audience question that we want to address.

EVAN ROSE: Yes, we had a question, as soon as I find it. And this was directed at Professor Mayer. Are we missing group welfare or disparate impact in the four lenses that you spelled out? Some groups are impacted more so than others. Or is this captured in one of the other lenses?

JONATHAN MAYER: I think you absolutely could view group welfare or disparate impact as a lens on dark patterns. You could also try to fit those concepts into one of the other lenses. And just to give an example from drawing on the example from earlier, so a buried consumer privacy choice you could imagine results in different privacy outcomes for groups with different backgrounds. And there was already some evidence in support of dark patterns having disparate impact of various sorts in Lior's work.

So yes, that is absolutely a viable lens. We talk about it in the paper in the context of attributes of designs. We didn't get into it as a lens. But you absolutely could use that as an approach.

EVAN ROSE: Great. And I think we may be getting into some of that material more in the next panel. But first, we have a little bit of time left. And in the interest of getting as much as we can out of this great panel, can I go to you, Mr. Myrstad, and ask an audience question about price discrimination. Is inherently wrong or manipulative? Or is there something different here about this kind of price discrimination, versus discounts for the military, senior, students, that sort of thing?

FINN LUTZOW-HOLM MYRSTAD: Yeah, no, I think that discounts that are based on very obvious attributes such as that are much more transparent, at least. I think the problem with price discrimination that's possible now is that you can maximize how much you can squeeze each consumer based on their data on behavior. And again, we talked about vulnerabilities and how all our vulnerabilities are tracked and analyzed to a certain extent today.

It means that we all going to be a victim to this because companies will try to maximize their profit. And I think, also, it will then, of course, impact vulnerable groups even more than today. And I think that will also create some sort of-- it will make it hard to compare prices, which, of course, is a very important market sort of dynamic that we have today.

And then finally, I think it will create chilling effects, which are really, really unfortunate for the internet. People will start modifying their behavior online. They will try to game the algorithms. And it will just create a mess. So I really think we should try to avoid price discrimination or price optimization based on personal characteristics as much as possible.

EVAN ROSE: Thank you. And with the minute or two we have left, there's another topic we'd love to get to. But maybe Professor Calo can just touch on it and kind of leave us-- leave them wanting more on the topic of empirical testing and how we go about drawing these lines from an empirical point of view.

RYAN CALO: Yeah, I mean, that's actually directly what I wanted to say, which is that I think that line drawing is a problem that we're going to encounter here. But it's also endemic to regulation and law, in general, and that the FTC, in that it polices against unfair and deceptive practice, is well positioned to identify certain practices as being unfair on the basis, for example, that they exploit vulnerability and inure exclusively to the benefit of the firm and not the consumer. In fact, one of my hopes is that dark patterns is a way for our society to begin to reclaim that almost moral notion of fairness.

And rather than to do more of the cost-benefit analysis, which unfortunately is now codified as being part of the unfairness standard, because some of these things are just unfair. But you know what? That is-- 100 years ago, that is what Congress told the FTC to do. Figure out which practices that companies are doing that are just too unfair to let go. They spoke in very overt, moralistic language about it. And so my hope is that dark patterns is a place where we can begin to reclaim that role for this great commission.

ANDI ARIAS: I know that we are running out of time. But I know, based just in our discussions before this panel, that everyone has so many more things to say. So I'm going to give an opportunity, at least, for Mr. Myrstad to give us a little bit of thoughts on regulation. I know you wanted to say something about that. And then we will kick it off to panel 3. Oh, you're on mute.

FINN LUTZOW-HOLM MYRSTAD: OK, so I know that we've been-- we have rules in place in Europe and in the US. We can use them, and we should use them. And we should enforce those strongly. In addition, we do need to look for additional safeguards based on all the evidence that's coming out, making the digital manipulation more prevalent. And then we need to also add especially rules relating to data protection by design, data protection by default, and not have rules-- not to collect more data than necessary.

Data minimization it's called, not to collect data for other purposes. And also work on data portability and interoperability so that people can actually move from their services if they're unhappy with a service. That's really hard today. And obviously antitrust and competition, when big platforms are abusing their power is really key to this. So my big message is, please do enforce now, but also work on additional safeguards.

ANDI ARIAS: OK, well, I think we have run out of time. This was such a terrific panel. I hope you all enjoyed it as much as I did. I want to personally thank our panelists for a terrific, informative discussion. Let's give them a virtual round of applause, all / and next up is Rosario with our panel on how dark patterns affect communities of color. So stay tuned. Thanks, everyone.

ROSARIO MENDEZ: Thank you so much, Andi and Evan, for a fascinating panel. And thanks to everyone that is joining us online. My name is Rosario Mendez, and I'm an attorney in the FTC's division of Consumer and Business Education. And I will be the moderator for panel 3, which is going to be focusing on how dark patterns harm consumers, specifically communities of color.

So I first will start by saying that if you have any questions, you can email them at darkpatterns@ftc.gov. I will present our panelists, who are just fabulous experts on the subject. I will give a very brief description of their work. But you can see the full bio on the web page of this event.

So I'll start with Mutale Nkonde, who is the CEO for AI for the People, which is a communications firm, a nonprofit organization that focuses on educating Black communities on the impact of racial justice-- the impact of advanced technologies and racial justice. Dr. Jasmine McNealy is an associate professor of Telecommunications at the University of Florida, where she teaches courses in regulation and policy. And her research focuses on privacy, data, and governance, among other things, in media and technology.

Dr. Kelly Quinn is a clinical associate professor in the Department of Communications at the University of Illinois in Chicago. Her recent work has centered on privacy literacy and its role in how individuals envision privacy and navigate online. We'll be hearing a lot about that. And Miss Stephanie Nguyen is a research scientist and Civic Science Fellow at the Rita Allen Foundation. And her work focuses on the intersections of dark patterns and design, data privacy, and improving technology for underserved and overburdened populations.

So thanks, everybody, for being here. I will say that, on behalf of myself and the panelists, our views today are our own and don't necessarily express the views of the commission or any particular organization. So with that disclaimer, let's start talking about how dark patterns harm consumers, specifically consumers of communities of color. And with that topic, I first [INAUDIBLE] saying as an introduction that I am very happy that we are talking about this topic.

I think we-- for many of us, racial equity is top of mind, hundreds of years of discrimination, of racial injustice that continues today in many ways. And so I think that, not surprisingly with the pandemic, we have also seen of that gap that already existed affecting communities of colors has even widened and become even worse than-- when we talk about technology, for example, we keep hearing the need for more and better connectivity, for example, or have communities of color have better access to better equipment, and all of that. And so even at the FTC, we have seen sellers and unscrupulous sellers and the dishonest people specifically target communities of color, in the real world, in the brick and mortar world, but also online.

So what does that all mean when we're talking about dark patterns? Does that translate, also, to dark patterns, the targeting of or impact on communities of color? And I want to start with Mutale, if you will, if you can give me some examples. I know that in your work, in your research, you have found some specific examples. And that would give us kind of an idea to set the tone of what is it that we're talking about when we talk about harm to consumers, how dark patterns harm consumers, specifically communities of color.

MUTALE NKONDE: Yes, of course, thank you, Rosario. And hi to everybody, and it's such an honor to be with such an esteemed panel. So what I'm hoping to do over the next few minutes is give you a case use, give you an example of how dark patterns actually impact racial harm. And that really comes from the work of AI for the People, where we're always looking for ways to advance technical systems, specifically those driven by machine learning operations, impact communities of color and other marginalized communities.

And so the example I'm going to go through is-- I borrowed from ProPublica, which is an investigative journalism organization, and into an investigation they did into TurboTax in 2019. So TurboTax, for those people who don't know, is a type of software that US consumers can use to file their taxes, and as we are in this extended tax period, something that I'm sure many consumers are using. And TurboTax are operating in a political environment where the IRS has actually made it available for Americans earning under $66,000 a year to file for free, which should be great.

Now, the way that becomes racialized is when we look at the median household income of Americans by race. So white Americans, for example, in a 2018 survey, had a median household income of $65,000. Black Americans, $41,000, so they definitely would fall into being able to file for free. Latinx Americans, $51,000, so again, we're going to have a population there. Native Americans, $48,000, and then Asian-Americans as a group, $88,000, so the rest of my remarks are really going to go to those populations for whom income is now becoming a proxy for race.

So what ProPublica did is that they created a fake profile where they were a TaskRabbit cleaner and earned $29,000 a year. They put into the Google search, IRS free filing. And the first search that came up was TurboTax. Now, this is really interesting for this conversation because that was a sponsored search. TurboTax had bought the ability to be listed first under that search on Google. And we know from SEO research that 33% of all click-throughs go to that top search.

Once they went through and put in the information, they were-- the front page had the word free five times. This is great, free, lots of people who potentially qualify using it. But guess what? The consumer had to click through that website five times before they get to a page where this virtual cleaner, this dummy account, this cleaner on TaskRabbit who's making $29,000 a year, is told that they are actually an independent contractor, and it would cost $119 to file those taxes.

When they went to look at the source code of TurboTax, they found that there was no free alliance code within the source code, which meant that, at some point during that input process, that particular search was never allowed to go through to the free portal. There was this kind of a side road, if you like, that hijacked the search, that put them into a situation where they would have to pay. So they went back, a similar search, Walgreens attendant making $31,000 a year, but said that they had no health care.

The search that came back up that in that instance was that, because they had no health care, they had to then pay $59,000-- I'm sorry $59.99 to file, which is in violation of this promise to the-- the IRS had given. And kind of just in closing, that's just one example of how this, what we call toxic code, this toxic code that's embedded into consumer systems that really derails us when we're on the internet, how stuff like income becomes proxy for race, because the word Black, the word Latinx was not used in that search. And I'm really looking forward to how my colleagues can really speak more about the rights that we have for consumers to actually know how the systems that we use are making decisions on our behalf. So thank you, Rosario. And I look forward to hearing from my colleagues.

ROSARIO MENDEZ: Thank you, Mutale. That was a great example, and wow. So I want to call on Dr. McNealy to follow up. We heard some talk now from Mutale, and she mentioned consumer rights. And I also want to ask you, in addition to talking about consumer rights, what kind of-- how do you see patterns being a different practice than predatory practices that we see, traditional predatory practices that we see often, especially when we're talking about consumers that have financial distress, and how some of the predatory practices target them specifically.

JASMINE MCNEALY: Yeah, thank you, Rosario, and thank you, Mutale, for setting it up. This perfectly illustrates how dark patterns exist at what Annabelle would call the level of the chips and the wire. So it's on the layer below what we see on the surface. And we know that the surface is what everybody pays attention to, for the most part. And we can say that people on this stage are experts, and we may pay a little more attention.

But the regular people, one, don't have time to be digging into code to attempt to save themselves from institutional logics, which are all about profit-making. We just don't have time. We don't have, one, also, the know-how to make the change. So if you find out that, well, the TurbuTax code is leading me to a different place, how do I make that change? Do I need to call them? Do I need to send an email? Can I program my browser to bypass? No, for the most part, that is not a thing.

So how are dark patterns different from regular or analog traditional deceptive practices? Well, I would say that the TurboTax case study illustrates that dark patterns are way more insidious. One, for the most part, dark patterns exist for an organization to gather something of value from the user. And that thing of value is usually something they wouldn't give up had they the reasonable choice to make related to that. If they were actually the decider, so to speak, if they actually were able to assert their autonomy and make a decision on what the thing is they want to disclose or not disclose, or want to pay for or not to pay for, or want to choose or not to choose, but dark patterns circumvent this or can circumvent that autonomy.

And I think, if we look at it in this case, so say you go to a store, whether it's like a tax preparation store or a furniture store, which you usually might say, have really unconscionable terms in their furniture rental kind of contracts or those kinds of things that we're used to. And you have a pretty, let's say even a savvy user or consumer. And they look at the contract and they say, no, and then they leave the store.

Well, that's a traditional idea. So you have this really savvy consumer. But online, even if they decide not to participate, guess what? The platform, the website has already taken the thing of value. And the thing of value, in this case, is data. And that data leads to all kinds of impacts. Even if we say, oh, it's just data, but just data now means so much.

Just data means, collectively, we get to make all kinds of predictions and inferences about you and your demographics and how better to target you for advertising and marketing, how better to serve ourselves, but to serve you to help persuade you to make certain kinds of decisions. So it's truly a different animal than the regular, analog, traditional deceptive practices that we're used to.

ROSARIO MENDEZ: That's very fascinating. And so I want to open it up to Stephanie and Dr. Quinn. Is there anything that you want to add with respect to what Mutale and Dr. McNealy have said, specifically or how this is impacting underserved communities and minorities?

STEPHANIE NGUYEN: Definitely. Just bridging off of what Jasmine and Mutale just mentioned, especially with going from analog to digital, these tactics can be carried across these devices. And we've seen it from the early 1900s with radio and TV advertisements, to the 1990s with the classic popup ads saying we won a random sweepstakes to spam us. These unfair and deceptive patterns have existed for decades in different forms and formats. So we're talking about whether you call it malicious, confusing, addictive, or just plain old bad design.

This is something interesting to note, is that often, we're talking about, especially in tech-related conversations, we're talking about hidden, algorithmic, black-box ways to discriminate. In this case, we're talking about visual elements that people can technically see. But as Jasmine kind of mentioned, you may not recognize it because of the shift in tactic and norms that designers and technologists who are creating the product have created. For example, green and red have cultural connotations. Green means go. Red means stop.

Imagine if you flipped that. Imagine if you moved the placement of buttons on the screen with how humans interact with these devices that changes how someone might be able to understand what's actually happening with that device. And I think, in terms of harms, whether you're talking about, oh, now I get spammy emails because I accidentally signed up for a service, all the way to financial loss and potential discrimination or opportunity loss, we can't discount even the mundane-seeming harms here.

And I think the reason why we can't do that is exactly what Mutale and Jasmine just underlined, which is, a harm for that person who got scammy emails might be very different for a person who doesn't speak English, who may have less education, who get the spammy emails, but now have to figure out how to undo sending their sensitive information that they actually sent because someone thought they should have actually complied with this email. And so I think that's really important to underline here, is that the harms related to dark patterns are clearly documented. And I think there's a lot more work that can be done here. But I think, in addition to talking about financial harm, in addition to talking about discrimination, we have to look at the examples where shame, embarrassment, denied choice, and wasted time are actually talked about, as well, especially for these communities.

