FTC BCP’s Age Verification Workshop

January 28, 2026

Mark Eichorn:

Morning. Welcome to the Federal Trade Commission's Age Verification Workshop. I'm Mark Eichorn. I'm an assistant director in the privacy division in the Bureau of Consumer Protection. We're looking forward to learning today from an expert group of panelists. And before we dive in, I'd like to thank all the panelists for sharing their expertise today and also for their patience. We had weather concerns and had to move to virtual, so I appreciate your patience. We also want to thank everyone involved in planning today's event at the FTC. In particular, I'd like to thank Diana Chang, who woke up at four o'clock this morning to start her day, Jamie Hine, Gorana Neskovic, Peder Magee, Liz Averill, Jim Trilling, and Manmeet Dhindsa in DPIP, my colleagues, and Bruce Jennings and James Murray. Thank you all so much. For those viewing online, thank you so much for viewing the event today.

If you want to share on social media, you can use #AgeVerifyFTC. So, let's get started. It's an honor to introduce FTC Chairman Andrew Ferguson to kick us off. Chairman Ferguson.

Chairman Ferguson:

Thank you, Mark, and good morning. I want to begin by expressing my gratitude to the organizers of this workshop. Our BCP director, Chris Muffarige and his team who have had to brave the challenges of managing the weather and moving this from an in-person on-site event to a virtual event, which is a lot more complicated than it sounds. And I also want to express my gratitude to all of the invited speakers and guests. The FTC regularly hosts workshops such as this one to gain a better understanding of emerging challenges in antitrust and consumer protection. By learning from policymakers, advocates, experts, and industry leaders, these workshops help us to identify the proper scope and application of our traditional enforcement power to address those challenges.

This is, in fact, the second workshop we've held just this week, and we have others coming down the pike. Today's workshop focuses on the interplay between the FTC's enforcement of the Children's Online Privacy Protection Act, otherwise known as COPPA, which requires covered websites and online services to provide notice and obtain verifiable parental consent, a term you'll probably hear a lot today, before collecting personal information from children under the age of 13 and developments in age verification technology.

In essence, the purpose of this workshop is to determine the best practices and possible pitfalls for the use of age verification technology as it relates to COPPA. Before proceeding with my remarks, however, I want to acknowledge my friend and colleague, John Schweppe. John has been working at the forefront of these issues for a very long time. For the past year, John has served as my senior policy advisor, and prior to joining the FTC, he spent 10 years with the American Principles Project, where he became one of our nation's foremost experts on the protection of children online and in the doctor's office of social media censorship and on the sorts of policies necessary to reign in the abuses of big tech firms. John has been one of my indispensable advisors and an advisor, not only to me, but to many of the lawyers on my staff who work on these issues.

We scored a lot of victories for parents and consumers in the tech space this past year, and John deserves tremendous credit for the agency's success. Not only has John been an indispensable component of my staff, but I think everyone in my office and everyone else with whom he's interacted with the FTC would agree that John has been a truly delightful colleague and friend to everyone here. On behalf of myself, my team and the entire FTC, I want to say thank you, John. You're a great guy and a patriot, and you've helped make America a better place during your time at the commission.

Today's workshop comes on the heels of a series of significant COPPA enforcement actions brought by the FTC over the past year. Iconic Hearts is involving an anonymous messaging app, which is currently in litigation, Apitor, a settlement of internet connected toymaker. And Disney in which the FTC alleged that Disney uploaded child-directed videos to YouTube without labeling those videos as made for kids. As Congress considers whether to adopt additional legislation to protect children online, which it's been doing for some time, the FTC must use every tool at our disposal, chief among them, COPPA and the COPPA rule, to empower parents who are the first and best line of defense to protect children online. COPPA enforcement is and will remain a top priority of the Trump-Vance FTC, and we will push COPPA as far as we lawfully can to protect America's kids. Our order requiring Disney to obey COPPA was particularly significant because it expressly acknowledged the role of age verification technologies as an emerging and increasingly important means of protecting children online.

Our complaint alleged that Disney did not label child-directed videos as made for kids, and YouTube in turn did not prevent the collection and monetization of personal information, without parental consent from children under the age of 13 who interacted with Disney's videos, as COPPA requires. Because of Disney's misdesignation of children's videos, children may have been exposed to age inappropriate YouTube features, such as the autoplay of videos that were designated as not made for kids. The FTC levied a $10 million fine against Disney for its COPPA violations and required it to implement an internal program that ensures a systematic review of each of its videos published to YouTube to determine whether it is child-directed and must be designated as made for kids. A systemic review of every single video uploaded to YouTube could impose significant burdens and costs on any operator. Our order therefore authorized Disney, which uploads an incredible amount of content to YouTube, to phase out the systemic review if YouTube implements and Disney uses age verification technology that can ensure COPPA compliance.

Higher costs are no excuse for breaking the law or for relaxing standards for complying with the law, and the FTC's order permits neither. It instead encourages technological innovation in COPPA compliance, which in turn expands the protection of children by reducing the cost of complying with COPPA. Or of voluntarily implementing other measures to protect children. And that's the broader goal of today's workshop, to discuss and work towards answering the question, how can government agencies and regulators, whether on the federal or state level, facilitate the development and adoption of emerging technologies that expand the protection of American children by reducing the costs of protecting them. Or stated a little more generally, how can regulators promote technological innovation that in turn promotes the common good, in this case, keeping our children safe online? Today's workshop will focus on age verification technologies, but that's a particular application of a general principle, aligning technological advancements with the common good.

I want to emphasize this broader goal to illustrate that there need not be tension between the FTC's mission to protect children and technological innovation. Every company operating in our nation ought to respect the demands of lawmakers and the public they represent. The task of innovators then is not to find innovative ways of breaking the law, as we sadly see so often in our consumer protection cases, but to develop and adopt new technologies or business practices that make compliance with the law and the company's service to consumers easier and more cost-effective. So too, lawmakers and regulators ought to encourage and incentivize this kind of technological innovation for the same reason, namely to advance the common good by making compliance with the law easier and more cost-effective. But we are not interested in technological innovation for its own sake any more than we are interested in compliance with the law for its own sake.

Rather, we are interested in laws and technological innovations that are directed to the flourishing of every single one of our fellow citizens. Within the context of COPPA, lawmakers, regulators, and businesses should be invested in technological innovation that makes it easier for businesses to protect the privacy of children online because we believe that the flourishing of our nation's children depends on the privacy of their personal data and on the capacity of parents to control who has access to their child's data and how those data are used. As I said in a prior workshop, COPPA and other laws governing online privacy for children ought to aim at assisting parents in exercising their right to exert meaningful control over their child's activities online and the data generated by those activities. When Congress passed COPPA more than a quarter-century ago, the internet was still in its incipiency. As the range of online services and users has expanded, we've seen significant advances in the protection of online privacy as well as in age verification technologies.

Today, it is much easier for businesses accurately to identify and fully protect the privacy of their child users than it was when COPPA was passed over 25 years ago. And that means it is much easier for everyone, lawmakers, regulators, and businesses to advance COPPA's noble intention of empowering parents, not only to protect the privacy of their children online, but also to oversee and guide their child's online activity. Empowering parents in this way is not good because of COPPA. Rather, COPPA is a good law because it is good to empower parents to shape and control their child's online habits and activities, or at the very least, it is better to empower parents to shape their child's online habits and activities than it is to empower big tech executives, obscure algorithms, AI chatbots, pornographers, or online predators to shape a child's online habits and activities. Indeed, the internet we encounter today does not look like one even modestly influenced by the choices of parents with small children.

It looks a hell of a lot more like Las Vegas than Little House on the Prairie. Absent an effective means of verifying an online user's age, parents must zealously and closely police their child's online activities to ensure that he or she doesn't fall headlong into age, inappropriate entertainment, gambling, pornography, or other forms of exploitation. As citizens and lawmakers, we can't eliminate all the dangers and depravity of the internet, but we can make it easier for parents to protect their children from it. By adopting robust age verification technologies, internet companies can demonstrate by deeds, not words, their own commitment to our nation's laws, and more importantly, to our nation's parents and the protection of their children. Now, why would anyone anywhere oppose or otherwise flat laws designed to empower parents to control their child's activities and data online? Because in individual data, whether a minor or an adult, is a precious commodity for advertising purposes, among other things, it's a source of profit and gain.

And because an individual's data is a source of profit, many online operators will do whatever they can to remove barriers of access to an individual's data, even if this means circumventing or ignoring altogether federal and state laws that require meaningful checks on an individual's age and verification of parental consent. Accordingly, many online operators are resistant to adopting age verification technology, not only because it prevents them from accessing children's data, but even more so because it might prevent them from accessing the data of adults. In the recent Supreme Court case of free speech coalition against Paxton, pornography industry lobbyists made this concern quite explicit. They did not dare to object to the principles that minor children should not have access to pornography. Instead, they argued that online age verification requirements would prove too chilling for adult consumers of online pornography, infringing on the exercise of their putative First Amendment rights to access pornographic content.

Because adult consumers of online pornography would fear their anonymity could be compromised by age verification technology, so the argument goes, they would be less willing to exercise their putative First Amendment right to consume pornographic content. In other words, pornographers fear that age verification technologies would reduce their customer base, but no industry, not even the pornography industry, is immune from laws that might reduce their customer base and profit margins. When deciding between pornographer profits and protecting children online, our nation's voters, by an overwhelming and consistent margin, have chosen to protect children. That's why the Trump administration has made it clear that the health and flourishing of our nation's children is non-negotiable, and that's why the Trump Vance FTC will not hesitate to use the full extent of its enforcement powers to protect children in the online space, even if this imposes financial burdens on companies or otherwise hinders certain forms of technological innovation. But we aren't here today to foment a conflict between technological innovation and common good, quite the opposite. We are here today to discuss the convergence between the two to identify how recent innovations in age verification technology can make the internet safer for kids and how providers and online gatekeepers can use age verification technology to protect children and their privacy. We're here to explore what steps the FTC can take to ensure that the COPPA rule does not unduly inhibit the implementation and innovation of effective age verification technology.

With that aim in mind, we've invited policymakers, academics, regulators, advocates, industry leaders, and product developers to participate in this workshop across four sessions. In the first session, we will learn about the current political, legal and regulatory framework governing age verification, both here and abroad. In the second session, we will discuss how the technology works, how it might affect users of online services, and how it can be tailored to fit the demands of policymakers or the needs of companies. In the third session, we will consider age verification as a tool of regulatory or legal compliance, including its implications for the privacy and free speech rights of adults online. And finally, in our fourth session, we will discuss with representatives of major industry leaders about the past and future implementation of age verification technology, as well as actions already being taken to assist parents in exercising meaningful control over their children's online activities.

To everyone presenting, I want to express my sincerest gratitude for your participation in this workshop. And I can assure you that the perspective you share with us today will inform the commission's work in this arena. Let me close with the brief word about what we hope will come out of this workshop. As the primary enforcer of COPPA, our discussions today will provide the commission with better insight into the interplay between age verification technologies and COPPA. To that end, I expect the fruits of this workshop will inform a future FTC policy statement on age verification technology, as well as a possible amendment of our own COPPA rule that would promote the use of age verification technologies in compliance with COPPA. In doing so, we hope to incentivize a wider adoption of age verification technology that would enable operators of online services to know whether children are visiting their website, and if they are, to ensure that the operators take the necessary steps not only to comply with COPPA, but also to impose safeguards for children.

In a time of rapid technological innovation, policymakers, regulators, and business leaders should keep in mind our common goal to promote the flourishing and its success of ordinary citizens and their families. We should not assume that technological innovations is at odds with this noble lane by providing a forum for constructive engagement between policymakers, regulators, advocates, and innovators on the possibilities and potential pitfalls of age verification technology, I hope we can forge a path that weds the promise of this technology with the purpose of serving our citizens and their families. Thank you so much, and I look forward to all the discussions today.

Peder Magee:

Thank you, Chairman Ferguson, for those opening remarks. Good morning, everyone. I'm Peder Magee, and welcome to our first panel of the day, Understanding the Landscape: Why Age Verification Matters. This discussion is intended to be something of a level set and to lay a foundation for the conversations that will follow in the later panels. We're going to start with short presentations from each of the panelists and then go to a discussion with Q&As. A quick note, in the interest of time, I'm going to ask each panelist to introduce him or herself and their organization, and then dive right into the presentation. The audience can find bios for the panelists on the workshop page on the FTC's website. Great. So now I will turn it over to Mark Smith. Take it away, Mark.

Mark Smith:

Hey, thank you very much, Peder. Delighted to be here, although I wish I were there in person and I think a lot of other folks wish they were there in person, but Washington and Snow don't mix too well. But that said, my name's Mark Smith. I'm with the Centre for Information Policy Leadership, or as we like to call it CIPL, C-I-P-L, which is an international privacy and data policy think tank. We are unique in that we reside within the law firm of Hunton Andrews Kurth, but our mission is basically to facilitate the building bridges between industry and regulators. And a core focus of our work has to do with organizational accountability. And promote organizations to use best practices in generating practices that show that they are good data stewards of the data they're collecting.

And Peder asked me to help out with a level set of what we're talking about today. And if you could switch to, I guess, the next slide, please. The name of this panel does make reference to age verification, but I just wanted to highlight that age verification is just one of the terms that fall under the umbrella known as age assurance. And I've listed the general categories of age assurance here. Self-declaration is a form of age assurance which simply asks a user to enter a birthday or click a box affirming that the user is over a certain age. Most people agree that self-declaration alone is considered inadequate, especially in higher risk services and higher risk situations. So nobody's here to promote self-declaration as folks lie, and it's very easy for kids to bypass any sort of self-declaration form.