ROSARIO MENDEZ: Great, thank you, Stephanie. And so do you want to expand, Dr. Quinn, on any of this before I move on to the next question?

KELLY QUINN: Yeah. I can't underscore more than, today we treat personal information as currency. It's an asset that people have. It's a depletable asset. As Stephanie mentioned, we are now trading that information to get all kinds of goods and services online. And insidious ways to trade that information is really a form of deception. It's a form of giving away something that's a valuable resource to me to somebody I might not want to give that to.

And so I think we need to begin to think about these mini harms-- and sometimes they're not so mini-- and look at them in aggregate. Because once personal information is traded, you can't get it back. And this is really the dilemma that we find ourselves in. So if you inadvertently provide personal information, how do you retract that? I'm not sure that there's ways for us to be able to do that. And it's the cumulative effect that really can create damage to individuals.

ROSARIO MENDEZ: Thank you. I want to also put you on the spot a little bit, Dr. Quinn, with regard to your work on privacy literacy. I think I want to move to the role of privacy literacy here and how-- can you explain exactly what it is, what privacy literacy is? And what does it mean in terms of dark patterns and in terms of how it may affect-- it is impacting digital divide that exist. And also, it could be impacting the communities of colors, the communities that are underserved that we are concerned about in this panel. Can you expand on that, please?

KELLY QUINN: Yeah, I'd be happy to. Thank you. Privacy literacy is both knowledge about privacy generally-- so knowledge about privacy regulations and legislation, the practices of online companies, and also kind of technical aspects of the internet-- but it's also an indicator of privacy-protecting skills. Privacy literacy connects to the digital divide in a couple of ways. So when we think about the digital divide, the primary focus has been on access. But more recently, we've focused on skills as a second level divide, and then as a third level divide, the way that individuals navigate and use the internet in ways that are socially and economically beneficial.

What we know is that higher levels of privacy literacy indicate an ability to engage in privacy-protecting behavior online, so looking for things like checkboxes that need to be unchecked so you're not signed up for those spammy emails, or being able to minimize the amount of tracking that is done on the individual as they navigate site to site. And our research has found that sociodemographic factors, the traditional markers of marginalization, things like being female, being less educated, being of color, are also indicators of lower levels of privacy literacy. So this means that having lower levels of privacy literacy also means that individuals are less equipped to protect their privacy online. So kind of the end result of this is, while the technology access gap is narrowing, disparity to navigate online and use technologies in capital-enhancing ways may be continuing.

ROSARIO MENDEZ: Yeah, I wonder also to what extent culture is a factor, in terms of like what is privacy, what is the meaning of privacy for different cultures, and how is that connected to dark patterns? Does anyone want to jump in on that?

KELLY QUINN: You hit on a really important notion here, that privacy is a cultural construct. In other words, we think about privacy in terms of the way that we navigate online-- we navigate our everyday relationships. And in a digital realm, the conception of privacy is very individualized. There is a lot of responsibility that's placed on the individual to navigate their own privacy. And that's not necessarily how privacy plays out in different groups.

For example, when we look at privacy controls, for example, on a cell phone device, they're very individually oriented. All your accounts are connected to that device. And it presumes that there's a single user. But in many communities, cell phones are a luxury commodity. They're shared among individuals. There's no way to protect individual users on a cell phone. So we have this bias in the way that we build these technologies, going back to what Stephanie is suggesting, that there are things that we should be doing to encompass or embrace multiple ideas of what privacy is.

STEPHANIE NGUYEN: And sort of to build on that point, I'm a designer who has researched and created dozens and dozens of user interfaces and products. I have also studied data privacy and how people understand it and how people interact with consent forms. My parents still have no idea what that is or what I do for a living. And I think when balancing business realities with empirical realities of just creating a product with screens and buttons, there's a lot of layers and communication and culture embedded in that process.

So I really want to underline like, we cannot point to users to learn and know better. I support users being informed. I support people understanding and weighing pros and cons. But we've seen this time and time again with research. Consent popup boxes with labor-intensive reading and the assumption that they can balance all of these needs is just not-- it's not feasible. Even a Consumer Report study done by Katie McInnis, who's a policy expert, she did a lot of work on the evolution of consumer attitudes toward data collection over the past 25 years.

And even though consumer awareness has actually increased, this doesn't actually translate to increased control. And so I think, on this point, we can't rely on self-accountability, and we can't rely on users educating themselves, to the extent that they'd be actually able to navigate these spaces. And so we need more regulation. We need people to step up and use the tools that exist in order to actually combat the industries and the business models that are created here that will eventually trump how designs and user interfaces are manifested in practice.

MUTALE NKONDE: And kind of just further to that lack of individualization, I know in our work with AI for the People, we're also really interested in social media as a site for online privacy, which is not completely tied to this. But certainly with my work with TikTok, one of the things that I'm always seeking to do, and my fellow members on that advisory board are seeking to do, is try and create a business culture in which the companies themselves that are enacting these harmful behaviors have a responsibility to protect their users as consumers. And I think this idea of whether we're a user or we're a consumer has an impact on the regulatory conversation.

I know when I did work around social media labeling in the Deep Fakes Accountability Act, we actually used consumer law in that particular statute to argue that American people going online should know that the information they're seeing is truthful. And we advocated within that bill for labeling, social media labeling. What did we see in the fall of last year? We saw Twitter labeling. We saw Facebook labeling.

In the example that I gave for TurboTax, the issue that I have, and where I think could be a key point of intervention, and really the topic of this workshop is bringing things to light. If people only knew that there are ways of-- there are technical ways of increasing the profit margin, then, yes, individually they may behave differently. But what if you're making $29,000 a year like the TaskRabbit person, who is a domestic worker in that particular scenario, and potentially does not have the time?

They have other limits on their time that aren't enabling them to do that type of work. And the statistics that I showed, they're more likely to be a person of color and more likely to be feminized. And so there is a role that government plays that goes beyond literacy. It really is our right as consumers to have these overall protections. Because as somebody who works with industry, just saying that they should be good people is not very compelling. We need stronger arms.

JASMINE MCNEALY: Just to chime in-- and I agree with everything that's been said so far. But to me, it seems to be pointing to the two things-- one, collective harms that arise from the collection, the aggregation, the use of data. Because we know that the use of one individual's data is not one individual's, but data is a networked thing. So it always leads to others in that person's family, their community.

And then, so predictions and inferences can be made about a zip code, about an area code, all of these things. And this is not just Jasmine making stuff up, but we have actual stories and reporting on this happening. But also, it points to the disparate power dynamics we're talking about between individuals and organizations. And that's where the power of governments like the FTC comes into play. Because the FTC is the premier, if not the only, consumer protection agency that we have to get involved in dark patterns and privacy and data governance, schemes like this, to ensure that the consumer is protected, and not just consumers or people using, but all people who are implicated by the trade practices of organizations.

And I should say it's not just corporations. But civil society organizations use dark patterns and these kind of design patterns, as well, that can collect data or get people to or persuade people to do things that they usually would not do. And so having that deceptiveness as an unfairness, quite frankly, as a point of regulation is really important to use to make sure, especially for already marginalized communities, that the consumer protection agencies are there and making sure that power dynamic is changed.

ROSARIO MENDEZ: Thank you.

KELLY QUINN: I would just [INAUDIBLE] one point on that to tie back what Jasmine is saying back to something Mutale mentioned earlier. Data inference is an issue. And this is where much of the collective harm ends up being. We have lots of things that serve as proxy for the kinds of things that we regulated in the past.

We have lots of characteristics from people's online activity that served as a proxy for race, ethnicity, gender, income-- well, not income, but income is one of those indicators. But the data inference from the online activity that people have created new categories that are not regulated in the same way that other categories have. And this is where, I think, we need to recognize that this inference has created these proxies, and this is where the harm really begins to exist.

ROSARIO MENDEZ: Thank you. And with that, I would like to take the, I guess, next 10 minutes that we have left in the panel to answer a few questions, and then end with some suggestions from you guys. So one of the questions that we have from the audience is related to if any one-- if you guys have any thoughts on California, the new California privacy law, the California Privacy Rights Act, that gives consumers the right to limit the use of sensitive personal information, which includes race, sexual orientation, gender, et cetera. Does anybody want to comment about that?

MUTALE NKONDE: I certainly welcomed it. And that was just because, I think on the federal level, what happens is that we look to the states right. And we look to the states for two reasons. First of all, we want to make sure that there is a community need and a community demand for that. So if groups that are advocates can get together in a state and advocate for this particular issue, it then puts pressure on their congressional delegation to go to the work of the Capitol and do the people's bidding.

And the second reason that I really appreciated it is that it happened in the United States. Now, we're in a country where some would argue that California is not representative of the whole country. But until it secedes, it is a part of the United States. Therefore, it becomes much more difficult, and it puts pressure on the other states to follow suit.

And I think that this centering of privacy is important. But I would also argue that there needs to be a centering of justice. We need to figure out who are these dark patterns, in this instance, harming the most, and how do we alleviate that harm?

I think Stephanie makes some really, really good points that I would like to underscore around people for whom English is not their native language. Imagine if, as native English speakers, we do not see this going on. Imagine what it is to then translate that into a cultural context, where I'm a new American. I have not been an American for a year. I naturalized last year. But when I am trying to be, quote, unquote, "American," I am going to follow rules because I want to fit into this idea of the United States that I have.

And so if I happen to be a new American in a state like California, where they have made a commitment to my privacy, I have additional layers of protection. And it really opens up what I call the policymaking imaginary, that it can and it should happen. It's not just something that is going to happen in the EU and the GDPR because they're European. And I love how strong it is. Illinois has really, really strong protections as well. And they've really helped in the biometric realm, which I know isn't this panel, but it proves that we can use the force of government to protect ourselves where there is the will.

ROSARIO MENDEZ: Thank you. Another question that we have relates to how dark patterns may impact older Americans. I know that this panel is focusing on committees of color. But that is another group that could be specifically impacted. And if anyone wants to share any thoughts on that, it would be great.

KELLY QUINN: I'd be happy to jump in on that one. I do work a lot with older adults and how they use technologies. And there are-- what happens as we age is we-- our cognitive ability peaks at around the age of 30. Sorry, everybody, but that's how it works. And we also lose some physical abilities too. We're not as quick to do things.

So dark patterns really can impact older adults, as well, because they, first of all, have not always developed the abilities to understand how technologies work. In my experience running workshops with older adults, I found that they have a really hard time understanding interoperability and how different things can be embedded into a website. And this limits the way that they can interact with platforms.

Things like low contrast in click buttons or icons, they miss things. And it really-- they are vulnerable in ways because they have not developed the same digital literacy and same digital skills as younger adults have. So I think dark patterns definitely affect older adults in ways, much the same as they do for things like being educated and being female.

ROSARIO MENDEZ: Thank you. Yeah, it sounds like there's a lot more that we need to do and a lot more that we need to find out about, the true impact of dark patterns and what we can do to combat the negative effects of it, especially in communities of colors, older Americans, in groups that are underserved and overwhelmed, if you will. So the last minutes that we have, I want to get your thoughts on, what do we need to do, what are your suggestions on what we need to do to move forward, specifically if you can tell us what kind of research do you think is needed and what other solutions, if you will, you kind of suggest. So I'd like to start with Stephanie, if you have any thoughts on that.

STEPHANIE NGUYEN: On this front, we have the tools available to combat dark patterns. There are endless taxonomies, definitions, and research showing that dark patterns are creating disparate harm. We have rules and emerging bills defining dark patterns, technologists, practitioners, designers, and engineers who are organizing researching and trying to outline interventions.

On the state level and the city level, there are different business bureaus and organizations who are already trying to ban hidden fees for things like delivery apps and things like that. There's been a lot of evidence. But I think we're at the point where we're missing opportunities for enforcement actions. We have the FTC. We have a lot of these other regulatory bodies in governments who really need to step up. And I think that between people and hiring and making sure you have the right leadership to be able to push for these interventions and to the actual process of being able to launch investigations and initiate legal action, I think that's incredibly important.

And on third point, just the outcomes, are we seeing the outcomes that we want to see with these types of interventions? I think having panels like this is incredibly helpful to be able to rally people together and be able to have these conversations. But talking about it and doing something is very different. And I think now is the time that we actually need to step up and start working on these issues.

ROSARIO MENDEZ: Thank you. Dr. McNealy, do you have any thoughts?

JASMINE MCNEALY: Yeah, so I agree with Stephanie completely. I think we have the law already there. It's part of the FTC Act. And that's anti-deception, anti-unfairness, as it relates to consumers or anyone who could possibly be deceived by these practices. So the law is already there. It just has to be enforced in full force. We know the harms are there, so we know enforcement has to be there, particularly as it affects already marginalized, already vulnerable communities, however we're going to define that, but especially these communities, and especially as it deals with children as well.

ROSARIO MENDEZ: Thank you. That will be the topic of the next panel. Dr. Quinn, do you have any concluding words, specifically if you have thoughts on research that may be needed?

KELLY QUINN: Well, I think it's important for us to begin to identify ways in which dark patterns are being leveraged against particular communities. And this is an underserved area in terms of the research community. And it's the kind of thing that I think would allow us to have that evidence that we need to really protect groups that are more vulnerable.

I'd also advocate for some literacy efforts to enhance consumer privacy. Although what I want to recognize is that, even with literacy efforts, the technologies are built in fairly opaque ways in this age where data collection is ubiquitous. Everything we do is being collected. And also, the ways in which it's being collected are really not visible to many of us. The deck is stacked against the individual.

So literacy efforts, while they're important, also place an increased burden on the individual to be able to navigate in this environment. And that's why it's so important to focus on the practices of those agents who benefit from dark patterns, benefit from reduced levels of literacy. We need to be able to recognize this and be able to minimize the harm that they cause.

ROSARIO MENDEZ: Thank you. Mutale, any final words?

MUTALE NKONDE: Yes, very quickly, on the research side, I would love to see more research which uncovers the different-- kind of creates the taxonomy for proxies for race. Because certainly, in my work, much of the pushback that I get from industry or government or whomever we're speaking to as an organization is, but they did not say Black, they did not say Latinx, they did not say Asian. And so going and doing some of the work of, zip code can become a proxy because what we're talking about is inference.

And it's not this idea of hard data collection that sets us onto the road. It's really an amalgamation of online behavior. And I think once we become accustomized to this idea that proxies can lead to disparate racial impact, it furthers the conversation towards justice, which is what we're definitely focused on. And then the second thing, in terms of intervention, is I feel that, as a legal practitioner, all the laws are there-- my colleagues have said that-- but the case law isn't.

So certainly in New York, we start to work with organizations, like Lawyers for Black Lives, where we can actually think about bringing cases to fore around racialized dark patterns, around other forms of technological discrimination, so that we can use the courts to try and then force the legislative conversation in the same way that *Brown v. Brown,* that court case really did pave the way for equality in education that we're both looking at. Why can't we use that same model? I think that we have the talent, and I think that we have the will.

ROSARIO MENDEZ: Thank you so much, excellent ideas. Unfortunately, I think we can keep talking about this for a long time. But unfortunately, we got to start wrapping up. So with that, I want to thank all the panelists. Your expertise has been great. And I hope that we can continue this conversation, or we can use you as a resource later on.

I also want to thank the organizers, my colleagues at the FCC, for bringing this topic to this workshop. I think it's really important, and I hope that-- I know that at the FCC we're going to continue talking about it and figuring out some ways in which we can continue to protect all consumers in all-- in every community. So next, I think we're going to be turning to Sam to lead a discussion about how dark patterns target kids and teens, which was brought up in this panel too. So stay tuned. Thank you.

SAM JACOBSON: Thank you, Rosario, for another great panel. And thank you to everyone joining us this afternoon. My name is Sam Jacobson. I'm an attorney in the FTC's division of Financial Practices. And I'll be moderating our fourth panel of the day, "How do Dark Patterns Target Kids And Teens?"

On this panel, we'll be discussing how dark patterns are targeted at children and teens, why kids are especially susceptible to these tactics, and the effects of this targeting on kids, teens, and their families. To discuss these issues, we're fortunate to have with us a distinguished panel of experts on child development, children's digital media usage, and children's advertising. First, Criscillia Benford is an educator, researcher, and media theorist who serves on the boards of the Campaign for Commercial Free Childhood and the Children's Screen Time Action Network. As a researcher, she investigates how different media, media environments, and interfaces impact human perception, action, and well-being.

Next, Dona Fraser is the senior vice president for Privacy Initiatives at BBB National Programs, where she oversees the strategic development, implementation, and ongoing efforts of the organization's privacy programs. Miss Fraser previously served as the Director of BBB National Programs Children Advertising Review Unit, CARU, the first FTC-approved COPPA safe harbor program.

Next, Josh Nelson is a proven social change leader with 15 years of experience using digital advocacy, online to offline campaigning, communications, and data to win campaigns for progressive change. He is currently directing a campaign to protect kids from big tech for ParentsTogether, a parents-led organization representing 2 and 1/2 million socioeconomically and racially diverse community members in the United States.

Last but not least, Jenny Radesky is a developmental behavioral pediatrician and assistant professor of Pediatrics at the University of Michigan Medical School. Her NIH-funded research focuses on the use of mobile and interactive technology by parents and young children and how this relates to child self-regulation and parent-child interaction.

Before we get started with our discussion, just a reminder that the views expressed today are our own and do not necessarily reflect the views of the commission or any one particular organization or company. If we have time, we will try to incorporate questions we receive from viewers. Please submit those questions to darkpatterns@ftc.gov.

OK, before we jump into our discussion of how dark patterns are targeting and impacting kids and teens, I think it'd be helpful to start by getting a better sense of how often kids are using digital technologies and what apps or services they are using. Jenny, can you help us out with this?

JENNY RADESKY: Yep, sure. So thanks, everyone, for having me here. This has been a great workshop so far. I'm learning a lot. So I will start with some pre-pandemic data. So we knew that the mobile apps and internet platforms that can contain dark patterns are widely popular among children.

So for example, surveys from the Pew Internet Center show that over 80% of parents report their children under the age of 11 use YouTube, many on a daily basis. My lab's NIH-funded study tracking the mobile device usage of over almost 350 preschool-aged children has shown that young children love apps and platforms that have been designed for adult users. So these include YouTube, which was over 1/3 of our sample of three and four-year-olds who use it, but other apps, like Facebook Messenger and TikTok and even gambling apps or [INAUDIBLE] apps like Grammy.

And then children who had their own mobile devices use their mobile tech on an average of two hours per day. But this was longer in children growing up in families where the parents have lower educational attainment. And data sharing from those apps sending private identifiers to third-party companies is also more common in lower education families.

So other research my lab has done has looked at how educational these apps truly are, finding that loss of them actually have distracting enhancements or ads that get in the way of the educational experience. We've also looked at YouTube viewing histories of kids 8 and under with Common Sense Media, finding that the early childhood content has the highest ad load, some of it camouflaged to look like recommended videos, but actually a sponsored ad.

So this was all pre-pandemic data. To give a sense of what's happening during COVID-19, we just recently surveyed over 300 parents of elementary school aged children here in Michigan. And we're still in the process of analyzing data. But so far, a few data points that are interesting is, 31% of parents report that their child who's 5 to 10 years has started using social media at a younger age than they had planned.

And then compared to in-person school, kids who are attending school in person, kids who are on remote learning have been doing more multitasking, according to their parents. They're spending more time every day on mobile apps and games, social media sites, like YouTube and TikTok. And they're playing more console-based computer games as well. So that's some data that we have about what's going on during COVID-19. But I know Josh, also, has some other interesting data.

SAM JACOBSON: Great. Thanks, Jenny. And Josh, your organization ParentsTogether also has surveyed parents, both before and during the pandemic, about their kids' technology usage. Can you briefly share some of your findings with us?

JOSH NELSON: Absolutely. ParentsTogether conducted a survey of thousands of parents in our membership last year and found a huge spike in the amount of time kids are spending online during the pandemic. Specifically, we found that nearly half of parents reported their kids spending six or more hours online each day. The average amount of time kids are spending online doubled during the pandemic from roughly three hours per day to roughly six hours per day. 26% of parents actually said their kids are spending eight or more hours online each day. And notably, despite Google's claim that YouTube is not for kids, as Jenny noted, YouTube was by far the most common platform used by kids in our survey. No other platform even came close.

SAM JACOBSON: Thanks, Josh. Dona, anything you want to add based on your work at the BBB and [INAUDIBLE] seeing? Dona, I think you're on mute.

DONA FRASER: My apologies. So I think that the statistics that everyone's talking about, with regards to increased usage, is consistent with what we've been observing as well. Personally, I do find myself questioning data regarding device usage of those under 13. And I'm not questioning the veracity of the data, the studies that Jenny and Josh have done, because I know that they're talking to parents.

But what we oftentimes find is that, if stats tell us at an average age of a child is getting their first device at 10 years old, then sometimes I think it's safe to presume that parents are sharing devices with their young children. And in our experience, we find that oftentimes parents are not changing settings regarding privacy data collection before handing their devices to the young child. So we just need to be, I think, careful about the veracity of some of the data that we're seeing. And again, I'm not talking about the data that's been presented here.

One of the other concerns that we've had, especially during this pandemic, and it seems to be less discussed, is the consequence of kids having become more comfortable being in front of a camera, whether it's out of necessity for schools or non-essential activity. Either way, the increased engagement undoubtedly raises concerns about their level of engagement and any possible privacy implications.

SAM JACOBSON: Thanks, Dona. Those are really helpful point, I think, for us all to bear in mind. So with that helpful context, I'd like to turn the panel's focus to the role of user interface dark patterns in children's media, apps, and games. And I'd like to begin by asking our panelists, why are kids potentially more susceptible to dark patterns? And what are some of the techniques that app and game designers use to target these vulnerabilities? Dr. Radesky, these are questions that you're focused on quite a bit in your research, and so let's start with you.

JENNY RADESKY: Yeah, thanks. So I have been trying to really articulate the mismatches between persuasive technologies and user experience designs and developing minds, kids that have not developed all of the cognitive or emotional skills that we hope adults have. So I'm not only trying to identify areas where design patterns that are considered manipulatives for adults are more so for children, but also areas where design patterns that might be pretty benign with grownups cross a line for children, based on their unique vulnerabilities, which, developmentally, are actually strengths. If you think of it from a child's perspective, they have to have a lot of these differences in the way they process the world just to learn and to grow.

So let's go through. I've outlined these five different-- five differences between the way children see things in the way they might be exploited through technology design. So first, kids have immature executive function. Executive function is the air traffic controller of your brain. They develop rapidly over the first five to seven years of life. It's where-- it's all in your prefrontal cortex, your frontal lobe, that through play and parent-child interactions, physical activity, sleep, and other sort of positive learning activities, kids develop skills like emotional control, impulse inhibition, mental flexibility, and attentional control.

In the first years of life, children's visual attention is strongly attracted to novelty. This is adaptive, actually, because they want to see the new things around them and categorize them, and build meaning systems and conceptions about them. But it can also be manipulated by lots of interactive visual and sound effects in media. So although an adult might be able to resist clicking on something that has sparkles or a countdown clock or candy or piles of gold, a child may not be able to possess that. So they can't really critically evaluate how their behavior is being persuaded by a designer.

OK, so number 2, after executive function, is children really form imaginative relationships with characters. They're really susceptible to pressure or role modeling from the parasocial relationships that they develop with their favorite characters. And so there's been studies showing that kids learn better from a character like Elmo, who they know and trust, and they're going to follow those instructions more than a character they've never met before. So we showed in our study of in-app advertising published in 2018 in *Journal of Developmental Behavioral Pediatrics,* that characters like Strawberry Shortcake are taking advantage of this by encouraging in-app purchases.

So the third developmental difference is reward susceptibility. So kids' behavior is really shapeable by positive reinforcement and rewards, which is why we use sticker charts for potty training. So this is an adapted part of social learning. For young kids, you learn to repeat behaviors that are pleasurable or that are reinforced. But it also means they will follow lures, they will follow rewards down [INAUDIBLE] or to next levels. And repeating this over and over will lay down some habits, or, we hypothesize, maybe make them less interested in everyday humdrum activities that aren't packed with coins or presents or other sort of gimmicks.

OK, so fourth, kids are concrete thinkers. They're not yet abstract thinkers. So they think of digital concepts differently than us. We just did a study interviewing 5 to 10-year-olds to see what they thought about digital privacy. And they understood things like, OK, the app saves by progress. It remembers what videos I like so I can watch more of them.

But none of them understood the scale of digital processes through which companies make inferences about them. And therefore, they weren't doing much protective behaviors when it came to online privacy. Even the older kids who did said they were mostly protecting themselves from bad actors, like hackers, not from commercial surveillance.

So OK, fifth difference is that children don't understand virtual currencies. They start to understand money concepts, like coins and counting, between ages three and five. But it takes years to develop the idea of value and how you exchange money for different goods. So you can't expect that kids are going to understand the value of different virtual currencies, especially when, game to game, all those different currencies are-- they cost different scales of money and different ratios. So those are the kind of five differences that I would want designers to be thinking about if they're going to create interfaces that kids are using.

SAM JACOBSON: Thanks, Jenny. So speaking of designers, Criscillia, you have some background in design and media theory and psychology. Can you talk a little bit more about how designers target some of these vulnerabilities or susceptibilities that Jenny was just mentioning?

CRISCILLIA BENFORD: Sure, and thanks. I'd like to thank the FTC for convening this panel and thank the workshop and thank you, everyone, for participating. This is such important work. Earlier in the panel, one of the first workshop panelists talked about how the designers are motivated by meeting different business objectives. And that's tied to their employee performance.

And so this is one of the reasons that you have some of the tactics that we're going to be talking about later, why they're so widespread. And I'd really like to underscore a particular vulnerability that Jenny talked about, and that is kids susceptibility to parasocial relationships, that kids-- we've probably all seen this, a kid forming an intense relationship with a dull, an emotionally-tinged, one-way relationship. So there's a lot of research that designers use that show that orchestrating social cues, creating a user interface that conveys social presence, opens a user, even an adult user, to-- it makes them more open to persuasion.

So a very popular technique for children, as Jenny was talking about, is to have a character, familiar is best, but even a character with a very strong personality that the child can connect with, that will open the child to persuasion from that character. Now, ideally, these techniques are supposed to be used to increase usability. But what we're going to see is that they are often used to increase time on device or to persuade the child to buy, make an in-app purchase.

And so when you have a design tactic like a character, that is trying to help the kid relate to it, you are opening the door to persuasion dynamics that are called social influence. And when those dynamics are combined with other dark patterns, or what I like to call adversarial design, like nagging, obstruction, aesthetic manipulation, and navigation interference, they become supercharged when you have a character involved with those sorts of visual design strategies.

So the other thing that I want to point out about using social influence on children is those tactics are even more powerful for children who may be lonely or alienated. And they will relate very strongly to that character. And so this is very much, I think, an extremely adversarial design practice.

SAM JACOBSON: Thanks, Criscillia. So I guess now that we've talked some about some of the reasons why kids and teens may be particularly susceptible to persuasive difficult technologies and how sort of, in general, designers exploit or target those vulnerabilities, I think it would be helpful to look at some specific examples of dark patterns in kids apps and games. And our panelists have prepared some examples and slides to it to show us here. So I'm going to turn things over to Josh to get us started here. And Josh, can you talk a little bit about how companies are employing dark patterns to keep kids using or returning to their services, and provide us with some examples?

JOSH NELSON: Certainly. So two examples I'd like to focus on today are YouTube autoplay and Snapchat snap streaks. With YouTube autoplay, after a child selects one video to watch, they are subsequently shown additional videos indefinitely. This design feature proactively promotes YouTube addiction and excessive screen time, particularly among children who may not realize they're being shown a continuous stream of videos.

As you can see on the slide, when a video finishes, there's a rapid countdown from 10 before the next video starts playing automatically. That countdown moves so quickly that many kids simply do not have time to read the title of the next video before it starts playing. Also note how the Play Now button is emphasized with a different background color, while the Cancel button is essentially hidden with a dark background color that nearly matches what's behind it. Next slide, please.

With Snapchat snap streaks, users are given a flame or fire emoji to indicate that they have snapped each other within 24 hours for more than three consecutive days. This design feature explicitly encourages children to use Snapchat every single day. And by removing the emoji for a broken streak, it creates a consequence for failing to do so. As you can see here, some Snapchat users, including some kids and teens, have streaks that last hundreds of days. Once you've invested that much time in a streak, there's a very strong psychological incentive to continue using the app every single day to keep the streak alive.