I know that we have providers here who can give greater detail on age estimation or age inference methodologies, but generally speaking, these methods use machine learning to infer or estimate a user's age based, for example, on selfies or on a phone number, or an analysis of a user's online history. Age verification, which is the term you hear most often, usually encompasses a more rigorous method to determine a user's age, such as supplying a website with a scan of a driver's license or passport, or using third-party databases or other government IDs to verify a user's age. For users who are minors, it may also encompass parental consent. All right, next slide, please.

In 2024, CIPL prepared a discussion draft of age assurance and age verification laws in the US. We took a look at the types and scopes of state laws, and the legal challenges arising from those laws. And we offered some recommendations for policymakers moving forward. Although ordinarily, I would be happy to supply a QR code here so that you can download this report. All of our papers are available on our website, but our website's in the process of getting a much needed overhaul so that any code I'd give you today would not be available when the new site goes live next month. So to eliminate the frustration, I'm going to forego the QR codes for now, but please reach out if you'd like a copy of this or any of our reports. Next slide, please. This map is something that we produced in that first report where we classified and grouped age assurance laws within one of three categories.

First were those that seek to prevent minors from accessing pornographic content or other content deemed harmful to them. And that's shown in pink on this map, and we use the designation of laws directed to content that's harmful to minors. The second category was laws seeking to prevent minors from creating or maintaining social media accounts without parental consent. Those are shown in blue. And the third category is those seeking to afford greater privacy and safety protections more broadly when minors are either likely to access or known to access a given site or service, and those are shown in yellow. This sort of law covers the age-appropriate design code laws and other broader laws like New York's Child Data Protection Act. The states that you see with striped entries have laws addressing more than one of the above categories. Please note that this map is not up-to-date.

It was prepared in 2024, so that's why some states like Vermont are still in gray. But I wanted to note that most US laws, especially those addressing the harmful to minors content, use the term age verification, or sometimes reasonable age verification, to describe the process of ensuring that individuals seeking access to that content or at least 18 years of age. So to clarify, businesses covered by these laws need not verify an individual's specific age, but simply whether an individual falls above or below a certain threshold, usually 18. That said, the age verification requirement applies to anyone seeking access, which of course means that it applies to both minors and adults. Others of the laws that we surveyed use other terms like age determination, age assurance, and age estimation. I wanted to highlight, in particular, Maryland's Age Appropriate Design Code Act does not include an age estimation provision, and it actually prohibits any processing of children's data for purposes of age estimation.

While most state laws with age assurance requirements generally adopt 18 as the threshold, there are a few exceptions. Georgia's law, for example, adopts a 16-year-old threshold for the portion of its law that addresses social media obligations. And 18 for the portion that covers pornographic content. Louisiana's social media law defines a minor as an individual who's under the age of 16 and not emancipated or married, which raises a question about how businesses are to assess marital status or legal emancipation. So this raises questions of collection of information. Businesses certainly do not want to collect more information than they need to, and I doubt any business wants to collect information on the marital status of a potential user. Next slide, please.

The US state laws set forth an array of elements that can factor into a reasonable or at least a statutorily acceptable age assurance methodology. As displayed in this graph, verification based on a government-issued ID is the most common element with verification based on a digital ID and transactional data filling out the top three. While some older teens may possess a government-issued driver's license or learner's permit, they would rarely have a history of transactional data. Also, I should note that most laws prohibit the retention of personal information after an individual's age has been verified, but other types of processing such as sharing the data with third parties are not addressed. Next slide, please.

Addressing the issue of age assurance on a global scale, CIPL has partnered with the We Protect Global Alliance to initiate a multi-stakeholder dialogue on age assurance. We first met in 2024, and this dialogue is still ongoing. In fact, many folks participating at today's event have taken part in these dialogues, including Michael Murray, my co-panelist here for this session, and anyone interested in taking part in future dialogue should certainly reach out to me. I want to highlight that this has truly been a multi-stakeholder conversation. Attendees have represented a diverse range of organizations, including child rights, privacy, safety, academia, regulators, civil society, and industry representatives from technology, entertainment, telecom, and financial sectors. The purposes of this dialogue are laid out on the slide here to promote a global dialogue, which we have because we have engaged with folks across the globe. We actually had one of our discussions here in DC, the October before last, I believe, to bring together experts from various sectors, including the ones that I just mentioned, and to advance a holistic and principle-based approach.

The conversations that we've had have been under the Chatham House rule, but we published takeaways from each meeting again, and all of those takeaways are available on our website. Next slide, please. Several working groups have blossomed from these meetings, law and regulation, one addressing risk assessments, and another on regional and global perspectives, and those reports are also available on our website. Next slide, please. I know this has a lot of information, this and the next three slides, but I wanted to highlight some of our high level takeaways from our discussion so far. I obviously can't discuss all of these, but let me focus first on the purpose of age assurance that it should be viewed as a process and not a singular one-off check. Also, it should not be viewed as merely a means of excluding children from inappropriate content. Age assurance can be used to provide tailored age-appropriate online experiences, and it can help businesses know when certain obligations kick in, for example, the COPPA obligation.

Second, age assurance measures should be risk-based and proportionate to the level of risk. Context matters, and we highlight that in our work. Third, as we often hear now, it's not just about privacy anymore, it's about safety too. And those two principles need to be balanced in any proposed solution. Next slide, please.

Next, as you can tell from the overview that I discussed earlier on the US landscape, there's a great deal of regulatory fragmentation with different requirements and different thresholds. So we need to focus on some baseline standards and consistent approaches to age assurance. The fifth point, age assurance is a technical solution, so it needs to be interoperable across websites and platforms and devices, ideally. Sixth, we can't forget about the end user. Age assurance should be user-friendly and accessible to children's, teens and parents, which I believe the chairman was highlighting in his opening comments. Next slide, please.

One of the key benefits of a multi-stakeholder dialogue is a greater understanding of the many players in the online ecosystem and their unique roles and expectations. So there needs to be collaboration among the different players in this space and most definitely clarity on the allocation of liability. On the matter of ethics and the right thing to do, an appropriate age assurance solution needs to address competing interests and unintended consequences, such as the free speech concerns raised in many legal challenges. Lastly, among the key challenges we identified, there is currently limited guidance on age-specific harm assessments leading to inconsistent evaluations across platforms. Since we are promoting a risk-based approach, we need to have a common understanding of what constitutes risk. And next slide, please. So where are we now? We concluded that the multi-stakeholder dialogue concluded 2025 with a draft framework that builds on the takeaways from the prior three slides.

So again, we hope to encourage the development of age-assurance solutions that build on fundamental principles, risk-based, proportionate context matters, incorporating privacy by design and allowing user autonomy and transparency. Also, it should be technologically neutral, like supporting various digital credentials. That could be a key benefit for users and the ability to reuse those credentials across different services and platforms. And also, as I mentioned earlier, to clearly define the roles and the liabilities of the players in this space, which includes age assurance providers, API providers, and the actual content providers, the app and the website providers. So, with that, Peder, I hope I didn't take up too much time, but wanted just to give an overview of our work and the conversations we've had today.

Peder Magee:

Great. Thank you so much, Mark. And now we'll turn it over to Amelia Vance.

Amelia Vance:

Wonderful. Thank you so much for having me. I'm thrilled to be here. I'm Amelia Vance, founder and president of the Public Interest Privacy Center, a nonprofit working exclusively on student and child privacy. I am a FERPA and COPPA geek, and I also teach privacy and EU data protection at William & Mary Law School. Very, very happy to be here. Thank you for inviting me. So, in many ways, I am giving the other half of Mark's presentation and focusing in a little more on the laws that don't explicitly say that there should be age verification or assurance, but say generally that you need to protect children. You need to take certain steps to protect children. And therefore, you need to have a way to show that they are children. So, in many ways, some sort of age verification, age assurance, whenever that guarantee is necessary, is going to end up being part of the solution or the answer to a law protecting kids.

So, the map on this slide is the combined everything as of six days ago, which means South Carolina has already passed something. These are all of the different laws that have shown up. So not bills, laws. As you can see, the landscape

Amelia Vance:

It is very complicated and messy right now. There's a combination of the different approaches. You have several states trying to work on data governance for kids in addition to privacy protections as well as levels of age verification for different purposes. But this is the landscape that everyone is looking at today in this legislative session. And as many of you may have seen, state legislatures have been far from slow to add to this list. So it will only get more and more important. And I think that is important for a reason I will come back to, but let's go to the next slide.

Taking a step back, why does knowing age matter? Might be legally required, but why does it actually matter? And it's because in many ways, it's a quote from a group who's been working on these issues for years, 5Rights Foundation, highly recommend their resources, very readable, "The digital world isn't optional for most kids. It's where they access education, health services, entertainment, build relationships, and engage in civic and social activities."

Kids are walking past companies that may be scanning their faces as part of day-to-day activities. They may be getting on a bus where they need to scan a fingerprint. All of the different ways that we as adults interact not only when we log onto a computer or use our device, all of those things carry over to kids. And so making sure the digital world is safe, also empowering, helping them to grow up to make smart decisions is absolutely essential. Much broader than any legal obligation.

Next slide.

So specifically, I mentioned I'm doing the other half of the landscape explicitly. Here are the laws that have implicit age assurance requirements. So as noted, a lot of states have introduced laws that regulate social media for minors. In particular, some of those do have strict age verification requirements or some level, I would say beyond just age assurance, whether that be scanning faces to estimate age, something more than just claiming you are a particular age.

And then the other is actually many ways more akin to growths from COPPA. It's more focused on data governance, empowering parents, making it so no matter what you are doing online, the internet is safe. And I should be clear, safe doesn't mean kid-proof. I think there's a lot of people who have been very concerned for years about the idea of kid-proofing the internet. I found out a few years ago that kid-proofing the internet was considered very concerning. You have these laws that are going to impact way more than just companies directed to kids. But one-third of internet users are under the age of 18. So when we talk about kid-proofing the internet, we're talking about one-third of all users, and so making sure you have these underlying protections is essential.

So, going to the next slide. So I am primarily here to add a little bit about the law, the case that the chairman mentioned in opening statements, giving caveat as a lawyer, not giving legal advice, et cetera, but also because I need to speak so quickly that I will inevitably miss some of the nuances. So feel free to reach out if you have additional questions.

What was settled in this? This case was focused on whether you could have age verification for people to access obscene content. So, very specific to that particular issue. And the Supreme Court said really in many ways, not entirely overturning precedent, but certainly dramatically changing their answer from the past 20 years to say, "We think age verification technology can work now, and therefore we are comfortable with states passing laws to require proof of age to access content that is obscene."

So they move down their level of scrutiny, so intermediate scrutiny versus strict scrutiny. Often if something is looked at under strict scrutiny, it's much more likely to fail. But they also noted that age verification requirements to access content that is obscene to minors triggers intermediate scrutiny generally, but the burden on adults. So bringing up that question from earlier, the burden on adults to prove they are adults, including verification methods that involve having to provide a government ID should, when it comes to obscene content, be considered incidental.

But I just went through a number of laws that aren't focused on obscene content. They focus on access to social media, which is often defined broadly as an interactive service. If you can chat with another user, then these laws would apply there. You would need to go through some level of proof that someone is an adult or is a minor, or have social media studies and any site where you have interactions possible restrictive without age verification.

And we don't know if those are constitutional yet. That was the question that was opened by this case. The court noted that, really in dicta, that social media falls outside of this statute, so they weren't going to get to it. The court assumed that social media companies have less than a third of their content as obscene to minors, and it was reasonable for Texas to not extend the age requirement to those companies. It noted that a burden on obscenity to minors in this case may not trigger strict scrutiny. Again, they went down to intermediate scrutiny, a little more of a balance between government's ability to protect children and burden on everyone else, on adults, et cetera.

And they noted that in one of those cases from 25 years ago, a burden on obscenity to minors may not trigger strict scrutiny even if a comparable burden on indecent speech would. And so an open question is, is indecent speech what we're regulating with some of these state child privacy and online safety rules? And if so, we have a constitutional issue. But it's a brand new ballgame with this, and we're going to see it play out for the next few years.

So going to the next slide, I think some observers were mildly concerned, but there are a number of... There's a lot of litigation, as you will see in a moment. But shortly after the Supreme Court put out this case, you had a denial of allowing for an injunction in one of those cases for Mississippi's law. And you had a concurrence from Justice Kavanaugh here saying, "Yes, NetChoice did not make its case that there should be an injunction, but they have, in my view, demonstrated that it is likely to succeed on the merits, namely that enforcement of the Mississippi law would likely violate its members' First Amendment rights under this court's precedence." And you'll note some of the caveats there, like, "Under this court's precedence." So, it's not like it's definitely unconstitutional, but again, this is an open and evolving question. This is from last August. We haven't had a lot of time to see what comes next.

Next slide, please.

I mentioned there's a lot of legal challenges. Here's a quick image of it. You have a massive number of cases. There are 13 laws enjoined, 11 in effect, but being challenged in eight circuits. And this is probably going to expand after some of the new laws passed last year. And we have zero final rulings. Overwhelmingly, judges have said that these laws are likely unconstitutional. We'll see.

Going to the next slide.

Mentioning a couple of the key issues brought up here, so there's that concern about whether it interferes with adults' ability to access content online, noting not just obscene content, but these laws would cover much broader swath of content where you need age verification. And then a lot of the other issues that have been raised are around concerns about privacy. Are you going to have this ID information being collected end up being an issue?

So, for example, even though it was prohibited in Australia, smaller companies retained age verification data for auditing. They were worried that they would be audited regarding the law despite being required to delete it. So, anything we see on age verification, any law or policy needs to make sure that these are the kind of considerations where we can learn from our peers, particularly in other countries that have already done this, to avoid some of these potential privacy issues.

So, I'll go to my final couple slides here. So, the next slide.