SAM JACOBSON: Thanks, Josh. Dona, once a company has a kid's attention and they're returning to the service and spending-- logging hours on a particular app or a particular game, it makes sense that they would want to find a way to monetize that engagement, either by collecting behavioral data about the child's usage habits or by selling them goods or services. Could you tell us about some of the design tricks that companies use to get kids and teens to spend money in their apps and games?

DONA FRASER: Sure. I do want to first note that the examples that I'm providing here are pulled from CARU investigations, all of which have been publicly announced. And 99% of the time, the companies we work with on these cases are unaware that they have violated CARU's guidelines, COPPA, or other laws. However, when they do find out they rarely work with us to come into compliance. And we do have the ability to refer noncompliant actors to the FTC or appropriate state AGs. And it's rare that we have to do so, which is, I think, proof positive that self-regulation does work.

So on this first screen, what you're seeing is a case we call My Talking Tom, which, in this case, we had to determine whether children would understand that the in-app advertisements were actually ads and not game content. So CARU's guidelines provide that if an advertiser integrates an advertisement into the content of a game or activity, then the advertiser should make clear in a manner that will be easily understood by the intended audience that it is an advertisement. In this particular instance, CARU determined that a child may not understand that the ads that pop up on the top of the screen are advertisements.

So for example, the ads appear on a block on top of the screen. A child can tap on them, and install them through the app store. It may be difficult to tell the difference between the content on the app and the advertisement because there are so many items on the screen. You can go to the next slide, please.

So there are also instances where popular and familiar characters in the app pressure children to make purchases in order to unlock certain features, to complete a task by expressing varying levels of disappointment, which could potentially lead to emotionally charged decisions to reverse course and complete the suggested purchase. So in this example [INAUDIBLE] app, CARU determined that the pet feature in the game created a sense of urgency to purchase virtual cash with real money in order to save the player's virtual pet. So we looked at the language and the structure of the game.

And when a player clicked on the pet icon to visit a previously purchased pet, a red circle with a line appeared across the pet with the word SPCA written on it. And when the player clicked on that message it appeared to state, your pet is going to be taken away by the SPCA for animal neglect. Pay a fine of 6 Cash to keep your pet. So obviously, we find that to be rather aggressive and egregious behavior, especially in a child-directed product. You can go the next slide, please.

So this particular app, which was the Barbie Mattel Sparkle app-- and you'll see this, there are two examples here. This first one is where we see advertisements that are embedded with rewards, such as earning coins, tokens, or being able to either retry or advance to the next level. These activities are particularly endearing to children, due to the positive reinforcement that's being provided.

And in today's digital world, advertising is increasingly integrated into entertainment and editorial content. So consumers may not be aware that they are being advertised to. And this is why CARU's guidelines require that advertising should not be presented in a manner that blurs the distinction between advertising and program editorial content in ways that would be misleading to children.

Our concern here was that children on the app would not understand that they are being presented with these ads. And there were two types of advertising that were here. First, we saw short video ads that automatically run during a game play between levels. Once these are started, the videos could not be stopped until they were over, so the child had to view it in its entirety, with no ad disclosures present at any point to let the children know that these videos were sponsored content.

Secondly, in the top right corner of the home screen, there was a glowing icon that stated, watch advert-free gems. So if the user clicked this link, another screen appeared which said, don't miss out, watch a video and receive a reward, 1 gem. Again, no further disclosure for this type of video telling children that the video they're about to watch was an ad. You can go to the next slide, please.

So again, in the same case what was presented was another example of a dark pattern known as grinding, where the free version of the game is so cumbersome and labor-intensive that it is practically unenjoyable until new features are unlocked, which makes in-app play more seamless. The interference of ads and/or the encouragement to watch videos or perform other tasks or earn points, go to the next level, or the like can make this a very poor experience for users. But it should be noted that many players enjoy this type of competitive experience.

And Sam, before you go on to the next slide, I want to just make a brief comment about Snap, because CARU, we did have a case with Snap a couple of years ago. And I want to be careful about how we're addressing-- or we're not even really defining what's a kid and what's a teen. I think we're all working under the scope of COPPA and a child being 12 and under.

We did open a case with Snap concerned that their platform was child-directed and/or predominantly child-directed. And we found that they were not. And we did a thorough review, met with their legal privacy and safety consumer teams, their customer teams, and learned, not only are they not directed to children, but they make great measures to discourage children. So I just want us to be careful about how we're, I think, using a broad brush to just talk about kids and teens on these platforms.

And I think it's also worth noting that, oftentimes, when companies release their platforms or their products, they have one intent, and then what we've seen is that users' behavior can drive a product in a different direction that the company never intended. And I would suspect that with these Snap streaks that may have been what happened, and that Snap is likely reconsidering how to address this, again, based on the behavior that the environment has turned this into.

JOSH NELSON: I just want to respond to that briefly to say that, I can't speculate on what Snapchat may or may not be doing to address the concerns that we and others have with their platform and its usage by kids. But I think we do need to acknowledge that the kids do use the platform.

DONA FRASER: Agreed. And I think it leads to a larger conversation-- and we'll likely get into this-- about, what are we doing as an ecosystem to help inform and educate kids and their parents about what's appropriate and what's not appropriate? I do know, in our conversations with Snap, they make every effort to discourage kids from joining. So there's only so much I think companies can do. We cannot lay all the responsibility on them. I think we as an entire-- and I mean we.

I mean COPPA safe harbors, I mean educators, I mean parents, I mean children. We all need to take a role in how we teach kids and teens about what they're seeing online, how they're engaging. We make great efforts to teach children about their behavior in an offline environment. We need to do the same thing in an online environment.

SAM JACOBSON: Josh, do you want to make one last point before we move on?

JOSH NELSON: While that's certainly true and parents absolutely need to be vigilant, there's a broader societal problem here, that obviously there's a role for regulators like the FTC, for policymakers like Congress to address this, and for the tech companies themselves. Some of what you're saying about sort of pushing the responsibility for dealing with these broad problems down to parents sounds very much like what we hear from the tech companies themselves when they're deflecting responsibility.

DONA FRASER: That was not what I'm suggesting. I actually did say the entire ecosystem and defined the entire ecosystem, as everybody from regulators on down. We all have an equal responsibility to make sure that children and teens are safe in an online environment.

SAM JACOBSON: I'm going to jump in. I appreciate the discussion here. I think there's a shared concern that these techniques can be harmful to kids, whether or not they're intended by the companies who are employing them. And there's obviously some disagreement about what should be done about it. And we'll have a chance to take that up a little more later. But I do want to give Dr. Radesky a chance to present her example. So I'll turn it over to her.

JENNY RADESKY: Great. Yeah, so my lab has been starting to study and describe dark patterns within the context of how children might perceive them and process them in mobile apps and platforms. And we have seen everything that Josh and Dona just described, in terms of the daily rewards. It's not-- it doesn't just have to be an example on Snapchat. We see the, come back every day and get a special reward or prize or pack or something else in so many different apps.

And so part of our goal is to try to describe these patterns so that they can be quantifiable, they can be measurable, they can be something that we either are recommending-- don't go into apps in the first place if we know there are child users-- or we can help families say, OK, here's one of these tricks, one of these gimmicks. Like if the apps that your kids are playing have these and your kids don't want to put the tech down, and you're fighting about it all the time, then this is something that you can change behaviors about. But we've mostly been studying these to quantify these dark patterns when it comes to tactics that are trying to get children to play for a longer time, download new apps from ads that pop up, or making purchases. So the examples I'm going to show largely relate to parasocial relationship pressure and some navigability constraints that will come in future slides.

So this main typology of parasocial relationship pressure is when a character that you've built a relationship-- in this case, it's Kick the Buddy, which is a really popular app amongst the three to four-year-olds. It's a general audience app, but it is being played by lots of kids. And it's advertised as a stress relief app, where you beat up a doll, Buddy. And so you beat him up in lots of different ways.

If you get in-app currency, you can buy things to beat him up with or to shoot him or other things. So this character actually creates a pretty conflicted relationship, disorganized attachment with the child in the first place. Because he's both funny and making you laugh and seeking out interaction, but at the same time, when you start to hit him, he tells the users to stop hitting him. He taunts and shames when you go to the store to encourage purchases, saying things like, don't just stand there, buy something, or baby need a new pair of shoes, or all these other kind of-- they're kind of funny, they're kind of coercive.

And at the same time, you get this array of different purchase-- to buy gold or virtual dollars at different varying ratios, which is confusing to a young child of what they would represent. So I played this with my seven-year-old, who went right to the, let's buy a big bank of gold. Like it's just-- it's such a lure. And it is that combined with the parasocial pressure. I'm sure it's hard for some kids to resist.

The next slide are some examples of characters urging a child to play more. So on the left is a screenshot of an ad for Kick the Buddy. That popped up while I was playing another app. So the tiny X in the corner hasn't shown up yet. And Buddy-- you're waiting for 20 to 30 seconds while Buddy is taunting the user to flip that switch and hang him. So this is a coercion to actually engage in an app that's not age appropriate, that is a new kind of design technique we're seeing a lot in advertisements, where they're kind of expecting you to play.

On the right side of this screen, this is a home screen of an app where these aliens are begging the user to play the game and push that big green button to rescue his alien friends from evil farmers who are torturing them. So it's similar to that like, don't let your cat go to the pound, don't let them kill my friends, please keep playing.

On the next slide, we call these navigability constraints, but it's just-- you feel trapped. You don't feel like you have a choice of where to go. So this is common when an ad pops up in the middle of your gameplay, which is pretty common. And you just don't feel like you have the autonomy to either navigate out of it. Or in this case, this is another prompt to get the child to play the app.

And in this case, you see this familiar swipe prompt that children follow all the time. When you do that, Santa, who's holding a gun, kills these onlookers. And then you're prompted to kill the witness, and then you're prompted to download Mr. Bullet. So again, it's another example of violent ad that has showed up in the middle of playing another general audience app.

And in many of these cases, even if you try to X out of it, it takes you right to the app store. So even though it's not-- it could be a free app. It's not asking for a purchase. It's asking for a download, which can extract data and feed into the metrics that allow mobile apps and app stores in this ecosystem to make more revenue.

SAM JACOBSON: Thanks, Jenny. So I do want to get on to a couple more questions, but I just do want to give Criscillia a chance. Do you want to weigh in briefly here on the examples that the other panelists have just been presenting?

CRISCILLIA BENFORD: Sure. Just briefly, one of the things I'd like to point out is that, in these examples, we see a combination of design strategies, which we learned earlier can, as I said, increase the power of an acceptance. So when a designer is making an interface on an app, they often have what are called OKRs, so objectives and key results. And I think that one of the things that is emerging right now that may be one of the things that we need to regulate, that if your key result is to increase time on site, that is probably not appropriate for kids. So if we go back to the YouTube autoplay example, we saw a design strategy that was using exploitative visual design to make-- to help entice the user to stay on site.

The design principle there, it's a persuasive design principle called reduction. So in the case of autoplay, the target action is, stay on site. And it is-- there's not even-- it's usually you would reduce to a single action. In that case, it's not even an action. You just stay in inaction, and you achieve the target behavior, which is to stay on site.

And as Jenny was talking about, these target behaviors can be, stay on site, they can be comment, they can be share, they can be download, they can be subscribe. And you can see different target actions for different ages of people. But I am very concerned about that are related to increasing engagement and increasing time on site that allows for data collection and also opens the door to overuse. It's been associated with physical problems and also psychosocial distress in some, depending on the app.

SAM JACOBSON: Great, thanks, Criscillia. I think that's a great segue into our next question. So some of the dark patterns we've been discussing have obvious financial implications for kids and their parents, especially if they lead to kids or their parents incurring charges or in-app purchases without their consent. But are there other harms we should be worried about here, such as harmful digital media usage, when it comes to dark patterns and kids? And Josh, I believe you have some numbers on this from some of the survey work that ParentsTogether has done, so let's start with you.

JOSH NELSON: Absolutely. Dependence on tech platforms and devices has real consequences for kids and their families, from conflicts in the family to declines in mental health. According to a Parents Together survey, 60% of parents are concerned that their kids are becoming addicted to devices. Another 43% of parents also think the excessive use of devices has caused conflicts in their family. And 39% think it is leading to declining mental health in their kids.

In total, for various reasons, a full 82% of parents are concerned about how much time their kids are spending online. Part of that concern stems from the fact that the YouTube algorithm, combined with autoplay, can push kids and teens down a rabbit hole of white supremacists and other extremist content that isn't suitable for anyone, particularly kids. The same survey found that 67% of parents reported feeling concerned about kids being influenced by white supremacist content they encounter online. And 80% of parents reported being concerned about which social media platforms their child visits in light of the January 6 attack on the US Capitol, which was planned in part on social media platforms.

SAM JACOBSON: Thanks, Josh. Dona, do you want to chime in on this question, the question of harm, generally?

DONA FRASER: Sure. I think one of the concerns that we have is about the forced creation of profiles in kids apps, which, in and of itself, might be considered a dark pattern of sorts, along with the possible long-term tracking of kids' behavior. So we work very closely with companies to discourage the creation of profiles. And obviously, tracking of child's behavior is not something that we're encouraging companies to do.

Or actually, I think where companies do find themselves tracking, it may be something that is not explicitly personally identifiable information. However, the combination of data and how some bad actors are using this data is obviously a concern of ours. So we're obviously discouraging companies to create any kind of profile that's related to a child, and I think, especially, if we're looking at 10 or 11-year-olds, 12-year-olds and how they're transitioning on platforms from being an underage child, as it pertains to COPPA, and then being 13 and up and how that profile follows them.

SAM JACOBSON: Thanks, Dona. And for those in our audience, I realize we keep saying COPPA. COPPA is the Children's Online Privacy Protection Act. And it generally requires parental consent to collecting kids' data on various services. And kids are defined currently in COPPA as being under 13 years of age. Criscillia, do you want to weigh in on that, as well, on the privacy harms that kids face?

CRISCILLIA BENFORD: Yes. In an earlier panel, I believe it was Ryan Calo who talked about a practice called persuasion profiling. And this is one of the things that I worry about with children. And again, back to apps that are trying to increase time online, that allows opportunity for surveillance of behavior, which then there are all kinds of things about a child that can be inferred by the way that they interact with the site. So this is something that's very difficult for kids to understand. For a kid playing the game, it's fun, or I'm watching YouTube, it's fun. But underneath in the back end, there is data collection that can be used, again, to supercharge other persuasive design and adversarial design practices.

SAM JACOBSON: Thanks, Criscillia. I'm conscious of time here, but I'll ask Jenny to just weigh in, as well, on this question of harms. And then we'll turn to our final question and maybe take a question from the audience.

JENNY RADESKY: Yeah, I have strong theoretical concerns about the social harms of getting our children used to or habituated to playing in digital spaces that regularly involve coercion or constraints on their autonomy. But I'm also a researcher, so I care about the empirical data. What we're finding in our cohort in Michigan elementary school kids is there's high rates of distraction from learning. So lots of kids are really just checking these high engagement platforms frequently throughout the school day. And that's detracting from their learning.

We're finding that time spent on mobile apps, specifically, is correlated with less sleep and more problematic media use behaviors that have more of that behavioral addiction kind of profile to them. And then we're also finding the more time kids are spending in digital spaces without the protections that they need, we found that only-- it was 3% of our sample, the parents reported the child had been contacted by an adult stranger online during remote learning. But 90% of those were on remote learning, not in school classrooms. So it's because they're just spending all this extra time online in spaces that don't always know that they're there.

SAM JACOBSON: Thanks, Jenny. So the last question I'd like to take up, which, in my view, is both the most important but also the most challenging, is what should kids, parents, educators, companies, and policymakers be doing to address or mitigate these harms? And I'll just ask each of the panelists to just weigh in here for a minute, because we're getting to the end of our time. And we'll just have to raise these questions for later. So, Dona, let's start with you.

DONA FRASER: Sure. So I recognize that, as you address what COPPA is, that my fellow panelists may advocate for new COPPA legislation to increase the age from 13 to under 16 or 17, that asks that we remain thoughtful about rushing to judgment on this. In the 20 years since the inception of COPPA, we still have a lot to do, a lot of work to do in educating the ecosystem. And I'd say, let's start with a concerted effort and funding for teaching digital literacy starting at a very young age and a real curriculum in schools. Again, I think that the ecosystem has a role here to play. We need to figure out who does what.

We're concerned about the teen space. However, let's not treat 16-year-olds like six-year-olds. We don't do it offline, so I suggest that we don't start doing it online. Teens are vulnerable. But again, I'd ask that we be thoughtful about the risks and harms that we're seeking to protect these vulnerable groups from, many which I contend can possibly be addressed by self-regulation models.

I'll close by saying, I think I'm biased, because for the past 15 years, I've had the distinct pleasure of working with companies that really want to do the right thing. They may not always know what that is, but let's not punish everyone, including the users, on behalf of bad actors.

SAM JACOBSON: Josh, I see you're eager to jump in here, but I'm going to give Jenny a chance first to take up this final question.

JENNY RADESKY: Sure, I would say I'm intentionally not saying what parents need to do to, because, as an equity issue, to ask each individual parent to navigate opaque digital spaces and manipulative design is just bad public health. It leads to more inequities. So I'd rather that there be more what the FTC is doing to enforce mobile advertising, mobile data collection, that's been found in recent research. And then to build upon what Dona said, is that, yes, there's plenty of people trying to do good things at tech companies.

But what we find in my research is that we get the feeling that no one is monitoring what's actually landing in kids' laps. You can have all the intentions in the world, but to not actually know and be able to monitor at scale what are kids using, what are the design techniques we use in those, what's the impact of that. If the FTC or other regulatory bodies could find ways to have developmentally sensitive ways to monitor what kids are doing, it would give us such a better feedback for tech companies, rather than feeling like we're flying blind.

SAM JACOBSON: Thanks, Jenny. Criscillia, do you want to weigh in?

CRISCILLIA BENFORD: I'd like to echo what Jenny was saying. And I'd like to add to it that, when we think about these regulatory bodies, one of the things that I think is needed is to have people who are conversant in these design practices. So interaction designers, human computer interaction specialists, and other technologists who can talk about the back end, we need them involved in how we design the regulations, and not just so that we can have laws that are contemporary with the state of the art, but also so that regulatory bodies can anticipate what is in store for us in the future, and make sure that we don't stay behind. Right now, we're behind.

SAM JACOBSON: Thanks, Criscillia. And Josh, I'll give you the final word here.

JOSH NELSON: Thanks, Sam. I agree with Criscillia and Jenny. As I started to say before, while I think there are things that parents can do about screen time, generally, and to control the types of platforms their kids are using, the use of dark patterns and other design features to manipulate kids is a structural problem that needs to be addressed by regulators and elected officials. So parents should absolutely be vigilant. They should make sure they're familiar with the platforms their kids are using.

But the responsibility for addressing this broader societal challenge of big tech companies manipulating kids online is a job that needs to be tackled by regulators, like the FTC, and policymakers like Congress. The FTC specifically should create an easy mechanism for parents to file complaints about big tech companies and do everything in its current statutory authority to crack down on these types of design features and keep kids safe online.

We've talked about COPPA a little bit. Congress also needs to modernize COPPA and amend the law to offer protections to all kids up to age 17. Congress should also pass the Detour Act, which would address dark patterns broadly. And while there are things big tech companies can and should do, like turning off addictive design features, like YouTube autoplay, and ensuring that kids aren't using their platforms with fake birth dates, there's no one I trust less than big tech companies to get this right. They have proven over and over again that they simply do not care about kids well-being. That's why the responsibility for cracking down on dark patterns that harm kids and adults really falls on regulators, like the FTC, and policymakers like members of Congress.

SAM JACOBSON: Thanks, Josh. So I think we're at the end of our time here. Unfortunately, we won't have time for questions from the audience. I want to thank everyone again. This was a really terrific panel. I hope you all enjoyed it as much as I did.

I want to personally thank our panelists for a terrific informative discussion. We're going to take a short break now, and reconvene at 3:00 PM Eastern for our final panel, which will focus on how we can continue to address dark patterns and potential strategies for dealing with them down the road. Thank you all.

JENNY RADESKY: Thank you, Sam.

REID TEPFER: Welcome back, and thank you for joining us for the final panel of the day titled, "How Can We Best Continue to Address Dark Patterns? Potential Strategies for Dealing with Dark Patterns." My name is Reid Tepfer. I'm an attorney in the FTC's Southwest regional office. This panel will focus on the work that the FTC and other organizations and agencies have done to combat the problem of dark patterns, enforcement challenges that dark patterns may present, and potential strategies for dealing with dark patterns in the future.

Now, if time permits, I'll try to incorporate questions from viewers. So if you'd like, please take the opportunity to submit your questions to darkpatterns@ftc.gov. To discuss these issues, we have an impressive panel of experts. I'll try to keep my introductions brief in the interest of time, but I highly recommend reviewing their full bios on the event page to learn more about their remarkable work.

First is Brigitte Acoca. She joined the Organization for Economic Cooperation and Development's Directorate for Science, Technology, and Innovation in 2005, where she supports the activities of the OECD's Committee on Consumer Policy. Next is Laura Brett. She oversees the daily operations of the Better Business Bureau National Program's New York office and leads the BBB National Programs' National Advertising Division, which is the advertising industry's system of self-regulation founded in 1971 to boost consumer trust in advertising.

Dr. Maureen Mahoney is a senior policy analyst at Consumer Reports. Her areas of focus include state data privacy, security, and data breach notification legislation, state right to repair legislation, and robocalls policy. Jennifer Rimm is an assistant attorney general of the Office of Consumer Protection at the Office of the Attorney General for the District of Columbia.

Professor Lior Strahilevitz is the f Austin Professor of Law at the University of Chicago, where he's taught since 2002. Professor Strahilevitz has written dozens of law review articles and several books focusing on data privacy law, property law, consumer contracts, Fourth Amendment law, and law and technology. And finally, Professor Lauren Willis is associate dean for Research and professor of law at Loyola Law School in Los Angeles. Her scholarly work concentrates on the intersection of consumer psychology, market structure, and law.

Now, before we get started, I want to note, on behalf of myself and the panelists, that the views we express today are our own. They don't necessarily reflect the views of the commission or any particular organization. Now, with that said, let's jump right into our first topic, which is the work that has already been done to combat dark patterns.

Thanks in no small part to the work of our panelists here today, there has recently been increased interest in digital dark patterns and how user interface design can be used to influence consumers' behavior, in some instances to consumers' detriment. But as has been observed throughout the day, dark patterns are not new, neither to the market nor to state and federal regulators, including the FTC. And while we have not always had the term dark patterns, the FTC has been working for years to combat deceptive and unfair design practices.

The FTC has done so using Section 5 of the FTC Act, as well as statutes like the Restoring Online Shoppers' Confidence Act, which requires that sellers of negative option subscription plans clearly and conspicuously disclose all material terms and provide a simple mechanism to cancel. In addition to these many enforcement actions, the FTC has published enforcement guidance on online disclosures, disguised ads and endorsements, all of which bear on a number of the patterns that researchers have documented.

So I'd like to start then by asking Dr. Mahoney and Professor Strahilevitz if they would please highlight for us a few of the notable FTC cases in this area. Dr. Mahoney, would you mind starting us off?

MAUREEN MAHONEY: Sure, thanks so much. And thanks so much for having me here today. I'm really honored to talk about dark patterns, which unfortunately are quite widespread, and more needs to be done to address them. But the FTC does have some tools in its toolbox to do so. For example, in 2011, the FTC took action against Facebook for engaging in dark patterns.

The facts of the case are that, in December 2009, Facebook moved forward with an update that overrode users default privacy preferences, including with respect to hiding your friend list and also hiding your profile from people who might be searching for it. Facebook also led consumers through a confusing interface that led them to believe that they could restore those defaults, even though that wasn't the case. So this has the potential to cause consumers real harm. For example, someone could deduce the consumer's location from that of their friends. Also, if an employer is looking for you on Facebook, who knows what they might find when they go searching.

So the FTC reached a settlement with Facebook, requiring them to develop a comprehensive privacy program and requiring consent of users for any updates to override the privacy preferences that they had established. However, bad practices continued. In 2019, the FTC took action against Facebook, including for engaging in dark patterns, for example, collecting phone numbers, purportedly, to enable two-factor authentication. But that information was also used to target ads, also with respect to the facial recognition technology that was used to tag consumers in pictures. In the privacy policy, it says it has to be done with the active consent of users. But in fact, oftentimes, it was on by default.

So I think Facebook is a good example, not only of a company using dark patterns to further its business practices-- for example, bending consumers' practices in order to enable their growth-- but also, even though the FTC does have authority to go after these companies, it can be hard to really reign them in.

REID TEPFER: Thank you, Dr. Mahoney. And Professor Strahilevitz, could you give us your thoughts?

LIOR J. STRAHILEVITZ: Yeah, sort of starting with the cases that Maureen talked about, I thought it might be helpful to identify a couple of Federal Court of Appeals cases, where the federal courts have generally backed the FTC when the FTC has brought cases alleging that dark patterns were occurring in ways that deceived consumers. Neither of these cases use the vocabulary of dark patterns. But they're both relatively recent published appellate opinions, where, essentially, it's the dark pattern that's generating the Section 5 enforcement action.

So the first such case is one that's been in the news lately, albeit for different reasons, which is the Ninth Circuit's opinion in the AMG Capital Management case. Now, I think people in this audience know, the Supreme Court just handed down a decision with respect to the kinds of remedies that the FTC is able to pursue under Section 13(b). But that's not the portion of AMG Capital Management that I want to talk about today.

There's really interesting language in the Ninth Circuit opinion with respect to, not remedies, but the question of what conduct violates the statute. And so a number of the things that AMG was doing that the Ninth Circuit felt were violations of Section 5, they were engaged in forced continuity. They had intentionally confusing prompts. They preselected default terms that were much worse for consumers than the other options available. And they made use of visual interference and fine print.

The Ninth Circuit looked at all of this and said, well, even if there is boilerplate language in the terms and conditions that is, quote, "technically correct," end quote, that that's not going to allow AMG to escape liability under Section 5. I think another really helpful precedent is from the Second Circuit that's Led Click Media. And Led Click Media involved an entity that was selling colon cleanses and weightloss products.

And their site was populated both with false testimonials for consumers that had been concocted by the company and with content that was designed to look like independent journalism, designed to look like newspaper, online newspaper articles, extolling the virtues of the product at issue, when in fact, this was fake news. It was advertising content disguised to look like news. And the Second Circuit had very little difficulty in finding that these kinds of false testimonials or misleading content constituted violations of Section 5, that this content was deceptive within the meaning of the FTC Act.

REID TEPFER: Thank you, Professor. States have laws prohibiting unfair and deceptive trade practices. And they've used these statutes also to address dark patterns. Now, Jennifer, your office specifically has brought several cases targeting dark patterns. Could you tell us a bit about the DCAG's recent work in this area.

JENNIFER RIMM: Yes, thank you. And thanks for having me today. I do want to preface my remarks by noting that any views express today are my own and not those of the Office of the Attorney General for the District of Columbia. Our office has brought several consumer protection cases that implicate dark patterns, although we have not necessarily used that vocabulary to describe the deceptive practices.

Broadly, these cases involve online transactions, where a dark pattern supports or is one of a number of omissions or misrepresentations about information material to that transaction. In most of those cases, the omitted or obscured information specifically related to the price charged for a product or service. I'll just also caveat, because these cases are in active litigation, I can't get into too many details. But I can describe the allegations in the complaints that we filed.

So right now, we have a pending lawsuit against the food delivery company Instacart, in which we've alleged that the company added a 10% default charge to customers' bills. And one of our allegations is that the way this charge was presented to consumers deceived them into believing that the added charge was a tip to the delivery person or delivery people. In fact, this charge was an added amount that went to the company. And interestingly, consumers could waive the added amount if they interacted with the link on the checkout interface.

But we alleged that, in many cases, consumers didn't take this step because they confused the 10% amount for a default tip, which, prior to the implementation of this practice, had also been 10%. Part of our allegations include a default or recommended tip-setting, as well as an obstruction pattern, whereby consumers wouldn't learn that the fee could be waived or what it was or that tipping was a totally separate option unless they clicked on a link. And even then, the explanation that the company provided to customers for this 10% fee was confusing and furthered consumers' misunderstanding. I'll just note, this isn't the only example of a lawsuit where we've alleged this kind of practice. Last year, we brought and settled a similar case against the food delivery company DoorDash, with a somewhat different fact pattern.