A vital part of this question of this whole debate is, what happens once we know? Because age verification identifies the child, and the next question is, then what? What are the obligations? How do we empower parents when we know... The FTC workshop last June discussed extensively that consent isn't enough. Parents need additional protections that layer in here. Also, inevitably, kids will find ways online. The question is whether they'll have protections when they get there. And of course, when honesty leads to better treatment instead of exclusion from a space, when kids are given a better experience on a website where they aren't automatically connected with strangers in chat rooms or in video games, well, kids are less likely to look for workarounds. Age can be used to empower kids and not just to take away access.

So, last slide here. Next slide.

So, some considerations that have been discussed with the litigation, with all of the bills moving, it's important to make sure that there are protections before the gate. Sometimes you don't need age verification and the potential privacy issues that can be raised despite legal prohibitions on keeping the information because many risks disappear when platforms build with safe defaults anyways. And this isn't kid-proofing the internet. This is limiting tracking, not automatically allowing people to talk with strangers, particularly when they don't want to. It means allowing people to have by default public profiles. Privacy and safety by design reduces what is otherwise riding on age assurance alone. And the FTC has always done a fantastic job both in COPPA enforcement and Section 5 enforcement at pointing out those things that undermine this, where it is fundamentally deceptive or unfair to not set things up with these assurances.

Once you get to the gate, making sure it's verification that doesn't create new risks, collecting only what's needed, deleting it quickly. Don't repurpose it. It shouldn't become another vulnerability for kids. In particular, for any methods that involve a company using data, it already has to estimate the name of a user. Let's make sure that the deletion requirements also apply there. After the gate, companies have to be able to do something different. And throughout, there needs to be enforcement.

And that is all. Thank you.

Peder Magee:

Thank you so much, Amelia.

Now, we're going to go to Michael Murray, and if we can shoot for around 10 minutes, 12 minutes so we can have time for some questions at the end. Thank you. Michael?

Michael:

Thanks, Peder. Glad to be here. Michael Murray, I'm head of regulatory policy at the ICO, and I lead on children's privacy and age assurance for the ICO, which is the UK's data protection regulator.

I'd like to first congratulate the FTC in a quarter-century of leadership in piloting children's privacy through COPPA. You set an initial standard that we have built on through our data protection legislation and safety laws in the UK, so well done on that work. Peder and colleagues asked me to deal with what happens for the services that have more of an international scope. Most of the large US companies will operate not just for American users, but operate internationally. And the data protection regimes internationally are quite complex, so I'll start by looking at the UK context and then finish off with a bit of an overview of what's happening internationally.

I want to start by this idea that protecting children isn't just a legal requirement. It's the right thing to do. And it's picking up on what Amelia has already set out. The ICO sets out through our standards of our Children's Code, which incorporates age-appropriate design code, a guide to industry of how they can meet the requirements of the UK GDPR in practice when seeking to protect children within the online world and not exclude them from it. This principle, protecting children where they are, recognizes the benefits that children tell us they get from being online, including learning, staying connected with friends and families, and engaging in online playing entertainment.

However, the online world is not designed with them in mind. Amelia already gave us an indication of how many kids are online, a significant percentage. We know that children are using services that are not designed for them and are indeed not suitable for them to use. Despite having terms of service of 13-plus for many social media services, games, shopping services, example, in the US, UK, or elsewhere, services do allow children to register, with over half of eight to 12 year olds in the UK holding at least one or more social media accounts, many of them with an age profile of an adult, and therefore they don't have any of the benefits of the higher protections and default settings that are designed into the age-appropriate design code. Everything that's there to protect them as a child is not protecting them because they've identified and the service identifies them as being 18-plus.

Our research shows that children who generally do not have access to credit cards or regular incomes consider data to be their only tradable commodity to gain access to services that they want to use. And too many services default to self-declaration to set up an account, which Mark pointed out in his presentation. Children will lie about their age, we all know this, and our research shows that half of parents are complicit with these lies to appease their kids.

So let me be clear, self-declaration is not age assurance, is not an effective mechanism to identify children and protect them online. The results of this self-declaration-focused current paradigm is the harm we see and hear about, two-thirds of teens encountering harmful content online. This includes content that is illegal, such as CSAM, child sexual abuse materials, pornography, or should never be shown to children, such as self-harm, suicide ideation, or pro-anorexia content.

Although the ICO is not a content regulator and the AADC is not about content per se, we leave that content regulation to our colleagues at Ofcom. We are concerned with how data is used to deliver content and contacts that are harmful to children, and we will consider the harms of content and contact when deciding whether to take action on data models that are causing those harms. Beyond the content and contact harms, we remain concerned about the core data protection-related harms, the loss of control of personal data, the psychological and physical harms that occur when children's data is inappropriately shared, is hacked, is accumulated to develop profiles that feed advertising and disadvantage their development.

Next slide, please. Okay. The UK regime, protecting children online is a priority under the UK regulatory law. Ofcom, the communications regulator, leads on the safety laws. And the ICO, the digital information regulator, leads on protecting children's personal information and data. Age assurance is core to both the ICO and Ofcom's children's regulatory duties. I'm going to focus more on the ICO at this point. If colleagues would like to learn more about Ofcom, I'd encourage them to look at Ofcom's website for more details. I'll talk in generalities about Ofcom, but I'm not going to go into detail, nor do I speak for them.

The UK's data protection regime is set out in the Data Protection Act and the UK GDPR. This regime requires services to take a risk-based approach when they use people's data based on key principles, rights, and obligations. The ICO's Children's Code is a statute code of practice under the Data Protection Act that applies to online services that are likely to be accessed by children. To support, organizations understand their responsibilities in terms of Standard 3 of the code, Age-Appropriate Application. The ICO has published opinions on age assurance in 2021 and 2024 that are available on our website if colleagues would like to take a look at those for some guidance. To provide regulatory transparency about what we are focusing on, we also published a Children's Code Strategy in 2024 and regularly update on progress every six months or so.

Ofcom's job is to make online services safe for the people who use them. Robust age checks are a cornerstone of the Online Safety Act to prevent children from encountering pornography or primary priority content and protect them from other harmful content. Requirements to have highly effective age assurance to prevent children accessing pornography came into force in July 2025 with a significant impact to date. And I want to congratulate my Ofcom colleagues on the work they've done. Age assurance is one of a number of collaborative themes where Ofcom and ICO work together bilaterally through the Digital Regulators Cooperation Forum.

Next slide, please.

So, the Children's Code applies to online services that are likely to be accessed by children, meaning anyone under the age of 18. It sets out 15 key principles of age-appropriate design that helps services comply with the UK GDPR. The likely to be accessed policy underlining the code and most European data protection requirements and online safety legislation is a significant difference from the COPPA approach of actual knowledge and is there to ensure that protections apply where children actually are and not where services deem them to be.

Online services likely to be accessed by children need to assess whether a significant number of children are, in fact, using the services. Significant here is a legal term rather than colloquial understanding, meaning the number is not insignificant. So, a service with a low number of users but a high actual number of users is considered to be... Or percentage of users, sorry. A high number of actual users is considered to be in scope. If a service provider concludes that the service is likely to be accessed, they need to complete a Data Protection Impact Assessment to spell out the risks to the data processing and how they intend to mitigate those risks to protect children, such as the use of age assurance.

Next slide, please.

This slide sets out the ICO's current priorities within the Children's Code Strategy. Services likely to be accessed by children need to know if children are using the service to ensure adequate protections are applied to them and to prevent harms from content that services have all a legal obligation to prevent. Children's data must be processed lawfully with the options for a lawful basis set out in Article 6 of the

UK GDPR. Our evidence suggests that many services use a combination of consent, contact, contract, legal obligation, legitimate interests for their lawful basis, all of which come with their own requirements. The strategies focus on the processing of data for under 13s. Our colleagues at Ofcom are focused more on the 18-plus. We have looked initially at services with no age assurance, as has resulted in the UK in regulatory action against Imgur. We have also worked with social media services to determine the efficacy of profiling as a form of age assurance, and I'll come back to that in a bit.

We are currently working up a project to engage with services to support a transition away from a reliance on self-declaration only as a primary form of age assurance at account creation, where services undertake high-risk processing. And high-risk processing here would be, for example, the use of innovative technologies like AI, profiling, or serving ads to children.

Next slide, please.

What are we looking at for companies to do instead of self-dec? On the left-hand side of the slide shows what Ofcom considers to be highly effective age assurance. This slide outlines the types of technology that Ofcom considers to be highly effective and those that are deemed not capable of being highly effective.

Note here that self-declaration is not considered highly effective. The ICO takes a technology-neutral approach but does require age assurance to be effective. We have clearly set out that self-declaration alone is not appropriate for our services. This slide takes you through one example of what a compliant age assurance process might look like at the 13-age gate. Self-declaration can be a useful point, but it should be backed up.

So, for example, the first step might be the child identifying what age they are, going through an age estimation, facial age estimation to verify that age, and then progressing onto the service if the age estimation proves that they're over the age. If there is doubt, then pressing down to potentially a second age estimation, age verification, or so on. At the end of this process, data should be deleted.

Next slide, please. The UK GDPR does not have explicit requirement for age assurance, but age assurance is implicit throughout. And services who are looking to be accessible for children need to be compliant with these principles. We are not regulating against the AADC per se. Our legal action, if this is undertaken, will be against the underlying principles of the UK GDPR. Next slide.

In Europe, the GDPR applies similarly to the UK GDPR in the UK. The Digital Service Act and Audio Media Visual Services Directive cover video service providers. Also impose requirements on age assurance and protection of children.

In 2024, the ICO developed a principles-based approach for age assurance through an international age assurance working group signed on by 10 jurisdictions globally, including Canada and North America. European Data Protection Authorities contributed to the design of those principles, but published their own similar guidance in 2025. Across these two documents set out a series of principles-based approach to age assurance that applied to about 40 countries and territories. The principles set out on this slide align to GDPR requirements.

Of note here is a need for services to recognize and prevent data protection risks associated with the age assurance process. So key principles include data minimization and purpose limitation of data collected in the age assurance process. Data collected for age assurance processes should be used only for age assurance, and deleted once the need for that data has been met.

Automated decision-making should not normally concern a child, meaning that any process that relies on AI-based determinations alone are not appropriate. They must allow for human engagement to address errors that would have a significant impact on a child, for example, a false denial of service. Both UK and European regulators expect data protection by design and default, and that services use state-of-the-art in age assurance technologies to protect children. The whole age assurance system, including the complaints processing, must be secure and services must be accountable for the decision. So, we're not looking just at the initial age gate, but the whole process of age assurance from the start to the finish of the decision-making process.

Next slide, please. And this is our last slide.

Let me close with an overview of children's protection initiatives globally. It started with COPPA and the UN Committee on the Rights of the Child and the GDPR that set out legal and legislative inputs to protect children's rights and keep them safe online. Over the last five to six years, we've seen a steady progression of international laws, starting with the UK's AADC, the Audio Visual Media Services Directive in Europe, and the VSP regime in the UK under Ofcom. The Irish DPC followed in 2021 with its own requirements for children's protection that includes guidance on age assurance, followed then by the California AADC.

The ICO, to help promote regulatory certainty internationally, formed an international age assurance working group in 2023. We have about 40 members from across the world that share information about regulation and supervision activity. 2024 saw the DSA come into effect, and then from '25, '26, a whole series of new laws, including the Australian Age Assurance Trial and social media ban that has seen 4.7 million children removed from social media access if they're under 16. There are regulations in Brazil, in Indonesia, Singapore, for example. And 2026 Australia's Children Online Privacy Code will be coming into effect, and ISO standards also coming online in 2026.

Coming up in the future, the EU and UK also looking at social media bans. There is app store and device-based legislation being looked at in California and globally, COPPA 2.0, which we'll hopefully hear about a little bit more coming up, and Canadian data protection legislation. So, this is an expanding paradigm where services need to be aware, not just of their obligations in the US, but also the obligations that apply in wherever they are serving children, wherever the children are likely to access their services globally.

Peder, I'll stop there, save time for questions, and hand back over to you.

Peder Magee:

Great. Thank you so much, Michael.

And now, we're going to hear from Bethany Soye about the South Dakota Age Assurance Law.

Bethany:

Yes, thank you for having me. I'm Bethany Soye. I'm a state representative for South Dakota. I serve District 9. I've been in for six years, currently a majority whip and the vice chair of the Judiciary Committee. I'm an attorney, but right now I'm staying home with three little boys, who are all four and under. So as a mom, I'm really passionate about protecting children, especially their innocence in this digital age and as I try to guide them as they grow up in changing technology.

And we really know the scientific damages that pornography does to the young mind, how it changes it, how addictive it is, and how it leads to really violent behaviors. And in 2017, the State of South Dakota passed a resolution that stated that pornography was a public health crisis. So, I looked at that and thought, well, we declared it, but then it's been seven years and we haven't taken any steps to address that. So, it seemed a little bit hypocritical to me and that we needed to take action.

I think before we've gotten into this debate that is growing across the states, the general view has been the digital world is different from the physical world and the default is, well, it's up to the parents. They need to be making sure they know what their kids are doing online. It's up to them to protect them. But my argument was we really shouldn't be treating the digital world any different from the physical world.

So, for example, we know that alcohol is damaging to the developing mind. So, we don't say to the liquor store, "It's okay. Go ahead and sell that to a 12-year-old. It's up to their parents to keep them from drinking it." And yet that's the way that we've been treating the digital world, especially when it comes to pornography. So, my argument was, like alcohol, if you are the one that is producing something that's dangerous to children, you should be the one that's keeping it out of their hands.

So that really led to this bill. It was first introduced in 2024 as House Bill 1257, and then finally it was signed into law last year as 1053. So, Louisiana was the first one to pass a bill in this area, and a lot of states followed their model. So we looked at that and then made some tweaks to it that we thought would be a little bit stronger. The biggest change that we made was what counts as a covered platform. So, the definition that a lot of states were following, Louisiana did and Texas did as well, was one-third. So that means if a website has one-third pornographic material, or in the law it's referred to as material harmful to minors, if there's one-third, then it has to be age-verified. But

Bethany:

... we saw in the Paxton case, and this came up and we were arguing the bill, just the impracticality of that standard and that, well, how do you measure one third? Is it the number of images? Is it pixels? Is it webpages? And during the oral arguments in that case, even one of the justices was asking the attorney, "Can you tell me what percentage of your client's websites have pornography?" And the attorney couldn't or wouldn't give an answer. And it really showed how unworkable that standard was. And the other question was, is there really a compelling state interest if one third of a website with pornography is dangerous to a child, but then you're saying one fourth wouldn't be dangerous to a child?

So, the important change that we made was we decided that a covered platform is the regular course of trade or business to create, host, or make available this material. So it's more getting at what is your purpose? So I think under this standard, if you looked at Pornhub, yes, their purpose, this is what they do in the regular course of business, is to produce pornography versus you look at Facebook. I think that was one of the big arguments. Make sure they don't get caught in there just in case one or two images slips through. But clearly that's not their purpose, their course of business. And then reasonable age verification was another definition that we added, which was spoken to by some of the other panelists. I think a really important part there at the state level is making sure that we leave the options on how to do it on the company. The state just wants to make sure that it happens, but we can't move as fast as technology develops. So we want to leave it a little bit open-ended so that our law isn't obsolete in two years.

So we did a list of, you could use these as options and then a catchall, that's any method that reliably and accurately indicates. So that leaves the burden on the company then to make sure that it happens. And then enforcement was a big discussion and a big political battle. We ended up being just the attorney general who will enforce our law through both criminal and civil penalties. I know some other states have done a private cause of action. And personally, I really think that's effective because they're the ones that are facing the damage. It's their child that has been exposed. They should be the one that's recovering. But there's a lot of negotiations that go on in the political process, so we ultimately just ended up with attorney general enforcement. And then of course we had the ... for the privacy concerns, there was also a crime to sell or retain any identifying information.

So a little bit of our coalition that we had. We did work extensively with the American Principles Project when we were drafting the bill and coming up with the language. But then after that, most of the support was actually in state. Lots of parents that could tell you about personal experiences when their child was just looking for something innocent on YouTube and stumbled across pornography. We had a school bus driver that talked about children sitting in the back of the bus and just watching porn on their phones and just how bad it was. And I think a lot of, especially the older generation, didn't understand how easy it was to access this and how bad it has become. We also had testimony from the criminal investigations here at the state and how they're the ones that are prosecuting these crimes against children and how a lot of the criminals started out with pornography at a very young age and how that damaged them.

Then we also had some lobbying groups from the state, Family Voice, the Catholic Conference, Concerned Women for America, just a really great coalition. And then on the opposing side, it was really focused on the free speech argument. So you remember we were talking about this before the Paxton decision came out. So it was the ACLU and the Free Speech Coalition, which is actually the lobbying group for the pornography industry. And then really more pushback came from an internal fear of litigation. I think that was the greatest fear among legislators. And South Dakota's a very traditional conservative state and generally we like to wait until things are certain and then we'll take action, but I don't think this is an area where you can do that with the way that technology is developing. And also my personal view is that the state legislator represents the people and it's your job to put out the policy that you think is morally correct. And then as it works its way through the court system, it does.

And also, we'll never get great decisions like Paxton if states aren't willing to challenge and to create new laws. So ultimately the first bill in 2024 failed in the Senate committee. It was smoked out on the floor and then turned into a summer study, which was a bit frustrating to me, but it did allow us an opportunity to educate more, to educate the public, and as was stated by the chairman, the public is very clear on this topic. They are very clearly on the side of protecting children. And the Senator who killed the bill in the Senate actually lost his primary election, I think a great deal because of this topic and because of how strongly the public feels. So brought it back in '25. So we had the Paxton case was argued in the fall and then we're in session January through March, so it was between the arguments and the decision.

But you could really tell, if you listened to the oral arguments, you could tell which way it was going. But ultimately that decision, we didn't know if it was going to be a compelling state interest under strict scrutiny or if they choose a lower standard. So our arguments were all based on this is the best chance to be upheld if it's strict scrutiny, if it's the highest standard. And then ultimately, as was stated, it ended up being a little bit lower at intermediate scrutiny. So, the next step here, especially for South Dakota, I think we're running the App Store Accountability Act this year. So another one of the biggest arguments against the bill that came up was the practicality of enforcement, and that's valid. Especially as a state, how much can we actually regulate the internet? How can you stop the internet at your state boundaries, especially when a website might be based in another country?

And one argument was, well, teenagers are just going to download a VPN and then it'll say they're in a different place and they can get around, which is valid. My argument against that was that, yes, we can't stop everything, but we're talking about the five, six, seven-year-olds that are looking for a Disney video and stumbling across this. They're not going to be working around the law. And those are the children that we wanted to protect first and foremost. So now getting to the issue of the VPNs, the App Store Accountability Act, I think will really help with that because it's based on a little bit of a different legal argument. When you download an application onto your phone, you have this long terms and agreement that I'm sure most people just scroll through and just click accept, but you're entering into a contract, and a minor does not have the legal capacity to enter into a contract, especially they're selling their personal data to that company.

So, this bill would require parental permission before a minor can download an application. So, then that solves the VPN problem. If you have a VPN can only be downloaded if the parent allows, so a child can't do it on their own and get around the law. That one hasn't been in committee yet, but will be pretty soon because we have a very short session. We have six weeks left, so we'll be running that one soon. I'm just really excited to be a part of this conversation and the way this is sweeping across the country, and I hope that we can continue to put more safeguards in place to protect our children and I'd love to answer some questions.

Peder Magee:

Great. Thank you so much. Thank all of you, panelists. Those were terrific presentations and the information was great. We've got about seven minutes left, so we can do a few questions. Maybe to start out, Amelia, you talked about what happens once a user is identified as a child. Can you talk a little bit more about what specific protections should follow that user once they've been identified? And others, feel free to weigh in after Amelia goes, thanks.

Amelia Vance:

I think it's a lot of things that the FTC has found in settlements to be problematic. So, turning on location by default. Allowing kids to enter a space in a video game or to engage in chat with strangers immediately after you sign up for an account. Showing targeted ads when you're not logged in and they know you're an adult. I think all of those baseline protections don't kidproof the internet in any way. They protect kids, and honestly, they protect the privacy of all of us.

Peder Magee:

Yeah. Michael?

Michael:

Just to add to that, so the age appropriate design codes would include turning off data sharing by default, profiling off by default, and setting the privacy levels high so that kids are not being contacted by people they don't know. Just from our learning to date of supervising in this area, a lot of companies do this. They have the high default levels there, but what's happening is kids are saying that they're 18 plus, so the defaults are not being actually applied to the children, and hence this is why the age assurance is such a critical bit to this. So, you can have all the defaults you want, but if you're not picking up that a child is actually on service, then none of it's going to be applied to the child.

Peder Magee:

That's interesting. Okay. Mark, maybe you could talk a little bit about the principal compliance challenges that industry is facing around age assurance.

Mark Smith:

Yeah, as I mentioned in my presentation, and I think Amelia followed up with that, there's a lot of regulatory fragmentation there with different standards and different expectations on businesses, especially international businesses. Michael went into the international aspect as well. So, we've got different ages that you have to deal with, different requirements, what's permitted in one jurisdiction versus what's not permitted in another jurisdiction. So, to follow up on the earlier question of what happens when age assurance is used, it should not be viewed solely as a way to restrict children from accessing inappropriate content. It can be used proactively by businesses to ensure that they know there's a minor behind on the other side of the wall there, and that certain obligations will have to kick in based on the jurisdiction.

Peder Magee:

Great. So, looking ahead, what do you all see in emerging digital tech spaces, things like augmented and virtual reality, AI, IoT, how do they pose challenges for age assurance and how can policy makers anticipate those challenges? I'll throw that open to anyone. Amelia.

Amelia Vance:

So, I think AI companions, chatbots in particular pose a lot of new concerns and a balance of you don't want companies looking at everything that your child might type, but you also want to make sure that your child isn't engaging in a conversation that ultimately ends in harm. So, thinking very carefully there about when age verification might come in without being too restrictive is going to be vital in figuring all of this out.

Peder Magee:

Yes, Michael?

Michael:

I'd add to that, that these services often should be looking at what data they are collecting about children. In some cases, it's going to be difficult for a service to know if it's a child or an adult using a connected refrigerator, for example, but does a refrigerator need to collect a child's email address? Does it need to collect their age? Does it need to record their voice? So first of all, encourage all these services to follow a data protection, a data minimization principle, only collect what's absolutely necessary, only hold it for as long as it's absolutely necessary, and then delete what's not needed. The more data they collect, the greater the chance of them doing something ... either something going wrong with that data or a hack leading to the loss of that data, for example. So, a very close examination of what's really necessary and not exceeding that.

Mark Smith:

Yeah, if I could just tack onto that. Yeah, the importance of privacy by design principles need to apply regardless of the technology that you're talking about for any of these newer connected devices or AI chatbots. So, the need for that, but also we recommend the use of privacy enhancing technologies as well, to the extent that information that is collected can be anonymized or otherwise pseudonymized. That could be also a factor in producing these technologies.

Peder Magee:

Michael, I think you mentioned in the UK the use of AI around age assurance and there being a requirement for human intervention. Can you talk a little bit more about how that works?

Michael:

Yeah, at the moment we're seeing a lot of social media or wider website companies starting to use profiling for age assurance. What we've noticed, for example, is that a service will allow people on and then profile them their activity online to determine whether they're a child or not, and then make a formal decision. First of all, there's problematic there because in most cases the service will default to self-declaration because it takes a while to build up enough content to know whether somebody is below or over 13, all of which time they're probably processing that data of the under 13s unlawfully. But we found is in the circumstances where you have that sort of profiling approach or an AI-based age assurance that makes a decision about a child without any chance for human intervention, and that was where they might run afoul of the automatic decision-making requirements within the GDPR.

None of these things is absolutely foolproof, so there should be opportunities for an error to be adjusted because we're talking about basically an equivalent of legal impact on a child that they would be denied a service that they should be rightfully able to use if they were more than 13 or more there, for example. But what's really important is that where services have that sort of backup, that second stage of appeals or complaints process, that's tied into the age assurance process. We've seen in the Discord case of last year where the complaints process was handled by a third party, and then there was data leaks associated with that. So the processing of data, the protection of children needs to run throughout the whole system. That's there to determine whether a child should be allowed on a service, whether they're the appropriate age, and the age assurance needs to build that in to make sure that everybody, the children, their rights under the GDPR or the UNCRC or COPPA are recognized and then they can appeal when necessary.

Peder Magee:

Great. So, another question, and I'll put this open to anyone, but maybe Bethany, from your work on South Dakota's legislation, how can these laws ensure that the parental consent requirements are empowering parents rather than simply shifting the burden to them? And anyone who has kids knows how it's hard to keep track of everything they're doing online, maybe you could talk about that.

Bethany:

Yeah, definitely. I think that's the most important part of these laws, is that parents, they are trying to protect their children, but everything is changing so quickly and there are so many different things to protect them from that it's hard. And that's how we had parents come and testify that I do everything I can to look at their phone and make sure they're not getting on this, but I don't have the tools that I need. So, I really think the App Store Accountability Act is really going to be helpful in that, because it's going to streamline it and say, "Here, your child wants to download this app, yes or no." And I wish that it would just happen naturally that technology companies would develop this and would make it easier for parents, but that doesn't seem to be happening, so I think that's why it's important for us to take policy actions.

Peder Magee:

Okay. Well, I see that we are a little bit over time, so why don't I wrap up with a final question, and it'd be great if each of you could try and answer this one. In the title of our panel, we ask why age verification matters. Maybe you could each take a minute or two and explain why you think age verification or age assurance matters. And I'll just go in order. Mark?

Mark Smith:

Sure. I think it matters because we're talking about a technology, the internet, that was not built for children, and it was never built with children in mind. So given what's out there now and the amount of content that's not appropriate for children, we need to have safeguards in place to protect children to ensure that they are consuming age-appropriate information and age-appropriate content. So, the only way to do that, we're coming in after the fact, but that's how age assurance can help solve that underlying problem of the internet not being built for children.

Peder Magee:

Amelia?

Amelia Vance:

I think age verification can be a step towards empowering kids, as noted. Kids don't necessarily want certain pop-ups or for them to type a letter wrong in a website and have things show up. Kids want to be able to access useful content. So, by making sure that some of those materials they're not ready for are behind age verification can be invaluable. The key is just remembering a lot of the worst abuses of companies, the scandals that we've heard about, were from companies who already knew their age. So, we need to make sure we cover that as well.

Peder Magee:

Great. Michael?

Michael:

Well, I've already mentioned the issue of having default settings and protections that aren't being applied because service doesn't apply an effective age assurance. So that's one thing. But I'll go beyond what Mark and Amelia already said, and part of it is because it's the law in many jurisdictions. You can't process children's information lawfully. Well, it depends on what the lawful basis is, but it might be that you need age assurance in order to show that you're not unlawfully processing the data of children who are not supposed to be on that service and which you don't have a lawful basis for. So having age assurance not only protects a child, it protects the company from breaking the law in jurisdictions where there are limitations on what they can do.

Peder Magee:

Great. And Bethany, why don't you take us home?