We also have an ongoing suit against Marriott, which involves mandatory amenity and resort fees that are not added to the advertised prices of a hotel room. So this is a classic bait and switch dark pattern that interferes with the consumer's ability to price compare. In addition, the complaint alleges a number of other deceptive design strategies that lend to the consumer deception by obscuring what the fee was.

So for example, we alleged that Marriott used a smaller and lighter type face when referring to the resort fee and used confusing language regarding whether the resort fee had been added or would be added to the room rate. And I'll just note, we also are pursuing a case against Facebook that touches on dark patterns, namely hiding privacy-related interfaces. We filed the suit in the aftermath of the Cambridge Analytica scandal.

And one of our allegations in that case is that Facebook varied a setting that allowed third-party apps to access your data if your friends used the app. So we think what you see in these cases is that our enforcement actions involving online platforms, while they continue to involve traditional statements, false statements, and omissions, it's increasingly necessary to also take into account design elements that work to deceive and harm consumers.

REID TEPFER: Thank you, Jennifer. Dr. Mahoney, in addition, many state legislatures are considering or have already passed legislation that specifically addresses dark patterns. Could you tell us a bit about these state laws and bills?

MAUREEN MAHONEY: Sure. So policymakers' willingness to prohibit dark patterns in some of this new and pending privacy legislation is probably one of the most exciting aspects of what's going on in privacy legislation in the United States, in my opinion. The California Consumer Privacy Act is, in my knowledge, the first privacy law in the United States to explicitly prohibit dark patterns in opt-out processes. So that went into effect in 2020, giving consumers the right to access, delete, and stop the sale of their information. Also, opt-in consent is required for the sale of kids' information.

So recently, the attorney general, which has the authority to issue rules to implement the CCPA, finalized rules saying that dark patterns are explicitly prohibited in the processes through which consumers are opting out of the sale of their information. In their words, you can't have processes that are designed with a purpose or have the substantial effect of supporting or impairing a consumer's choice to opt out. So that's good.

Another thing that the legislation does is that it directs the AG to design a uniform logo to help draw consumers' attention to the opt-out option on the homepage of the website. They're also pretty prescriptive about what the opt-out link can look like. So I think that's a good thing, as well, something that policymakers may need to do in the context of an opt-out law, put real guidelines on what companies can do in terms of their design choices.

And this is really important because Consumer Reports has found patterns in CCPA opt-outs. We had hundreds of consumers try to opt out of the sale of their information at data brokers listed on the California data broker registry. And oftentimes, it was really hard. About 14% of the time, burdensome or broken opt-out processes had the effect of preventing consumers from stopping the sale of their information, which was really frustrating for consumers.

They were upset that companies asked for things like their social security number in order to opt out, even though identity verification isn't required. And consumers felt really uncomfortable giving their SSN to a data broker that they didn't know. Also, confusing popups, asking for consent to cookies, the consumers really didn't understand and were frustrated by.

Another thing I want to point out is that Proposition 24-- that's a ballot initiative to amend the CCPA. And that was ratified by voters in the fall and hopefully will go into effect in 2023. That explicitly prohibits using dark patterns in obtaining consent to opt back into the sale of your information after you've opt out, or for providing consent for kids to share or sell their information. So I think that sets important precedent. But it's also important to note, this only relates to the California Privacy Law. It's not going to do anything for other dark patterns, like hidden fees. And it's also only going to be as effective as enforcement.

REID TEPFER: Thank you, Dr. Mahoney. And Laura, next, if you could tell us, the BBB National Programs and other advocacy groups are playing an important role in this space, so could you tell us a bit about the work that the BBB National Programs is doing concerning dark patterns?

LAURA BRETT: Sure. So first, let me tell you a little bit about our role. So BBB National Programs, the National Advertising Division and some of our sister programs function as independent advertising self-regulation. So what that means for the National Advertising Division is that we open cases, and we review advertising, and hold it to standards that are primarily set by the FTC, that advertising be truthful and transparent. We do this by opening cases on our own initiative and also opening cases brought by competitors as a challenge to competing advertising claims that are skewing the marketplace.

It is a voluntary process, but we see 90% to 95% voluntary participation and compliance with our recommendation. In large part, that's due to the strong support that we receive from the FTC. The FTC regularly follows up on referrals from us when companies do not participate or comply with our recommendations. So we really do appreciate that. And I appreciate being invited here today to talk a little bit about the work that we're doing, so thank you.

So it's important to remember that independent self-regulation, like we have at NAD, does not prevent fraud. But it can be really effective at preventing widespread misleading practices from creeping across a marketplace, like what we're seeing, these dark patterns. And it is a good method to expand compliance with FTC guidance on these issues.

What it does is it allows both, thinking of the forum itself, to hold companies to FTC standards, but also to have companies that are trying to comply with FTC guidance to have a forum to challenge practices that their competitors are using in the marketplace that are having an impact not only consumers, but also on competition. A good example of what we've done in dark patterns is a case we opened recently. And again, we never called it a dark pattern.

But we were looking at, specifically, advertising for discounts on fitness wear for VIP members. Now, a VIP member was not-- it was not a loyalty program, but, in fact, was a subscription program. So instead of getting points for purchases, you actually had to subscribe to monthly shipments of fitness wear at $49 a month. So obviously, if VIP membership requires a subscription, that's a material term that requires disclosure.

We think there's FTC guidance on that that's fairly clear. But what this company was doing was they were [INAUDIBLE] their ads to consumers, sometimes in emails, but often online, telling them that they would get this great discount with a VIP membership, and then bring them to their website, where you'd see the kind of purchase flow we've been talking about, where you're presented with items to purchase, but first you actually have to take a quiz. So they kind of draw you in and keep you on their site for a while and talk about your fitness practices and your preferences for fitness wear. And then ultimately, you'll be given a choice for products to buy.

And when you go to click on one of them, you can check out as either a VIP member or as a regular customer, with the VIP membership being highlighted in bright red, where the regular customer purchase is in gray, and the VIP membership price obviously being much more appealing, like $19 a month as opposed to like $59-- or like $19 for the outfit as opposed to like $59 for the outfit. On the final purchase screen is where you get disclosure about what the terms of VIP membership were. And in fact, even that disclosure wouldn't be right up front at the top of the screen, but it would require, at least in some interfaces, scrolling down to see.

We asked this company to come in and talk to us about their marketing practices and got them to comply with what we felt was FTC guidance about clear and conspicuous disclosure. So I do think that dot-com disclosure guidance the FTC put out several years ago provides a lot of guidance on this issue. And we got this company to participate in the process and comply with our recommendations and make disclosures up front at the time they make the discount offer about what a VIP membership means.

Now, we'd like to see competitors challenge these kinds of dark and misleading practices, these deceptive interfaces. We have seen some challenges around that. We have seen challenges around that misleading initial approach, like luring in a consumer with a misleading interface or not disclosing that the speaker was an advertiser, making it look like editorial content. We have actually seen plenty of challenges in that space.

We've also seen some challenges around disclosures and failure to-- obscuring or failure to disclose conditions and fees. But we have not seen as many as you would think, given the widespread use of these practices in mainstream commerce. And that is, of course, a concern for us because we would like to see industries police themselves instead of dropping down to the bottom, so that that practice to self-regulation to police the marketplace and bring in better practices that are not misleading to consumers, and frankly, more fair for competition.

REID TEPFER: Thank you, Laura. The problem of dark patterns does not stop at our border, of course. Brigitte, could you please tell us about the international regulatory landscape and the work that the OECD and our partner agencies have done to tackle this issue?

BRIGITTE ACOCA: Sure. First, let me thank you for giving me the opportunity to discuss what the OECD commission has recently launched on [INAUDIBLE] online. And thank you for having me today. This is preliminary work. And from our discussions with our committee members, we've seen that most OECD jurisdictions have general consumer laws in place, or sometimes specific laws, as well, covering many of the techniques associated with dark patterns. And while some of these practices are new, consumer authorities around the world have addressed for quite some time several practices that may be considered dark patterns, such as, for example, [INAUDIBLE] pricing or subscription [INAUDIBLE].

And I just wanted to mention about, for example, the European Commission, which has recently indicated that many dark patterns can be in breach of the EU's Consumer Rights Directive or the [INAUDIBLE] Commercial Practices Directive, which includes a blacklist of specific practices that are prohibited in all circumstances. And such blacklist, interestingly, specifically refers to a few examples of dark patterns, for example, falsely claiming that a product is available only for a limited period or for continuity purposes. However, to provide further clarity about how EU consumer laws may apply to a broader range of dark patterns, including new forms of dark patterns, the European Commission has decided recently to update its guidance in this area.

I also wanted to note that a number of consumer authorities have actually made use of existing consumer and data protection laws to take a number of enforcement actions, again, in many different countries. For example, in 2019, a coordinated action was initiated among several EU member states and the European Commission, again, against several booking sites, who use misleading selling techniques, involving scarcity cues, even charged with misleading discount claims.

Another example in 2017, the Australian Competition and Consumer Commission took action against a ticket reseller for using misleading scarcity cues. And I also wanted to mention about an important internet sweep of almost 2,000 retail websites and apps that was conducted in 2019 by the International Consumer Protection and Enforcement Network, which found that close to 1/4 of these websites and apps featured dark patterns. And the top three were being-- were pressure selling, drip pricing, and design issues, such as hidden terms and conditions.

This sweep confirmed the need for more guidance to business on dark patterns and for education of consumers. And I wanted to mention that several authorities have, since this sweep, actually issued guidance, in particular, guidance for businesses-- for example, the Dutch Consumer Authority, which last year issued guidelines for businesses on how consumer law applies to dark patterns, and which provides examples of practices that may or may not be permitted. Another example is guidance issued in 2019 by the French Data Protection Authority containing the [INAUDIBLE] psychology of potentially deceptive design practices.

Finally, I wanted to note that some agencies, such as, for example, [INAUDIBLE] Consumer Authority, or again, the European Commission, will be conducting, very shortly, studies involving a mapping of dark patterns, which will be followed by behavioral experiments to test the effects of [INAUDIBLE] patterns on consumers. And we at the OECD, as part of our new work on dark patterns, will also carry out empirical work in this area. We'll take stock of the situation in different countries. And our goal is really to identify those key dark patterns that should be a policy priority globally, to precisely avoid a fragmented international regulatory landscape in this area.

REID TEPFER: Thank you, Brigitte. So let's move now to our next topic, enforcement challenges. As we learned in the earlier panels of the day, many dark patterns have analogs in the brick and mortar environment. That said, the scale, sophistication, and personalization of dark patterns may present new or different enforcement challenges.

Jennifer, I'd like to start with you. Could you please tell us about some of these challenges?

JENNIFER RIMM: Sure. I'll start by saying that I agree that the sophistication and scale of these practices presents enforcement challenges. As some of the panelists have observed, dark patterns are a relatively common problem in contexts like online shopping and privacy settings. And the proliferation and sheer variety of dark patterns on its own makes it more difficult for government enforcers to effectively identify the most severe offenders and to address these offenders through targeted enforcement actions.

Added to this, as has been discussed today, a feature of dark patterns is that they are covert or deceptive, meaning that consumers may not realize they are being manipulated, or they may be inured to these dark patterns due to their prevalence. So it's possible that these dark patterns are less likely to spur consumers to make reports that they've been misled or deceived. And this could potentially be an even greater concern with respect to dark patterns in the privacy context, where the consumer is not subject to something like an unexpected credit card charge that would trigger some awareness of having been opted into something they didn't intend.

So as a result of this, as well, problematic dark patterns can be harder for enforcers to identify. And this is one reason why the research that's been discussed today and work done by organizations like the Norwegian Consumer Council is very important. Related to this, because of their sophistication, dark patterns also present challenges when demonstrating that a particular set of practices was deceptive.

So as has been pointed out today, dark patterns can be refined from user data to make them more effective and harder to perceive. And the deceptive nature of a transaction, when it involves a dark pattern, can come down to the particulars of the wording and presentation of an interface and other related representations that are part of the entire transaction. Because the mechanics of the deception can be subtle, these can be cases that are going to benefit from expensive expert reports and potentially even consumer surveys or studies to show that design elements themselves have the effect of misleading consumers.

And if consumer if the consumer experience involves interacting with a number of screens, and if the screens change over time or across modalities, the challenges involved in this work are only compounded. And this is even leaving aside the more personalized targeting of dark patterns that we discussed today or the need to dig into underlying code to understand a deceptive practice, which can add an even greater level of complexity.

I'll also note that a number of dark patterns that we've discussed, depending on the case, might not squarely involve false statements or material omissions. So for example, [INAUDIBLE] the nagging dark pattern. Here, a successful prosecution could also require experts in fields like behavioral economics to get a court comfortable with and understand how consumers' cognitive biases are exploited or their behavior is shaped by design.

So just to sum up, I think these cases put different burdens on enforcers, in terms of the work needed to identify where to devote resources, and also the level of resources that might ultimately need to be brought to bear in order to really elucidate the deceptive or unfair nature of dark patterns.

REID TEPFER: Thank you, Jennifer. And Professor Willis, could you give us your thoughts on this question?

LAUREN E. WILLIS: Sorry, I just had to find the unmute. So first I wanted to just thank the commission for putting on today's event. The entire day has been really illuminating. And thinking about enforcement challenges, I really urge the commission to develop an approach that's forward-looking in nature. For tech, in the tech world, things are here before you know it. So I really see three enforcement challenges that are posed by unfair and deceptive digital practices. I've written about these in an article called "Deception By Design." And they really echo a lot of the things that Jennifer, from her on the ground enforcement experience, has talked about.

So the first is the proliferation of unique online marketing materials and sales interfaces. So AI today can produce and disseminate thousands of iterations of digital marketing materials, homepages, app flows. The Trump campaign, even back in 2016, apparently used over a million unique iterations of its marketing. And of course, enforcement agencies can't test every iteration with consumer subjects to assess deceptiveness. They can't even engage in a facial analysis of a representative sample when we're getting to these numbers.

I think some people think, well, let's just grab a few web pages at random and test them, and that should be sufficient. But as Jennifer was saying, subtle design changes make a difference. Was the key information at the periphery of a screen or in the color blue, both of which are harder for elderly consumers to visually perceive? What was the exact timing of an animation that might have distracted consumers from, say, warnings about the risks of the transaction?