Bethany:

Yeah, thank you. I think this area is connected to so many other policy problems that we see. Especially here in South Dakota, I mean, we have so many debates over juvenile justice, over behavioral problems in schools, over mental health. All of these things are really coming to a crisis, and we're having to extend a lot of states' funds to fix these problems, but really what are they stemming from? A lot of it is from interactions online. A lot of it is from children accessing things that they're not developmentally ready for. So, we're having to deal with the fallout further down in our justice system, in our public schools. So, if we could get to the root of the problem and protect the innocence of children, protect them developmentally before they develop these problems, that would just be a huge saving for all of our communities and our state.

Peder Magee:

Okay. Well, that's great. I want to thank you all again for your participation, it was a terrific panel, and it is now 10:38. We're going to take a very quick break. We're going to start back up at 10:45, and thank you.

Mark Eichorn:

Welcome back everyone, and thank you for a great first panel. Thank you to Peder Magee and to all of our panelists. One observation, the volume of legislative activity at the state and international level is really just stunning. So thank you everyone for setting the stage. Now I'd like to introduce FTC Commissioner Mark Meador for the morning remarks. Commissioner Meador.

Commissioner Meador:

Thank you, Mark. Thanks so much to the panelists for being here and to the chairman and the commission staff who helped make this event possible. It's an honor to be here with you all for this very important conversation. Whenever I read about the different generations in the mix, millennials, Gen Z, and now Generation Alpha, I'm struck by a curious phrase that keeps coming up. It's the phrase digital natives. And what's so interesting about that phrase is what it implies about much of the world of high technology we've made. Talking about a generation of digital natives implies that we live in a world transformed by impersonal historical forces that a whole generation of young people happened to be growing up within. The world just changed almost by itself.

It wasn't that the world was changed, that someone might've been responsible for why things turned out this way. But of course, someone is responsible. The online world in which my children and many of your children are growing up is a world profoundly shaped by the decisions of powerful people in high places. For the last two decades or so, these same people have been running an elaborate set of economic, psychometric, and socioemotional experiments on America's young people. Those experiments are meant to ferret out information, what they're anxious about, what they're hoping for, what keeps them hooked on their phones, what will make them customers for life as they grow up and enter the workforce. All of this has been carried out on the same population that we euphemistically describe as digital natives. So we might equally as well and perhaps more accurately describe them as digital subjects. Earlier this month, I gave a talk in Palo Alto about the limits and possibilities of the idea of innovation, and I mentioned there's something that's kept coming to mind for me lately, just how much of a gap there is today between the tech future we hoped for in the 1990s when I was coming of age and what we actually got. We thought we'd connect with each other. We'd create new things. We'd learn to talk across our differences. Mass monetization of America's children is not the future we hoped for, nor are the predation and extremism that seem increasingly to define our encounters online. Overwhelmingly, the cost of these failures aren't borne by the adults responsible. They're borne by the children and teenagers subjected to them. Let's look at the figures since 2010.

Suicide rates have spiked, increasing by 91% for adolescent boys and 167% for adolescent girls. Emergency room visits for self-harm among adolescent girls have increased by 188%. Depression rates have surged and things don't look like they get better with age. Anxiety among individuals between the ages of 18 and 25 is up by 139%. Now, correlation isn't causation, of course, but this is a pretty suggestive pattern. And what it suggests is that the more we've been connected digitally, the worse off we've become. Bitter irony. The age of the smartphone and social media is for too many children an age of suffering. Now the fact that we're here today having this conversation is a big change, a positive change from even just a few years

Commissioner Meador:

Years ago. We've recently seen a state and national level movement to take smartphones out of schools, not because we hate smartphones, but because we realize there's a time and a place for them. And middle school classrooms are definitely the wrong place for Snapchat and TikTok. In some quarters, you're starting to see a backlash to the backlash. Folks are arguing that everything is fine, that we're simply living through a new era of kids connecting with each other through new media.

If you walk just a few blocks over to the Smithsonian's National Museum of Natural History, they've got a special exhibit up on the second floor called "Cellphone." Funny, I thought the point of this museum was natural history, like dinosaur bones and gems. But in any event, there's a display in that exhibit that essentially claims that concern over smartphones is just a moral panic, that people raised hell about the invention of writing, the telegraph, the TV, and so on. So of course, that's what's happening today. Everyone just needs to stop worrying and get with the times.

But again, this is the same logic of inevitability that lets us speak of digital natives rather than digital subjects. It's a logic that denies that as a society, we have a moral responsibility to keep young people safe from harm. We don't accept this in other contexts. For one thing, we don't treat other addictive or potentially dangerous goods this way. A 12-year-old can't walk into 7-Eleven and buy a pack of cigarettes. A 16-year-old can't stroll into a liquor store and buy a fifth of bourbon. And it wasn't just physical, consumable substances that had restrictions.

It was content too. When I was growing up, GameStop wouldn't sell M-rated video games to preteens. In fact, I talked to the head of the ESRB last week and they still don't. AMC wouldn't let 13-year-olds buy tickets to R-rated movies and so on it goes. Now you might've noticed that I haven't used the phrase age verification at all yet. And that's because while age verification is something of a hot new topic in the digital world, the basic principle behind it is nothing new at all. For decades, even centuries, we have had community standards around certain kinds of products because we recognize that giving young people unrestricted access is a bad idea. This is not authoritarianism.

It's not a violation of anybody's free speech rights. It is an acknowledgement of the lived reality of life, maturation, and growth. Anyone who actually has children, rather than just having opinions about them, grasps this intuitively. Now, over the years, whenever the topic of age verification comes up, we tend to hear a certain cluster of criticisms. Some of these criticisms, in my view, are made in bad faith. Here's the biggest one. "Just parent better. Moms and dads out there, if you were really doing your job, we wouldn't need age verification technologies." So the argument goes. But here's the thing.

We don't accept this argument in other contexts. We don't get rid of the requirement that you show your idea at the liquor store on the grounds that parents should just parent better. No. We acknowledge that as important as it is that parents parent well, there are social backstops that still matter. It's not about usurping the role of parents. It's about making their lives easier. There's another criticism that often turns up. Age verification, we're told will make the lies of adults harder. We hear that we won't be able to download basic apps like calculators without having to submit to an owner-ish age check.

There's no reason this process needs to be cumbersome and messy or invasive. Earlier this month in Palo Alto, I spoke about how when we think about innovation, all innovations aren't the same. There are bad innovations like dark patterns that addict and deceive, but that doesn't mean innovation as such isn't worth supporting. And when I look at the landscape of age verification technologies today, I have to say, I'm incredibly impressed with what entrepreneurs are coming up with. Just as policymakers have grown more interested in these measures as a way to keep kids safer online, the market has responded to do what it does best, meet the needs of the moment in efficient and sophisticated ways.

With the new age verification systems that are emerging, you don't need to hand over your personal data or your child's personal data to a company you might not trust. Instead, these systems rely on third party providers who keep that data secure. Third parties who can contract with social media companies or other online service providers to simply verify whether a user is old enough to access a product without turning over any raw personal data. This is elegant, it is efficient, it is secure, and it is the future.

But there are other possible futures too. Behavioral age verification, that is ascertaining a user's age by the way they interact with an online platform or system has always been a major challenge.

And this strikes me as one of the best uses cases for artificial intelligence. Machine learning can help detect patterns in browsing and usage behavior that consistently indicate whether a user is too young to be on the platform. American companies can help lead the way in pioneering this. That is what real leadership will look like in the years to come. Now, one of the watch words of this administration is, and something I'm personally committed to fighting for is affordability. As I see it, American families have enough on their minds right now.

The last thing they need is to have their kids' data harvested and monetized by multi-billion dollar tech companies or watch their kids suffer from premature exposure to the worst the internet has to offer. Age verification offers a better way. It offers a way to unleash American innovation without compromising the health and wellbeing of America's most important resource, its children. It is grounded in practices of responsibility and stewardship that extend across our entire history. It is a tool that empowers rather than replaces America's parents. Really, I don't know that we can afford to forego it. Thank you for your time.

Elizabeth Averill:

Good morning. Thank you, Commissioner. My name is Elizabeth Averill and I'm an attorney in the Division of Identity and Privacy Protection. I'd like to introduce our second panel, which will be focused on the discussion of different types of age assurance technologies. I'd like to first briefly introduce our panelists. Their full bios are available on the FTC's website.

First, we have Iain Corby who serves as the Executive Director of the Age Verification Providers Association. We also have Sarah Scheffler, who is an assistant professor at Carnegie Mellon University. Go. Okay. Thank you, commissioner. My name is Elizabeth Averill. I'm an attorney working in the Division of Identity and Privacy Protection. I'd like to introduce the second panel, which will be focused on a discussion of different age assurance technologies. I'd like to first briefly introduce our panelists. Their full bios are available on the FTC's website.

We have Iain Corby, who serves as the Executive Director of the Age Verification Providers Association. We also have Sarah Scheffler, who is an assistant professor at Carnegie Mellon University with CyLab, the CMU Security and Privacy Institute. We have Jim Siegl, who is a Senior Fellow with the Future of Privacy Forum. We also have Rick Song, who is the CEO and a co-founder of Persona, a global identity and age assurance platform. I think we might be joined a little bit later by Denise Tayloe, who is the CEO and co-founder of PRIVO, a COPPA Safe Harbor organization. We'd like to start the panel by providing an overview of current and emerging age assurance technologies. If we could start with Jim Siegl. Thanks.

Jim Siegl:

Great. Thank you. I want to start out this morning by sketching out a roadmap that I hope will be useful as this panel unpacks the technologies in the age assurance landscape. This is based on a recently updated infographic from the Future of Privacy Forum, which is available on our website. So if I can have the next slide. Great. So really, this combines a lot of the questions that we're going to be talking about in the next 75 minutes. So I want to start with some key questions which are highlighted on the left side. And that's first really, what is the goal of age assurance?

Typically, we think of this as to place an individual within an age threshold like 18 plus or an age band like 15 to 19 in order to deliver an age-specific experience and provide or restrict access to an age-restricted service, which in the case of COPPA, that age-specific experience could include triggering verifiable parental consent. So another important question is for a given service, if there is no need to restrict users or deliver an age-specific experience, is there a need to collect data for age assurance? And I think this was something that was raised in the previous panel.

Some other key questions when thinking about age assurance methods is, is the level of assurance balanced with the privacy risk? We talk a lot about proportional age assurance risk. Where in the technology stack does age assurance happen and how does that information flow? When evaluating specific age assurance mechanisms, we need to consider both the accuracy under different conditions and the privacy impact so the method is balanced against the level of assurance. This is an important point that Amelia Vance highlighted in the first panel. Can I have the next slide?

So let's look at the four categories of age assurance method that are currently in deployment or emerging. This was briefly raised at the beginning of the first panel. And the first thing to understand is that age assurance is an umbrella term. There's no single technology here. It's a spectrum of methods. Each one has different levels of accuracy, privacy implications, user friction. And the challenge I think for policymakers and for platforms is matching that level of assurance to the risk.

Unpacking each of these, we start with declaration, age gating. While this is the most common age assurance method where a user self-asserts their birthdate, I think it's problematic to call this an age assurance method. As Michael Murray noted in the first panel, while the privacy risk is arguably low, it's most appropriate for very low risk situations. It's easily bypassed by children or adults posing as children. I think it's also worth considering that exact birthdate, especially when it's combined with additional data like name and zip code, can be a unique or near unique data point for distinguishing individuals and makes a primary target for identity theft.

Inference, an emerging area and one that in the previous version of this demographic we had combined with estimation. We've broken out now as a separate method. Inference draws reasonable conclusions based on contextual, behavioral, transactional, environmental signals. So for example, a long-established email or certain financial transactions might strongly infer that someone is an adult. This isn't just a point in time method when age is determined. So both recently, OpenAI and TikTok have implemented AI-driven inference systems designed to operate continuously using behavioral signals. Estimation uses AI and machine learning to deduce a likely age based on biological traits like facial image or voice or typing patterns.

And verification is a high assurance method that references typically authoritative verified dates of birth from government IDs or databases, and it usually or often involves the scan of a government ID matching it with a live selfie or some type of facial recognition. I think a critical concept here is that these aren't isolated. The idea that these can be layered or used in a waterfall approach, so that platforms don't need to rely on a single method. They can deploy successive validation, starting with lower friction, more privacy-friendly approaches and escalating as needed. We'll look at this as an example in a little bit.

Going on to the next slide, I want to briefly highlight a few emerging technology concepts that are reshaping age assurance and that I suspect my fellow panelists will be touching on. First, age signals and tokens. A signal is a real-time communication of a user status that over 18 are placing a user in an age band while an age token is that signal or proof that's locally stored in a user's browser or on a device. We're seeing a real potential shift between one-time checks, each site requesting age assurance from a user to reusable age credentials. So there's a lot of friction in repetitive identity checks.

This creates a lot of data exposure. So reusable credentials can store a verification in a secure token enabling age assurance confirmation across platforms. However, it's important to consider that that interoperability is currently limited to specific trust frameworks as universal standards for platform, cross-platform recognition are still evolving. The idea of a double-blind architecture, that an external service verifies your age without knowing where that data's going to be used, and that the website confirms you're of age without learning your identity, so there's neither party gets the full picture.

And with user binding, ensuring that this age assertion is linked to the actual person that's using the device, not just to the device itself, so that if a verified adult hands their phone to a child, that the binding mechanism is using biometrics to prevent unauthorized access. So on the next slide, we'll see some of these concepts in action. So let's consider the scenario of Miles, a 16-year-old accessing an online gaming service. So the initial experience begins with a low assurance age declaration to play the game.

When he tries to enable a 16- plus feature like video chat, the system triggers age estimation, so for example, with a selfie. And thinking about the concept of accuracy, applying an age buffer of three years creates a gray zone, so a range of 15 to 19. And because Miles falls within this buffer, we need stronger verification. And typically, most 16-year-olds, about 75% of 16-year-olds don't have a driver's license. So offering another option like parental vouching where a verified parent confirms the age. And then lastly, that binding verified, once verified, the age credential is bound to say a device pass key.

So that ensures that if Miles shares his phone with a 15-year-old friend, that friend can't access those 16-plus features without the correct local biometric or a PIN. And as we go to the last slide, want to look at risks and challenges and mitigations. We face challenges like loss of anonymity and secondary data use and potential for false negatives and positives. And my fellow panelists are going to be talking about this in more detail, but I want to highlight that as part of any risk management strategy, it's important to match the potential risk with an appropriate mitigation. And with that, I'll hand it over to Iain to dig more into the methods.

Iain Corby:

Thank you, Jim. And hello, everybody. My name's Iain Corby. I'm the Executive Director of the Age Verification Providers Association. So we're a global trade body. We've been representing the age verification industry since 2018. We're not for-profit, we're politically neutral. We believe strongly in standards. We've got about 34 members around the world, a third of which I think I would say are broadly speaking headquartered in the USA. And we say that our mission is to make the internet age aware. And I'd just like to emphasize that's age aware, not identity aware. Now, I agreed with Jim beforehand that he would effectively handle the theory and I would try to demonstrate the practice.

So there will be an element of repetition here, but I imagine the audience today is divided between those who are extremely familiar with this topic and then for many, it's all brand new. So I hope those who are familiar will forgive us with a bit of repetition to help people understand some of these important foundational concepts. The main one being that the essence of age verification is proving your age without disclosing your identity. So there's a whole separate identity verification industry, which is all around you opening a bank account and proving your identity online to get government grants and so on.

That's not me, that's not us. That's a different part of the world. In fact, it's quite often very different technology as well. So let's focus today on age verification. And when we started this sector, in fact, it was really encouraged by the adult industry who seeing, particularly in the UK, a requirement for age verification coming in, they recognized very quickly that their users wouldn't be comfortable sharing personal data directly with those sites. So effectively, they encourage third parties to step in and say, "Look, we'll handle the proofing process, and then we'll just say yes or no, are you over 18 or not?"

And that would be the only information that those sites get. So in the most basic way, identity was protected just structurally by having an independent third party handle that process and then make sure that none of the personal data was passed on. Now, obviously that's moved on a long way since and technology has allowed us to put in even more protections than that, so it's not just a structural protection. But we've maintained this recommitment to data minimization. Starting again in the European context, we have data protection laws such as GDPR.

These are not at the federal level in the US and many states don't have data protection, but what you do see invariably is age verification laws at the state level requiring equivalence to data protection provided in Europe and in particular immediate deletion of any personal data after an age is established. Jim also mentioned the double blind option. This is something which was really inspired by the French regulator, but has been increasingly adopted as a required option at least in legislation where it's impossible for the website to find the identity of the user.

And likewise, it's impossible for the age verification provider or whoever's doing that proofing to know which website the user wants to access. Because we're not only trying to prevent the website finding out who you are, we also don't want to create a track record of which websites you're looking at, which might be compromising and could be used perhaps to blackmail people. So double blind has become an increasingly important option, which is delivered through the privacy enhancing technologies mentioned earlier. All of this technology can be independently audited and certified.

We have international standards, apologize for a couple of typos here. It's actually IEE289.1 and 27566-1, but yeah, these are the standards on which our assurance processes are based in the age certification scheme issues certification around the world for that, which is not just about accuracy, it's also about privacy and data security. And finally, there will be debate later today, I'm sure, about where best you should do the age check, should it be done in the operating system or the app store, or should it be as with common at the moment on the publisher, on the digital service itself, the operator?

I suspect the answer is in most places we believe in a layered approach so you get as much protection as possible, but also different use cases with different levels of risk and particularly different liabilities for the websites concerned are probably going to use a different approach. So next slide please. So what I'm going to do now is just counter through a number of the different methods. And what I'll do is just ask the slide operator to hit play on each of the videos as we get to them rather than just asking me to hit play. So they're just there for you to look at while I talk.

So this first one is the classic document verification where you're just seeing essentially that you need to find a driver's license, a passport, military ID or something. You're going to show that to the camera and then you're going to follow up with a selfie image so you can confirm that the picture on the ID is the same as the person who's showing it to you. The ID might be read with optical character recognition, or you might, if it has a chip in the ID, read the chip using mobile phone.

And then there are things like anti-spoofing checks that are involved to make sure that neither the ID is a fake generated by AI, for example, nor the person is a fake, also potentially generated by AI. And these days, that technology can be shrunk down so the whole process is actually handled on your mobile phone. It's not ubiquitous that that's that way, but in theory, it could be done at that level. Next slide, please. So the next method that we'll just have a look at here is reusable digital identity. This is becoming increasingly popular and it could be done with a state issued mobile driver's license.

It could be a privately issued Digital ID app, which is an example we're looking at here from a company called Luciditi, but others like Yoti also offer those. And you're keeping a verifiable credential in your wallet. So you can actually also keep that in your Google Wallet, for example, or your Apple Wallet. And then these are shared as a selective share just of the fact that you're over 13, you're over 18, whatever the age is you're trying to prove. There may be authentication mechanisms involved. That could be just a pin number, which isn't as secure as a biometric.

Obviously you can share your pin with somebody. Biometric, which is your face or your thumbprint, you cannot hand that off to somebody else. We do have to also be a little bit careful that device-based biometrics, for example, face ID on your Apple phone does sometimes allow you to have a second face. So you could hand that to a younger sibling, add them to your phone and send them down to the store to buy beer. So some of the more sophisticated apps will actually record your face and image at the time that you create that credential and the only way that you can use it is with that original image.

Next slide, please. Authoritative data is the other way we go about trying to find this information. And typically that came originally from credit reports. Oh, sorry, we've got a bit of background noise here. Apologies for that. Just let that play out. It's quite short. So what's happening there is the users selected their bank. They've chosen the bank that they bank with and payment, for example, to make sure that that to the credit reference agency, as I said earlier. We need to make sure that it belongs to you, and so that can be done through knowledge or logging into online banking, logging into your online bank or a micropayment, as I said.

This is also going to be important for children. And one of the things we are looking at in Australia is how do we find the ages of 16 year olds who don't have credit and driver's license and passports. And one of the things I think we do need to see is joined up government here and where government is asking us to do a strict verification at 13 or 16. They're going to have to find ways of giving us access to health data, health insurance data, schools data, places where we can get access to this. Apologies, my signal may not be quite as good as we hoped. The next slide is facial age estimation. This is... Just moving on. Yeah.

So this is where we create a mathematical map of the face, images of people whose ages we do know of ages that we've looked at in the training data. We don't use enough data to identify the user. Again, this is something that some companies can do entirely locally on your device. You're not even having to share that facial, that mathematical map of your face with a server, but the image is never stored anyway. And in fact, so it's not even held in on a disc.

It's just processed live in order to do that calculation. Next slide. Just to give you an idea of sort of state of the art about accuracy here, we're looking at plus or minus two years of the real age. It does vary slightly by age. Obviously, we focus a lot on the ages that matter. So six to 12, for example, this is an example from Yoti where their true positive rate for 6 to 12 year olds is estimating them as being under 13 is 99%. So bear in mind the baseline today is 0% because we're just asking people to self-declare.

Now, obviously some people think this should be perfect technology, but what we would tend to do is use this technology in partnership with a buffer age so that you wouldn't actually check that somebody, for example, is 21 if you're asking if can they buy alcohol. You check that they looked at least 25, for example, which would much increase the level of accuracy and the certainty that somebody is actually 21 if they're estimated to be over 25. Next slide, please. So the next method that we'll just consider is quite a novel one. This is where you just move your hand on the camera. It turns out that the way we move our hand is very much affected by age, the tendons and so on change.

This was a discovery where people were actually looking to figure out whether people were taking performance enhancing drugs in sport and inadvertently discovered that they could also figure out whether you're over or under 18. And that is now being adapted for different ages as well. Again, pretty accurate, a false positive rate of around 2%. The organization makes sure that they, it's called need demand. They make sure that this product isn't good enough to be able to read your fingerprints, for example. There was an example of how a German politician who appeared on camera had their fingerprints stolen just because they were broadcasting their hands.

So again, that was something which we had to make sure you did breach that anonymity. And if you put your face in the picture, the whole thing stops because they don't want to actually see your face. Next slide. This is a relatively new approach, which is based on metadata. So how have you been using your email or your mobile phone number, your cell phone number? Have you used it to lease a car, to buy a house, to get credit? And if you find a hundred different examples of how somebody's used their email address, for example, then you can be fairly certain that they're an adult. Obviously that person needs access to the inbox. Likewise, if you're using a cell phone number, you need to have access to the mobile phone to be able to get a one-time password and just plug that in. And this has proven to be quite a popular way with a lot of users who are happy to share their email address or just their cell phone number in order to very quickly prove their age. Next slide, please. This is a very similar option developed by PRIVO who are a FTC COPPA Safe Harbor. And when they do their age verification... Oh again, sorry, there's a bit of noise here.

Actually, just stop the presentation if you wouldn't mind. Oh. Thank you. Sorry, again, apologies for the sound there. What they actually do is they just check that you have an email address associated with a company that only employs adults. So therefore, if you work, for example, for the FTC, chances are there aren't many 12 year olds on the FTC payroll with access to the FTC domain for their email address. And so that's a quick and easy way to check your age. The next option, I don't have a video for on the next slide, but this is around age inference, something that has been referred to earlier by [inaudible 02:16:41], also Michael Murray at the ICO.

This is where you look at basically user generated content. What has the user been doing online? Who are their friends? Have people said happy 12th birthday? How many candles were on that particular post about their birthday, for example? You could even be slightly more sophisticated with natural language processing and looking at how people are interacting with the app. This is something which we've seen being used a lot in Australia and with what has been a very phased introduction of a 16 plus social media delay as they call it. One of the problems, as Michael pointed out earlier, is you don't at the start with a brand new account have any data on which to do inference.

So really you need to be doing an age check for all new users. But this is one way to check over time that those who are perhaps signed up previously with a fake age, you can weed them out by using some of this inference technology. Next slide, please. This is then the EU's example. We will play, it does have some sound, but we can listen to that. So please go ahead. Oh, even better without the sound. Thank you. Great. So this is just how the EU works, very similar to the Digital ID version, the reusable Digital ID we saw earlier, but the EU is trying to roll out as an adjunct to its European Digital Identity wallet, which each member state has to provide by the end of this year.

A system based on the open ID for verifiable presentation protocol where essentially an age restricted site, in this case, it's a cinema. We just put a QR code up. You have an app on your phone where you've created a batch of, say, 50 age tokens, which prove your age. They've had to create this batch in order to try and approximate to that double blind process we were discussing earlier because the EUDI Wallet wasn't originally designed to deliver anonymous age verification.

But anyway, this is something which is being added to EUDI functionality by a lot of member states and the EU is hoping is going to be a standard for age verification. Next slide, please. And then finally, I'm going to just come on and talk about a couple of the latest developments. And I think 2026 is going to be the year of interoperability. The first one comes from the OpenAge Initiative, which is given to us by a company called k-ID based out of Singapore. And they're using past key technology, which you may well be familiar with where you can, once you've done an age check, accept a pass key, which they obviously call an age key onto your device.

And this is the system where you may be familiar with it. Occasionally you're asked to put your thumb on the thumb reader or the fingerprint reader on your device or you do a face check on your computer and it just allows you to reuse a saved credential. Now that is not the same as making a device check because it's not associated to the device per se, it's also associated to the individual who created it.

So you can obviously therefore have a biometric authentication about that. With this particular solution, it's up to the digital service to decide which age keys it wants to accept. So it would perhaps consider what method was used to create that key, which issuer created that key. So how good is that issuer? It might look to see if that issuer has been certified to decide whether it's willing to

Iain Corby:

... accept those AgeKeys. Next slide, please. And then the alternative interoperable solution is AgeAware, which was developed originally by an EU-funded, project and then by Safe Online, which is a United Nations fund, but is now effectively an alternative to the previous one that we saw. And this is, again, a similar tokenized solution, where having done an age check once, where you can choose the method that you prefer, and then having chosen the method, you can then choose the supplier or the issuer that you trust.

And then, having done an age check, in this case it's using a mobile phone number, you're able to accept a token, which is in this case held on a progressive web app in your browser. So it's not using the passkey technology, it's using slightly different technology. In this case, they do have a minimum standard, and they're effectively telling the relying parties, the websites, that if you're using one of these tokens, we will only allow you to use a token that is fit for purpose in the jurisdiction where you're trying to be compliant for the use case you wish to be compliant with.

So on the next slide, there's just an example of what happens when you go back to the website, and effectively here, you won't have to go through the process, because you've already been authenticated. So you go to another website, could be served by a completely different age verification provider. And when you agree to verify with that AgeAware solution, that interoperable solution, instead of having to do a new age check, it's just going to allow you to go straight into the website that you're trying to access.

So those are the examples. So just in conclusion, there's a whole wide range of methods here to give consumers choice, so they can choose something they're comfortable with. This was something the Supreme Court thought was important in the Paxton case. Cryptography can guarantee privacy. So you're not reliant on the goodwill or the good behavior of the providers and the websites, that cryptography can guarantee it.

There are government ID options, but also private sector options. Some people are pretty distrustful of government, they prefer to choose a private sector option. And proportionate solutions mean that you'll have different approaches, methods, interoperability networks, different places in the tech stack, depending on the use case and the level of risk. And finally, you can have certification underpinning this entire system to guarantee privacy, security, and accuracy. So apologies if that took a little longer, but hopefully people have now got a slightly better idea of what this looks and feels like. Liz.

Elizabeth Averill:

Thanks so much, Jim and Ian. Denise, do you want to add something?