And the FTC has recognized this in the Pom Wonderful case. The commission would not accept evidence of deceptiveness of print ads as evidence of deceptiveness of similar billboard ads. So proliferation is a major barrier to the FTC and other enforcement agencies, demonstrating the scale of deceptive practices. It's already proven to be a barrier in the case that the FTC brought against DIRECTV.

The second major barrier is microtargeting-- so microtargeting of unique iterations to particular consumers or consumers with particular associated data points in real time in real contexts. So different iterations of a web screen or an app screen isn't produced for the reasonable consumer in a neutral context. It's created for narrow types of consumers at their most vulnerable moments.

And the commission or an enforcement attorney can't fully appreciate the position of particular consumers at particular moments. Was the consumer in the flow of playing a video game, habituated to think a green button means continue, not buy now? Was the consumer incapacitated? So intoxication or cognitive impairment can be discerned from keyboarding patterns or mobile device handling. A commissioner, hopefully, when they're looking at these things will not be incapacitated at that time. And microtargeting also means that experts can't really reproduce the context in which the screen appear to consumers, I can't know which iteration would have been shown to which consumer, which type of consumer, and therefore can't sort of match the experimental subjects with the specific iterations they would have been shown. So the real world is likely worse than Professor Strahilevitz's experiments showed. And so extrinsic evidence of deceptiveness, which some case law requires when deceptiveness is not facially obvious, may be impossible to obtain.

The third barrier is automation. And I think this was discussed a little bit in the very first panel of the day, the automation of the microtargeting process and the digital design process that is made to optimize what [INAUDIBLE], I think, in the first panel called OKR, the objectives and key results, that the business seeks. So using big data and machine-driven experimentation, businesses, in real time, can determine what combination of data points about users and their context and specific design features of an app, a web page come together to create the consumer response that the business seeks. And when businesses optimize the process for sales, the digital materials will produce sales, even if the sales are based on deceptive or unfair practices.

And that matters for enforcement. Because not only are automated systems likely to be a bit better at deception than humans are, or learn to be better, but AI systems lack intent. Now, intent isn't required to prove a Section 5 violation. But it's very helpful in proving a violation. And as more of the control over the digital design process is given to the AI, humans at the business may not even know that they're violating the law.

So there's not going to be any internal ethics control, no whistleblowers coming out to say there's a problem here, no way to stop deceptive and unfair practices, unless the business is forced to proactively look, to metaphorically open that AI black box and determine, well, was that particular web page design directed at consumers with visual perceptual impairments? Was it designed in a way that exploits habituation of the consumer or distraction? And so those three things together are making and will increasingly make enforcement difficult.

REID TEPFER: Thank you, Professor. And before we move on, Professor Strahilevitz, I'd like to get your thoughts on this too.

LIOR J. STRAHILEVITZ: Yeah, I think maybe I'll just start right where Professor Willis left off, because I think she makes a number of important points. I guess I'll sound a somewhat more optimistic note about how much data can inform the regulatory and enforcement strategies that the FTC makes. I don't think the burden on the FTC is to show that each particular consumer was deceived or treated unfairly. And I don't think Professor Willis was suggesting that that's the burden.

What I've seen from our research and from our results is that, when we replicate our results-- and good social science should try and replicate, and we've replicated our research on dark patterns-- we do make minor tweaks. We're going to adjust our obstruction pathway, or we're going to change the nature of the confusing prompt. Or we're going to make hidden information a little bit more hidden or hidden in a slightly different way.

And generally, that doesn't have a big impact on the results. So what our data, at least, tells us-- and for sure, more data would be useful. And I know a number of social scientists are out there working on dark patterns, and that's fantastic. But at least when our data tells us so far is that it's possible to generalize and say things like, hidden information is likely to cause welfare losses for large numbers of reasonable consumers. And we can show that experimentally and mathematically. Or obstruction-- regardless of whether you tweak the obstruction pathway this way or that way, obstruction works.

And it works in terms of increasing purchase rates among our consumers every time we've tried it. And so I think that should alleviate at least some of the anxieties that Professor Willis rightly points to and suggests that, if the standard is, would reasonable numbers of reasonable consumers be deceived, that's an answer with-- that's a question with an objective answer that experimental results can shed light on.

I do think-- picking up on the point that both of my panelists just made, I do think that the context in which the information is presented is important. So in some of our research that Jamie Luguri and I have done that we haven't reported, we have tested whether our subjects are encountering dark patterns on big screens, like desktops and laptops, or on small screens, like smartphones. And in general, as you might expect, techniques like hidden information are much more effective on smartphones than they are on iMacs, for example, which have very big screens. Why? Well, because on a small phone, information towards the bottom is going to require a lot of scrolling. And so it's even less visually prominent than it would be on a laptop or a desktop monitor.

And then finally, I do think one other complexity that I'd like to throw into the mix are the parts of this regulatory challenge that are easy from a First Amendment perspective and the parts that are a little bit harder. So there's some kind of dark patterns that, I think, when employed have no First Amendment protection or minimal First Amendment protection. Let's say there's a false testimonial or false social proof. A company said, 3,000 customers have just bought our product, and it turns out that they just made that number up. OK, that's easy. There's no First Amendment interest in lying to consumers about how many widgets you've sold.

Some of the dark pattern strategies, though, do fall into grayer areas. And so there's just not a lot of case law on whether regulating obstruction would run aground of commercial free speech protections under the Central Hudson line of cases. Or is nagging protected by the First Amendment as a sales strategy? We just don't have a lot of precedent there. We know that under certain circumstances, nagging isn't protected by the First Amendment.

Someone who asks out a coworker on a date once, probably protected, unless there's a power disparity there. Someone who asks out a co-worker repeatedly despite refusals clearly isn't exercising their free speech rights. They're engaged in sexual harassment, if the requests are pervasive. So we do have some guideposts that we can look to from other areas of law. But I do think the First Amendment issues surrounding the regulation of obstruction, of nagging, of confirm shaming, and of certain kinds of subtle visual interference, those are questions that the FTC should spend some time thinking about and consulting with First Amendment experts as they try to regulate in this area.

REID TEPFER: Thank you, Professor. Let's move now to our final topic, which is future strategies for dealing with dark patterns. In recent years, a number of organizations, such as darkpatterns.org and the Norwegian Consumer Council, have sought to raise awareness of the negative effects of dark patterns by calling attention to particularly egregious examples in the market. So I'd like to ask the panel, do you think these naming and shaming efforts are having an effect on firms' behavior? And more broadly, can we expect self-regulation or the market to do any of the work of mitigating the harmful effects of dark patterns? And then lastly, is this an area where increased consumer education or technological interventions, like browser extensions, can be effective? Lauren, would you mind starting us off?

LAURA BRETT: Sure. So as I said when I spoke earlier, independent self-regulation can be effective at preventing dark patterns from spreading and cutting them back where they exist now. It's important to remember that if the FTC takes more action and there's more regulatory enforcement or has more tools, they're always going to have a limitation on the resources they can devote to any particular issue. And so then we have to look at the marketplace and see if the marketplace can expand those-- can adopt better practices. I think good, strong guidance from the FTC and good, strong enforcement from the FTC on some of these issues supports the use of self-regulation, generally, with companies.

What we've seen, generally, is when there is clear FTC guidance in what we see is an FTC enforcement priority, cases do come to us between competitors. It's important to think about how you leverage a competitive marketplace to expand better practices. So companies that are fearful of enforcement from the FTC may change their own practices. But we know that when they change those practices, they may make less money. They're engaging in these practices because they're profitable.

So if they see that they're following FTC guidance because they're worried about enforcement, they may bring challenges against competitors, and use this advertising self-regulatory forum that's been around for 50 years to do that. It's important to recognize that there are limitations on regulatory enforcement. And so we should be thinking about how to get marketplaces to adopt better practices on their own.

And self-regulation has been an effective way to do that for a long time. And we have seen that, where the FTC issues clear guidance, where the FTC sends out a lot of warning letters, we will start getting a spate of challenges from competitors. Because they clean up their practices and then use this independent self-regulatory forum to do the same thing for their competitors.

And that benefits-- that certainly creates a more fair marketplace and more level playing field. But there's a benefit to consumers too. So I hope that, in the spirit of calling for greater regulation or regulatory enforcement, we think hard about how to make sure that that expands beyond what we know regulators can do.

REID TEPFER: Thank you, Laura. And Dr. Mahoney, same question to you.

MAUREEN MAHONEY: Sure. I think it's an important first step, but it's not entirely enough. We do think it's really important to bring dark patterns into the light, consistent with the goals of this workshop. And that's been a motivating goal of what Consumer Reports has been trying to do since its inception. And we have, in the course of our product testing, encountered a lot of dark patterns in the marketplace.

For example, in testing smart TVs, we saw not just for consumers to share information, also in testing payment apps such as Venmo. We also looked at Facebook, following the steps of the Norwegian Consumer Council, in addition to other investigations. So that's trying to bring dark patterns into the light. That's the naming.

In terms of shaming, I mean, we've also called on companies like Facebook to improve their practices, like adopting ethical design procedures and to no longer use abusive practices to keep users engaged online. We've also engaged in a lot of consumer education through our magazine, for example, providing information to consumers on how to close your Amazon account or how to stop location tracking on Facebook.

Another thing that we've been doing, which I'm really excited about, is trying to help develop browser privacy signals, similar to Do Not Track, called the Global Privacy Control, so that consumers, in a single step, through their browser, can indicate to every website that they interact with, send them a CCCPA compliant Do Not Sell signal. So consumers don't have to go to sites one by one and potentially encounter dark patterns in order to opt out.

So I think all of those things have a role. But in our experience, progress has been limited, and it's just not enough. You do need clearer standards, prescriptive guidance from the FTC, obviously, strong enforcement. And naming and shaming really isn't enough when you're talking about the strong incentives for companies to use dark patterns. We're talking about billions of dollars here.

REID TEPFER: Thank you, Dr. Mahoney. And Professor Strahilevitz, could you give us your thoughts, please?

LIOR J. STRAHILEVITZ: Yeah, I would start with the core idea that consumers are generalists, not specialists. And they just don't have time to be on the lookout and constantly be on their guard whenever they're interacting with a device. And so I really think that consumers are going to be empowered, and frankly, more willing to spend more money online if the FTC is smart and aggressive about trying to rid the marketplace of the very worst dark patterns, those that are most likely to cause consumers harm.

And I think I'd-- for that reason, I do think self-regulation is valuable. Self-regulation hasn't done nearly enough so far to police the worst abuses of dark patterns. And so I'm not at all optimistic that self-regulation is a big part of the solution. Though, I think for reasons Laura mentioned, it will be valuable as a supplement to the work that the FTC, state attorneys general, even class action plaintiffs' attorneys are doing to keep companies on their best behavior.

The other thing that I think I can speak to, with reference to some of the studies that we've done, is consumer learning. So I do think consumers learn to a degree about dark patterns. And one of the results from our paper that I referenced earlier suggested that consumers at least have built up a defense mechanism against one type of dark pattern, which is the, if you don't act within 60 seconds this great deal is going to disappear. I think consumers over the years have come to associate that kind of sales practice with sort of dodgy, fly-by-night operators. And it may well be that, when consumers see that kind of blatant manipulative behavior, they react negatively, and it makes them less likely to purchase. Our data is consistent with that story.

On the other hand, consumer learning can actually make the problem worse rather than better. So I'll give you an example, which is nagging. A lot of us have gotten nagging prompts on our smartphones, with an app repeatedly asking us to share our location or repeatedly asking us to authorize push notifications. And if you say yes, it stops nagging you. And if you say maybe later, which is the only option you've got-- no, is not an option-- It'll ask you a week later and a week later and a week later, until you say yes. And consumers get worn down by this.

And eventually, they just say, well, they're going to get me eventually, so I may as well go ahead and agree to push notifications now. So sometimes then, consumers learn, but what they're learning is a learned helplessness. And that's really an example that cries out for aggressive government regulation. Because to characterize consumers as having consented after repeated nagging to do something that they don't want to do is to sort of make a mockery of the idea of consent.

REID TEPFER: Thank you, Professor. And before we move on, Brigitte, could you weigh in, please?

BRIGITTE ACOCA: Sure, thank you. Just on the first part of your question, with respect to raising awareness of the negative effects of dark patterns, I think it's really important to note that civil society and other organizations can really play an important role in alerting jurisdictions worldwide about the problem of dark patterns. And the international campaigns that the Norwegian Consumer Council conducted in this area with a number of other organizations in the European Union in 2018 have helped to raise awareness internationally about the problem of dark patterns, and even generated legal action in a number of countries. This is-- I think these are very important initiatives. And they prove to be actually very important during the-- since the beginning of the pandemic.

Second, we have seen bubbling up in different countries a number of interesting co-regulatory initiatives across countries developed recently that may help, in a way, to address dark patterns. For example, corporate digital responsibility initiatives, which are aimed to foster commitments from businesses to go beyond existing legal requirements and provide digital services in line with trust and fairness and transparency, these corporate digital responsibility initiatives have been put forward in several countries, including, for example, Germany and France. So these are preliminary, but it's interesting, I think, to follow the development of these initiatives in different countries.

We've heard, also, earlier today during this workshop that there is growing interest among user interface designers in some countries in developing standards for the ethical design of user interfaces. And finally, we've also heard about the promising role that technological solutions could play. And I just wanted to mention about a project that is being funded by the German Ministry of Justice and Consumer Protection, which involves the development of an app to detect dark patterns using AI.

I also wanted to mention work that researchers at Oxford University have developed. They've developed the tool allowing consumers to apply a variety of patches to remove dark patterns from their apps via using a simple [INAUDIBLE]. So these are interesting technological developments that might be useful to also follow. And as we've heard from our OECD members and nonmembers, there may be scope for consumer authorities, again, to collaborate and enter into partnerships with academics or businesses in developing these technological solutions, which we know are resource-intensive and can be very costly and difficult to put in place.

REID TEPFER: Thank you, Brigitte. So next, I'd like to discuss how the FTC and other regulators and enforcement agencies should prioritize their resources concerning dark patterns. Professor, Strahilevitz, are their dark patterns that you think are particularly harmful or otherwise deserving of scrutiny or attention?

LIOR J. STRAHILEVITZ: Thanks, Reid. I would be data-driven in answering this question. So at least the preliminary data that we presented earlier suggests that hidden information, that obstruction, that intentionally confusing prompts, that these are probably the most pernicious dark patterns, are the ones that are most effective in getting consumers to do what they otherwise wouldn't want to do. I think there's also some dark patterns that I haven't been able to test, like sneaking purchases into the cart.