Denise Tayloe:

Hi, Elizabeth. Thank you. Great job, Ian and Jim. I just wanted to make sure that we level set on age verification is not equal to verifiable parental consent. Because you'll note from all of what we've looked at, there is nothing here that measures a reasonable method in light of available technology to ensure the person providing consent for the child is the parent or guardian.

So relationship verification is really the next horizon, and a piece that I just want to make sure people walk away from this understanding that age verification is a component of all of this. Secondly, when we talk about reusable identity, I've been a big proponent of that. Ian and Jim both know me for a long time in that sense, and offering people choice. But we have to remember that a method that can be used in the home, so let's take an email, for example, you'll see that we have an email method as well. We stuck with where you might work, because we found in our testing that parents didn't leave their work email widely available for their children to use.

But when you're talking about doing a method where the child is in the home with the data, or they can leverage the parent in the home without the parent knowing, and then creating reusable identity from that point forward, we just have to take into consideration the fact that in another home, somebody's not going to use my email address. So your child is not going to use my Gmail to get through a system, but my child in my home can get access to my cell phone pretty easily. So I think that we still have to consider levels of assurance along the way. Thanks, Elizabeth.

Elizabeth Averill:

Thanks, Denise. Let's move on to our second cluster of questions, and we're going to start with Rick Song. The questions are, what are the relevant performance metrics to consider, and which technologies are most accurate, and which are most susceptible or resistant to circumvention, particularly from children? Thanks.

Rick Song:

Thank you so much, Elizabeth. Next slide. So hi, everyone. I'm Rick Song, co-founder and CEO of Persona. We're an identity verification platform that verifies billions of identities every year, and partner with many of the leading enterprises across a variety of industries to help build these secure online experiences for children.

My background's is an engineer, and public speaking isn't my natural habitat, so I'm especially grateful to be supported by my fellow panelists today, who bring a lot of their deep expertise in this space of age assurance. But fortunately, I'll be talking about the deployments of the technologies behind all of this, which I feel significantly more comfortable to talk about. Next slide.

The first thing I'll say is that Persona's technology-agnostic, in that we actually believe, as Ian kind of shared, there's a lot of different approaches to how you verify someone's identity, how you verify someone's age, and ensure that someone's the proper age they assert to be. And our perspective is that depending on what use case, depending on the target demographics and the kind of access that you're trying to offer, it's really important that you apply the right technology.

And for us, we've actually seen every single one of these technologies out in the wild. So we want to discuss a little bit about the relevant performance metrics, what matters, what has worked, what people opt into, and how to best apply all these technologies to balance usability, assurance, and most importantly, privacy.

And one thing before getting at it all is that the goal of a lot of today's age assurance approaches is not to be perfect, but to meaningfully improve on self-attested age. As Jim earlier had spoken about, we're really thinking about, self-declaration of age is easy to circumvent. If someone's just attesting that I'm of this age, for both adults pretend to be children or children hoping to attest to be adults, it's easy to just put in something and attest that you're something different.

Age assurance is really trying to figure out how can we balance that, improve meaningfully above self-declaration, and balance the conversion and usability with a level of assurance that's better. So here we've listed out all the approaches, Ian's done an incredible job sharing them. And as we go, next slide, one thing we really want to call out is that we oftentimes think about this in three major ways from a metrics perspective, the first of which is coverage. It's incredibly important that the technology, first and foremost, is applicable and usable by as large of a population as possible, that as many people in the world are able to leverage the technology and have access to it. The second is assurance. We need to think a lot about, does this apply the level of confidence that the user is the age they claim to be, especially for the use case that we're targeting for?

And lastly, the usability. Is it easy to use? Can people get through it? Is it not deterring and meaningfully harming the existing experience that users would otherwise be going through? Next slide. And as we talk a lot about age assurance, we oftentimes think about things from ensuring that kids do not get access to adult content. So age gating, ensuring that age is above 18. But one thing that we really want to discuss, and as we think about the assurance of things, is also designing age-appropriate experiences. This means that, I think many folks all believe that children should continue to have access to online experiences. It's ensuring that children have safe online spaces, that folks who are the age that they are and that children are interfacing with other children who are actually the age that they attest to be.

And we think of this as applying to a variety of things. For oftentimes for age gating, we think of it from the perspective of inappropriate content, making sure that children aren't getting access to adult content, making sure that they're not purchasing things or services that should be age restricted. Whereas oftentimes on the age-appropriate experiences side, we think of things a lot more from protecting online communities, ensuring that children are playing games with other children. Next slide.

And there's different performance of these technologies depending on what goal that is. There's a lot here, but we want to call out that right now, every one of these have different trade-offs. For selfie age estimation, we see rather high coverage, assurance, and usability. For government IDs, it might be a bit worse on the usability, but high coverage and assurance as well. As we go further down, we're really excited about the newer technologies emerging right now, like reusable age tokens and digital IDs, which right now, although still novel, we believe have a really great balance between assurance and usability, and on top of that, the future of privacy as well.

For age-appropriate experiences, this is where we oftentimes find a lot more challenges as well, because a lot of these individuals don't have access to the technologies that would otherwise be available to children right now. Denise, I think, will be covering this in just a bit, but parental consent is an incredibly, incredibly important topic. And right now, the ability to tie the relationship between a parent and a child is more important than ever, as we think about designing age-appropriate experiences.

I'm going to talk a little bit now about common circumvention techniques for each of these, for age gating versus age-appropriate experiences and how folks oftentimes try to be able to bypass these, as we've seen in the wild and in real-world use cases. Next slide.

So for age gating, we oftentimes see these three as the most commonly attempted circumvention techniques. And I want to call out that each of these circumvention techniques within this space, folks like Iain, Jim, and others have spent a long time thinking about how can we make sure that we're designing the right technologies to ensure that we're balancing? And there's a lot of ways to prevent these circumvention techniques as well.

But speaking of the circumvention techniques, the first is really around parental impersonation. We see this a lot in which a kid will present their phone to their parent, let them know that this is just, "Please quickly scan your face," and then they're able now to pretend to be a far greater age than otherwise. Oftentimes, borrowed government ID is also a very, very common approach. This means going to a sibling and requesting that they get their government ID, asking their parent if they could just see their ID very quickly, and scanning that to be able to attest that they're a certain age.

And then lastly, spoofing their location. Right now, age assurance is not required across the world, so pretending, rather than being in a country or state that requires age assurance, modifying their device, downloading software to be able to try to spoof it such that it appears to be in a different location than it actually is. Next slide.

That said, for each of these techniques, there's already a lot of technologies out there to help protect against it. For example, parental education and designing experiences to clearly warn parents the intent of the verification while scanning the face, letting them know that this is allowing children to be able to access certain applications or being able to verify their age is a really, really powerful way to let the parent be aware of exactly what service that they're allowing their kids to get access to.

Proof of ownership, if you're presenting a government ID, binds the individual to the credential that's being presented. So if you're scanning your face on top of the government ID and comparing, and making sure that these are the exact same individual, it helps deter a tremendous number of these attempts. And lastly, risk signals. You can balance a lot of behavioral device network signals to determine that this individual is likely location spoofing. Depending on how you want to handle these situations, you can escalate it and be able to continue to provide some degree of age assurance if you believe that this individual actually is in a different location than where they attest to be.

Speaking of age-appropriate experiences, next slide. We oftentimes find right now that the challenges here are far more challenging. Persona is a bit unique in this space, since we started specializing in fraud prevention, and we borrowed a lot from our background in that to build our age assurance solutions. And what we often find is actually designing age-appropriate experiences presents a lot more fraud challenges than actually designing age gating. From an age-appropriate experience perspective, oftentimes these are adults, who are far more technologically sophisticated, attempting to be children. And as a result, the technologies and approaches they leverage are different as well. For example, oftentimes they'll leverage deepfakes in order to create AI and generate face masks or digital identities to impersonate minors. Account selling is very, very common as well. This means a child who's already verified their current account, selling the account to someone else, such that they can then purport to be a child or a minor on some of these digital services.

And lastly, faking a parental consent, pretending to be the parent of a child who may or may not exist, or have even no parental relationship to the impersonated child. So right now, especially in the world's digital communities, online gaming, this is really, really a huge topic, something that they're all thinking about. Fortunately, for a lot of the partners that we've had the opportunity to work with, there are also a lot of prevention deterrence techniques to prevent these as well. Next slide.

To prevent deepfakes, we oftentimes need to think a lot more about multimodal fraud defenses. This means adopting a multimodal approach to prevent this type of fraud, making sure that you're checking liveness, that you're leveraging additional signals to ensure that the presentation of the face is who they say they are. And of course, testing against third-party assessors. There's a tremendous number of incredible third-party assessment committees who've done a lot of work to ensure that you're building the best fraud deterrence techniques out there.

In order to prevent account selling, continuous verification helps a lot. This means that verifying that the individual is the same owner whenever anomalous behavior is detected, combining the four different approaches that Jim had spoken about earlier in order to alert that maybe additional verification or additional kind of assurance is necessary to ensure that account is not being sold.

And then lastly, parental relationship proof. It is more important than ever to verify that the parent and the child have a relationship, ideally through some form of child in the loop confirmation, to ensure that there's a tie between the two.

Elizabeth Averill:

Thanks so much, Rick. Iain, is there something you wanted to add?

Iain Corby:

No, no, carry on.

Elizabeth Averill:

Okay, great. Let's move on to the next question, which is what privacy risks are associated with different methods, and do some technologies present greater risks? And if we could start with Sarah on that question.

Sarah Scheffler:

Yeah, sure. Hi, can you hear me? I hope so. My name is Sarah Scheffler. I'm an assistant professor at Carnegie Mellon University in CyLab, the Security and Privacy Institute. I study privacy and encryption. I've been researching these topics a lot recently in the context of age and identity verification. So starting with privacy, I want to start by reemphasizing that privacy is really explicitly baked into the policies here. So almost every US state law on age verification has some kind of explicit privacy requirement. Usually this is something like delete the identifying information after access is granted, although they vary. And so I just want to start from the point of, I think we're all in agreement that this is a worthy goal, protecting children, but we're also talking about building an infrastructure for children and adults that poses a lot of really tricky privacy challenges, even when everyone is trying to do the right thing.

And also, it can feel very invasive and disproportionate to users. So if you could go to the next slide, we have a tech report where we return the results of a study where we ... Oh, can we go ... I think this is not my slide.

Elizabeth Averill:

Yeah, we're not quite at that question yet.

Sarah Scheffler:

Oh, gotcha. Okay, that's fine. I'm going to still give a little bit about this here, because I think ... On privacy, I want to talk both about user perceptions of privacy and about underlying technical architectures here. I want to start talking about privacy difficulties seen by users. We ran a web experiment and a follow-up survey where we recruited participants, told them they needed to age verify for the experiment, but in reality, we were just trying to see whether they would actually do the age verification process. This was sort of a fake age verification process set up by us, but it looked real to the users.

And we would basically give them one of seven age verification conditions. So checkbox, ID upload, ID upload plus liveness check, a lot of the things that Iain talked about, and we saw how many users would go through with that process, and then we followed up with a survey to ask them about that, and people had a lot of things to say about privacy. So I guess we'll get back to this later, but first of all, users are overall pretty uncomfortable with these methods. People said a lot of their reasons for discomfort were identity theft, tracking, surveillance, abusive data practices, and people voiced that it really only takes one bad age verifier to ... face data is also already sensitive now. It's likely to be even more sensitive in the future, if biometrics sort of catch on more.

So Iain mentioned a couple of technical architectures that do help ... where you only share age rather than the full ID information, but these still require some kind of trusted wallet or some party that can be trusted by the user to hold the ID and can be trusted by websites to verify the ID.

And of course, there's the cryptographic approach. So I believe Jim mentioned zero-knowledge proofs to prove possession of an ID. So these are both ways to minimize the parties to whom the information is presented, and minimize the information presented to many parties. So you show age and not the full ID contents. These are a big step up compared to methods that don't do this, as far as privacy goes. But in my view, these are really a starting point, so they're really a baseline. There are quite a few other pieces of data that get collected in the process of doing this that aren't necessarily the ID contents.

So first of all, right, a lot of data that's not necessarily explicitly referred to as identifying data in the context of the age verification laws. You've got network metadata, you've got device identifiers, you have fingerprints based on ID information that aren't reversible, but they are still linkable. So these still do sort of de facto build some kind of linkable identifier often, even if the literal name isn't stored. And there are compelling reasons to look at these for stuff like anti-fraud mechanisms and stuff like this. So really, I guess from a privacy perspective, if there's no requirement to avoid collecting this information, then mostly the incentives often weigh in favor of collecting it.

I guess ... users later, the last thing I guess I just want to say is that, in addition to being sensitive, there's also a lot of upstream issues here ... [inaudible 02:40:21] of reasons. Even before these laws, the privacy situation was often seen as pretty grim. People don't want this to add tracking or targeting for advertising. They want to remain secure, and a lot of these things push against that. And we should really take this into account when considering this infrastructure.

Elizabeth Averill:

Thanks, Sarah. Denise, did you want to add to that, to mention any privacy risks associated with different methods?

Denise Tayloe:

Sure. Thanks, Elizabeth. And hi, just a quick introduction. I'm Denise Tayloe, the co-founder and CEO at Privacy Vaults Online. We're known as PRIVO in the marketplace. We are an FTC safe harbor under COPPA, but more importantly, we've been innovating in areas of verifiable parental consent for many years, and we've seen a lot of what Sarah just talked about, how people feel.

I think it's going to come down to trust/ you're ultimately going to need, and I'm sure Sarah's going to touch on this later, but you're going to need for individuals to be able to choose what they're doing and establish that verification upfront, I believe, so that they feel comfortable with it, as opposed to relying on each company that they run into.