We couldn't come up with a good way to test that. But my strong prior is that those are quite harmful to consumers, as well, and that other strategies, other dark pattern strategies like notions of scarcity, must act now, probably ought to be a much lower priority for the FTC.

The other thing I'd say is, you can't be everywhere as a regulator. And so to the extent that you see very large-scale, very well-respected entities employing dark patterns-- Ticketmaster, Sony-- going after them will get more attention than going after an entity that may be using dark patterns but isn't using dark patterns on nearly the same number of consumers as these large and very prominent entities are in e-commerce.

REID TEPFER: Thank you, Professor. Laura, do you have thoughts on this?

LAURA BRETT: Sure. I think I'm going to echo a lot of what Professor Strahilevitz just said. Very often, we see regulars go after kind of the worst practices. But here, I think the widespread harm is really coming from the widespread use of some of these practices. So going after mainstream companies that are engaging in these practices that are widespread, I think, will send a message, a strong message to companies broadly that these are practices that are against the law and the FTC is taking seriously and taking enforcement action against.

But I've also seen that FTC make very effective use of guidelines, issuing guidelines, specific guidelines directed at this, that can bring a lot of compliance around, and also can bring companies to police their own industry by using self-regulation. And I also should say that the FTC warning letters have been very effective in a lot of different contexts to send a message to industry that these practices are under scrutiny and the FTC feels that they're deceptive. So those are my two cents. And we would like to see more companies use self-regulation to enforce FTC guidance that we think is out there already, but could be strengthened and more particularized to the dark patterns we've been talking about today.

REID TEPFER: Thank you, Laura. Finally, Dr. Mahoney, could you weigh in on this?

MAUREEN MAHONEY: Sure. So we'd encourage FTC to take a look, particularly, at hidden fees. That's long been a concern of Consumer Reports. They're really common in online ticketing and in hotel and travel sites. We think that if a fee is mandatory, it should be included in the upfront advertised price. We also think it's important because we've heard from consumers that it's a priority for them.

We help direct over 6,000 complaints about hidden fees in online ticketing to the Federal Trade Commission in 2018. They're really frustrated that, when they begin this process, they have no idea how much it costs until it's too late. And they're afraid to cancel the process for fear of missing out on an opportunity.

We're also concerned that these practices have persisted despite action from the FTC warning letters sent in 2012 and 2013 to a number of hotels and online ticketing, online travel agencies, for hidden resort fees. Consumer Reports took a look in 2019 to see what was happening, and most of these entities were still engaging in these practices. So we think that a rule banning hidden fees and drip pricing under the FTC Section 18 authority would be appropriate.

REID TEPFER: Thank you, Dr. Mahoney. For my final question-- and I apologize, I don't believe I'll get to everybody on this one. But could you-- could the FTC and other regulators' efforts to combat dark patterns be strengthened by the adoption of additional rules, guidances, or other policies? And if so, what would you recommend? Professor Willis, would you mind starting this one off?

LAUREN E. WILLIS: Sure. I think there's a lot the FTC can do with current authority. One thing you could do is educate the judiciary. In the *FTC versus DIRECTV* case, the court said, all told, the FTC's theory of the case requires the court to attempt to determine the net impression of more than 40,000 advertisements. And in dismissing the case, the court said-- or didn't end up dismissing, by anyways-- but basically causing the [INAUDIBLE], the court said that the variation among the DIRECTV ads precludes generalizing. And so there's some education to be done among the judiciary.

And the way in which the FTC applies Section 5 really needs to evolve to meet the challenge of the current unfair and deceptive digital design practices. So one evolution would be the use of presumptions, that once the FTC demonstrates a significant number of a business's customers hold roughly the same false belief about a transaction, the FTC should apply or ask courts to apply a presumption that the business that benefits from that false belief is responsible for it. Now, if customers think they committed to paying $29.99 a month for cable for two years, when they really committed to pay twice that amount in the second year, or they believe they opted out of a subscription when they opted in, the commission shouldn't be forced to prove which specific iterations, or that all 40,000 iterations, or even a representative sample-- because that gets to be quite a large number-- of the marketing produced the false belief. And instead, there should be a legal presumption that arises that the business that benefited engaged in deceptive practices.

Alternatively, the FTC could use Unfairness Doctrine. It's unfair for businesses to take advantage of customer confusion about really material facts, regardless of the source of that confusion. Unknown high fees in year 2, unwanted subscriptions-- know I think we've all spoken that that really collectively causes substantial injury. And consumers cannot avoid that injury because they don't know what they agreed to. They didn't intend to agree. And there are really no benefits to competition or consumers of sales based on confusion. And applying Unfairness Doctrine in these cases would obviate the need to identify which specific designs cause consumers to be mistaken.

In addition, I would suggest that the FTC be a bit more creative in crafting effective injunctive relief. So we often see that there are consent decrees that say, well, the firm from now on is going to clearly and conspicuously say this or that. But then we also see those same companies being brought back again and again for similar practices.

So one way to attack this is to require businesses that have engaged in unfair or deceptive digital practices to periodically demonstrate that, when their customers clicked Buy Now, they weren't confused. They understood the material facts about the transaction. And businesses would, of course, have to hire an independent third party to periodically test random samples of their own customers to demonstrate compliance with that, sort of along the lines of what the expert did in the LendingClub case.

The commission I think rightfully disagrees with the expert's conclusions for some other reasons. But the general approach of testing customers to see, did they know what they were agreeing to, is the right one. As for additional authority, one place where the commission could use additional authority would be in standardizing some aspects of online consumer transactions. So imagine if every business selling subscriptions posted on the first visible screen in the same location on the screen the identical digital Stop button for customers to automatically unsubscribe, or a standardized No button, or don't collect my information, or a standardized link to a standardized list of fees, a clear price tag.

So the commission could develop these with industry and standard-setting bodies, but would probably need some additional legislative authority to do that. The internet, it shouldn't be the Wild West anymore. There's just too much traffic. We need infrastructure. We need Stop signs and street signs to enable consumers to shop easily, accurately. This would benefit consumers. And of course, it would benefit fair competition.

REID TEPFER: Thank you, Professor. And Dr. Mahoney, could you please give us your thoughts on this?

MAUREEN MAHONEY: Sure. I'll try to be brief. Dean Willis already made some really good points. So clearly, the FTC should use the fullest extent of its authority to crack down on these dark patterns. It already has a number of things it can do--robust enforcement, updating the dotcom disclosures, recommending more guidelines and standardizations to businesses, and like I said, pursue a rule prohibiting hidden fees. But we are concerned that the FTC does need more authority, especially in light of the recent Supreme Court case that Professor Strahilevitz mentioned, *AMG Capital Management versus the FTC,* which struck a blow to its Section 13(b) authority to seek refunds for consumers.

So we're supporting recently introduced legislation that would restore that authority, better enable the FTC to ensure victims get timely relief. Further, in the context of privacy, we've long advocated for comprehensive privacy legislation to further enable the FTC's authority in this area. And we've also urged policymakers to go beyond the traditional opt in and opt out consent model and put strong limits on what companies can collect, use, and share in the first place through data minimization, to ensure that the onus isn't always on the consumer to opt in or opt out, which can provide an avenue for dark patterns. So I'd say that those are our priorities.

REID TEPFER: Thank you, Dr. Mahoney. Laura, I'd like to go to you next on this one.

LAURA BRETT: Sure. With regard to what the FTC can do, I'm going to reiterate what I said a little earlier. Specifically, I think targeted enforcement together with specific guidelines on dark patterns that they're seeing that are against the law and warning letters, widespread warning letters, could have a huge effect on the marketplace. We've seen it work before with emerging issues in advertising and certainly seeing it work pretty effectively right now with regard to COVID-19.

The FTC's actions on COVID-19 claims has been pretty remarkable. And it has encouraged a lot of compliance with self-regulatory guidance on those similar issues. So we think if the FTC undertakes those three actions, we could see more widespread adoption of better practices in this area.

REID TEPFER: Thank you, Laura. Professor Strahilevitz, do you have thoughts on this?

LIOR J. STRAHILEVITZ: I do, thanks. So I think there were a number of really great suggestions just offered. And I'll just agree with and elaborate on a few of them.

I do think federal legislation, along the lines of the legislation that was suggested and has been proposed by Senator Warner and Congresswoman Blunt Rochester, would be great, would be needed. I do think AMG Capital Management tells us that the federal courts are going to be real sticklers about the FTC's enforcement authority and whether it complies with the text of the relevant statutes. Deceptive and unfair has a lot of case law that developed what that means and some guidance from the commission about what that means.

But the Wyndham case, which I didn't think even should have been close in the Third m, actually turned out to be close. And I think the federal courts are worried about whether the industry is adequately on notice with respect to what it can and cannot do. And I think that's one reason why a specific statute would be especially helpful given that kind of judicial skepticism.

The truth is, industry follows what the FTC does closely. They follow the consent decrees. They certainly follow the cases. They follow the regulatory guidance. But the federal courts don't necessarily realize the extent to which all of this is closely tracked by industry. And so you see some worries about whether there's adequate notice built into the existing statutory and regulatory framework. And a clear statute would address that.

A clear statute would also give the FTC greater authority to go after first violations, as opposed to just having to enter into a consent decree after an initial Section 5 violation. And that would be really powerful deterrents for companies that are toying about employing dark patterns. And then just a couple of other thoughts, sort of in addition to what Professor Willis said about the consent decree process what the FTC does enter into a decree, I think another way to think about those consent decrees is to employ audits, where a company has a user interface that presents a choice between not yes and no, but yes and maybe later.

FTC imaginative consent decrees could authorize the agency to tweak that code and, on that particular platform, examine and analyze what the difference would be between switching yes and maybe later to yes and no as the option set. In other words, whenever an entity has struck a consent decree with the FTC, the FTC could be much more creative about using that as an opportunity both to deter misconduct and evaluate how powerful particular design choices may be in altering user behavior.

And then I think the final point a mention-- and this elaborates on an idea that Dr. Mahoney suggested-- is really think about providing guidance on what consent means. Even if you go to common law, there are doctrines, like undue influence, which is covered by Section 177 of the restatement second of consumer contracts, that give us a vocabulary already from the common law to talk about how, under certain circumstances, consent, that's checking a box or advancing to the next slide, might actually be voidable consent from a consumer's perspective if the consumer's actions are not fully their own, but rather are being dictated by the design choices made by the party on the other end of the transaction. And so when the courts are thinking about what constitutes consent, they want to see data, and they want to see hard data. But to the extent that us social scientists can help with that, it's something I think a lot of us are very fired up about doing in a way for us to help make the FTC's job easier and improve enforcement in this space.

REID TEPFER: Thank you, Professor. Unfortunately, that's all the time that we have this afternoon. But I want to thank the panel for joining us today for what's been a fascinating discussion. We greatly appreciate it. And thanks to everyone for joining us. I'm going to hand it over now to the acting director of the Bureau of Consumer Protection, Daniel Kaufman, for some closing remarks. Thank you.

DANIEL KAUFMAN: Good afternoon, everyone. And I am sorry for that delay. I want to close today's workshop first by thanking you all for joining us and for participating in today's timely, excellent discussion. After listening to today's panel. I am reminded of the saying, everything old is new again. For decades, direct mail marketers have relied on well-tried psychological tricks and design tactics, like prechecked boxes, hard to find and read disclosures, and confusing cancellation policies, to get consumers to part with their money. Today's discussion confirms what we've been seeing in our cases for some time. The old tricks and tactics have moved online and been repurposed for the digital marketplace.

But our panelists also highlighted important differences. In particular, new technologies like A/B testing and artificial intelligence now allow marketers to quickly and cheaply test, refine, and personalize user interfaces to maximize click-throughs, consents, and purchases. And increased tracking of people's online preferences and behavior means that advertisers can now target us when and where we are at our most vulnerable.

Our panelists also have pointed out that, like their offline predecessors, digital dark patterns can and do harm consumers. For example, some dark patterns cause people to purchase goods or services they didn't want or intend to buy and to give up personal information they would have preferred to keep private. They also run counter to our commonly held notions of individual autonomy and free choice. And they may disadvantage competitors who choose to operate fairly and transparently.

Moreover, our panelists have reminded us that these harms, like so many in our society, are not borne equally by all of us. Those in marginalized communities and people of color, in particular, have long been targets of unscrupulous companies looking to score a quick buck. And dark patterns gives companies a potent tool with which to prey on these communities.

Separately, our panelists have given us examples of how digital platforms are using dark patterns to gain and keep children's attention, raising concerns about unhealthy usage habits, and to sell them in-app merchandise, in some cases without their parents consent. Many of the dark patterns discussed today already are illegal under Section 5 of the FTC Act and state laws prohibiting deceptive and unfair practices, as well as Under the Restoring Online Shoppers' Confidence Act. And the FTC, along with its state and international partners, have been and will continue to be active in investigating and bringing suit to stop these unlawful practices.

But as our panelists have flagged, the scale, diversity, sophistication, and personalization of dark patterns also present regulatory and enforcement challenges. Our panelists have given us a lot to consider, and I would like to thank them for participating. And thank you to the groups and individuals who have filed comments. For those who have not yet had the opportunity to do so, the window to file comments will remain open until May 29.

So what is next? No crystal ball here, but one thing I can tell you to expect is continued aggressive FTC enforcement in this area. A number of panelists also have argued for additional rules, policy statements, or enforcement guidance. I can't make any promises, but I can assure you that we are carefully considering all options and nothing is off the table.

I'd like to close by thanking our organizers and moderators, Sam Jacobson, Katharine Roller, Miry Kim, Min Hee Kim, Evan Rose, Andi Arias, Rosario Mendez, Reid Tepfer, Sandy Brown, Jason Adler, Dan Salzberg, Kate Moody and Brad [INAUDIBLE] for paralegal assistance, Danielle [INAUDIBLE] for our workshop and logo work, Leslie Fair for our business blog posts, Julianna Gruenwald Henderson and Nicole [INAUDIBLE] from our Office of Public Affairs for our press engagement and live tweeting, Linda Hodge and Anita Thomas for coordinating public comments, Bruce Jennings and our web team for today's webcasting, and last but definitely not least, our event planners Crystal Peters and Arissa Henderson. We also thank you for attending today. Have a wonderful evening. Thank you.