So here's just a couple things I want to add. The privacy of the child is really at issue, right? Once you get through the age verification process, if, for instance, it was a situation like Rick put in place, where the child may have been able to create, for even a temporary time, a global fake ID with reusability, then what are the privacy risks that are going to happen? It's all the same stuff that we've talked about in the past. It's the reason that COPPA exists to begin with. So privacy doesn't stop just at the front gate of using some sort of age verification.

Iain's ... No, I'm sorry, not Iain, but Jim's infographic, you note that vouching is in that. And we're not really talking about that too much here, but I think as a parent, take, for example, the school knows who you are, you get the alert when ... verified in the offline world, and help people leverage themselves onto a digital credential that they can use, then we can establish some really great parent-child relationship when it's already in place.

Think about all the times you've sent in your child's birth certificate to sign up for soccer, and dance, and Cub Scouts, and all the different places that you're already verifying those relationships, getting your kids set up in school. So as an example, if we're dealing with an education company, we might, even though they may not be in a position to grant consent on behalf of the parent, because they don't have a contract, or the school isn't demonstrating its FERPA controls, that doesn't mean that the Department of Education for a state couldn't give parents a lift, and overnight help them enable their verified parent credential, and sort of get that behind them. So I think Sarah did a great job of touching on all of the privacy risks, and that's really all I have to add on this question.

Elizabeth Averill:

Thanks a lot, Sarah. And Denise, if we could just move on to the next question, and that's if companies do collect information for age assurance purposes, what conditions or steps should apply to minimize privacy and security risks? And if we could maybe start with Rick.

Denise Tayloe:

[inaudible 02:44:20] this, I want to give a bit of a high level. I mean, Sarah and Denise has touched on this already a ton, but privacy must be considered upfront. As you're choosing through these technologies, as you're thinking about how to balance it all, appropriate data handling usage is absolutely critical. And earlier I talked a lot from a circumvention and fraud prevention perspective. There are mechanisms out there to prevent all the forms of fraud, but what we have to think about all the time is balancing the fraud with the privacy of the children, privacy of the individual, ensuring that the minimum ... age assurance method is appropriate given the demographic of use case, and risk-appropriate as well.

Ensuring that the method is appropriate for the given use case and context, for example, risk of childhood success accessing adult content versus the risk of an adult posing as a minor. If data must be collected, you have to assess the data management policy to ensure that data's being redacted as soon as possible. As Sarah called out earlier, there's a lot of hesitation in terms of adopting any of the technologies in this space, because if one leak happens, especially for sensitive information like biometrics, government IDs, it'd be catastrophic. You can't change your ... stored for any longer than it absolutely necessarily needs to be is critical for any sort of a setup out there.

And lastly, data minimization. Is there a way to collect even less data, minimize the amount of data, and that has to be collected and stored to achieve the use case? There's some really exciting things. Right now, we've discussed already a lot about the idea of double-blind approaches, in which the provider and the website are not sharing any information across. On top of that, there are innovative technologies right now like zero-knowledge proofs, age keys, tokens that allow even less data to be collected while still guaranteeing some degree of assurance that may be appropriate for many of the use cases out there. With that, would love to hand it off to Jim to maybe provide some additional thoughts as well.

Jim Siegl:

Great. Thank you. I really like the if companies collect information part of that question, because I think it acknowledges that it's not a given necessity, and that there could be a scenario where data could be processed securely on the device and not collected, as Iain mentioned in his deep dive on the methods. But if it is collected, I think it's important to examine and mitigate those risks around retention and secondary use. I think it's very useful right now because there's a lot going on globally around age assurance. So one finding from the recent Australian age assurance technology trial I think provides a cautionary example around this. The trial found that a minority of the providers were retaining full biometric or document data beyond what was needed just to verify, apparently, in the report, because they were anticipating future requests from regulators or from law enforcement, even though that there wasn't any such requirement.

So I think this kind of over-retention creates some unnecessary privacy and security risks because of the retained data. So retention policies should be clearly tied to that purpose of age verification with automatic deletion or de-identification, where any audit evidence is required, that it should be clear that those companies, so that companies aren't keeping the data just in case.

As the first panel brought up, many of the laws in this space require immediate or timely deletion of the data. They also think it's useful, in addition to technologies like zero- knowledge proof, to look at international technical standards that speak to this retention question. ISO 27566-1 emphasizes three things. First, purpose limitation, second, data minimization, and third, retention limitation, that personal data should be retained only as long as necessary, something that we see echoes in COPPA's requirements as well.

IEEE 2089.1 looks at this from a systems design standpoint, and encourages that architectures avoid retention by default. I think it's also important to think about both the data and the metadata. So thinking about, does a user's digital footprint disclose things? This was to the mention of the double-blind architectures.

So where that age assurance data is collected from a child, I think we can have conversations about how this relates to COPPA's internal operations exception. Age assurance really fits only when it's confined to that compliance and site functionality, and no secondary use and limited retention. And once you cross those lines, it would seem to fall outside of COPPA. So excessive retention and secondary use are two of the risks that we flag in our updated infographic. All technologies have risks, and there are ways to mitigate, but often not completely eliminate risks.

But it's going to come down to trust. And I think this is an area where, especially in certain high risk uses, audits and certifications may have a useful role.

Rick Song:

And I want to add one last thing as we speak about high risk use cases, because I think that's a great point from Jim's take there, which is that I think as we're designing and thinking about policy for this overall space, it's incredibly important that as we limit the amount of data that's being collected and how we're minimizing the access of it, that we're not also creating vectors to allow impersonation as well.

So privacy's absolutely critical, and there are a lot of technology innovations on the horizon, but I think as of today, all technologies do have certain risks. We need to take a balanced approach that focuses on user consent, immediate data redaction, while ensuring the security of each of these technologies. And this is something that for all the players in this space, I think it's incredibly important that we're continuing to think about, which is the idea of how do we ensure that even as we [inaudible 02:50:42] these technologies, we don't pitch them as completely risk-free, but there's a lot of approaches in terms of how you manage technology that can make it as minimal risk to the individual's privacy and security as possible.

Elizabeth Averill:

Thanks, Rick. Denise, did you want to add something? Denise, did you want to add something?

Denise Tayloe:

Oh, yes. I'm sorry. I didn't hear my own name. That's funny. I just wanted to touch on a couple things, because it's problematic for us. So of course, retention is an issue. We all talk about retention, and we certainly don't want companies or third party providers retaining very sensitive data. But when you flip over to the world of how this is being used with verifiable parental consent, and you have something as strict as no secondary use, then people remain confused on, well, how do I ensure that the same data isn't used over and over and over by a whole classroom of kids to create an account? Or how do I establish a way to communicate new consent requests to the parent, or material changes in the privacy policy that the company has to communicate to the parent who provided consent, making sure that the consent, the second one, is actually the same person who provided the first level of consent.

So there comes a point where you may have to take some amount of data, hash it, have it available as a lookup to enable yourself to see, wait a minute, this same data was used 10 times prior, I'm not going to let it get used again on this service, or retaining some amount of online contact data for the parent to provide the rest of the requirements under the law. So that was one. And then the second one was the impersonation. I mean, let's get real. We need a law that says it is illegal for an adult to pose as a minor for the purpose of committing fraud and causing harm. If every state would pass a law that adults cannot impersonate their children, that would be an interesting thing to have to leverage and help us all.

Elizabeth Averill:

Thanks Denise, I just want to make sure we get through all the questions today. So I think we touched on this earlier, Sarah, but one question I wanted to cue up for both you and Rick is what do we know about user willingness to utilize different options?

Sarah Scheffler:

Yeah, so I'm happy to talk about this one. So if we can go to the next slide, next slide again. So first of all, if you want the more detailed version of all this, we have a tech report out. You can go to the link at the bottom or you can scan the QR code.

So we ran an experiment to see, basically, just how many users are really feeling these particular comfort issues, how many people are being turned away from websites because they don't want to do the age verification. And this was really focused more on the particular kind of age verification that the user would see. So this is about whether the user was facing a checkbox, or an ID upload, or an ID plus video selfie liveness check, or AI facial age estimation, or a couple of other variants, including one email version in which we told users that this email was going to be used to look them up using employer information to verify their age.

So we recruited participants for a study on 18+ romantic web media and told them that they needed to age verify before they could do the study. But in reality, we just wanted to see whether they would do the age verification process for each of these different processes. Can you go to the next slide please?

So we would send users to what looked like an age verification site. On the left here is the ad we use to recruit participants, or one of them. On the right is something sort of along the pipeline. So before we could send them what they thought was the study, we would send

Sarah Scheffler:

Send this, you know, "You must be 18 to show this content," and then we would send them to something that looked like an age verification website, but in reality, was also controlled by the researchers and did not actually collect any personal information. Next slide, please.

So, this is a screenshot of one of the conditions. This is what the study participants would see when going through this study, and just looks like a regular old age verification system, but once again, controlled by us. Okay, next slide please.

So the main takeaway here is that users ended up being... Next slide again, please. Users ended up actually turning away from websites quite a lot. So, the top bar is the checkbox. About 99% of people went through the checkbox, no problem. The next four lines are all variants of an ID condition, so we asked people to upload a picture of their government ID with some different text. The one that says "Gov ID plus liveness" is one where we added a video selfie to the liveness check. And you can see that for each of those, only about 17 to 20, I can't quite read that, 27-ish% of users went through these conditions, and the remaining roughly 80% just quit the study, which either means they hit the exit study button or they timed out, or they closed the tab. The second to last bar is the AI condition, so you can see about 50% of users decided to go through with the facial AI. And the last line is the email condition, where we'd said, "Reenter your email and we'll use it to look up your age information based on financials and employers." Can we go one slide farther, please?

Users also had some interesting stuff to say about effectiveness overall. Sorry, I should say all of these results are coming from a survey that we asked after age verification for users that completed the age verification process. We just sent them straight to the survey. If they closed the tab, we sent them the follow-up survey via email. Users, for the most part, were pretty skeptical of the effectiveness of all of these approaches. So yes, you see very high users broadly agreed that the checkbox was very ineffective, as you can see from the top bar. But also, there was a general perception that even many of the much stronger age verification systems here were not likely to be very effective. So, even the strongest thing that users thought would be the most effective only hit about 40% of people saying it was at least somewhat effective or very effective. Can we go forward two slides, please? There should be two graphs on slide 10.

And the last thing here I want to talk about is comfort. So, even among people who went through and completed the age verification, we saw very high numbers for user discomfort. So with the non-checkbox conditions, on the left side, we see comfort among those who completed age verification, and even that is at least 50% for all of the non-checkbox methods, up to 80% sometimes. Among people who did not complete age verification, a lot of that somewhat discomfort turns into very uncomfortable, and you see the general numbers being closer to 80 or 90% of users who are uncomfortable with this. And people voiced reasons for discomfort that included risk of identity theft, surveillance, tracking, abusive data practices, and once again, this "One weakest link is all it takes," and stuff like this. And if you want to see more information about this, again, you can go to the link at the bottom. And I think if you go a couple more slides, we can put the QR code back up.

Elizabeth Averill:

Sarah, I don't know if that QR code is there, but it'll be in the slides available after the presentation today.

Sarah Scheffler:

Gotcha.

Elizabeth Averill:

Rick, is there anything you wanted to add, in terms of talking about particular user reactions to particular methods?

Rick:

No. I know that we're over time, so I'll keep this brief. What we find is just really two additional findings. I think so much of what Sarah shown here is accurate. The first of which is that oftentimes users will opt for whatever's absolutely the most convenient. For us right now, age estimation via biometrics has become by far, the most, because it doesn't require an individual to find an ID or access. Email, for us, in practice, once it requires confirmation, and using a real email versus maybe a temporary email, in terms of a conversion, dropped pretty significantly from what we see in the real world. I think in practice, that makes a lot of sense. If you are accessing one of these, you probably don't want to unveil something that is tied maybe to your identity. An email is oftentimes perceived by a lot of the users that we've done studies on, to see that they believe that email is actually much more identifying of them than just a biometric.

And the biggest perspective of all of this, I mean, we continue to find that right now, age estimation has become the most dominant form of identity. Government IDs, I think as a supplemental form, creates quite a bit more assurance. We actually find to be very effective from perception perspective, but in actuality, we find it to be probably the most powerful form right now. And then for a lot of the newer forms, like digital IDs and authentication tokens, the coverage is still rather low, but we're seeing those pick up pretty quickly, at least there are early signs of adoption there.

Elizabeth Averill:

Great. Thank you both. We're just going to move on to the final question, and that's can you discuss specific challenges related to estimating or verifying the ages of children? And I'd like to start with Denise, and then hopefully Iain, we'll get to you as well. Thanks.

Denise Tayloe:

Thanks. Okay, just real quick. I mean, obviously, training data is the issue, but I think that we're going to quickly get there with folks like Persona, and others, doing mass verification. I do think the challenge, however, remains kids using somebody other than themselves to get through the process to establish what will become perhaps their more modern global fake ID, so we need to be very careful on that front. Over to you, Iain.

Iain Corby:

Yes. One of the ways to deal with that is authentication, so repeatedly asking people to do checks. It's not a once and done, it's not a lifetime pass. And maybe they had an older sibling or a parent around when they first did it, but that person is not there when you surprise them with a request for it two weeks later. The Chairman began the day by looking at what we might need to clarify around COPPA as well. I think one thing it'd be really helpful to clarify is, which we already have with GDPR, when you're processing data for the purpose of implementing a law, then you have quite broad latitude to use that data. And I think people panic with COPPA, that, for example, if a child has lied and said that they're 15, and then you do an estimation and reveal that they're under 13, maybe that facial age estimation gets you into trouble because you process their data and they turned out to be under 13, and you didn't have consent.

Obviously, you stop and delete everything immediately at that point, or seek that consent, but people get into a great panic about that, and I think it'd be really helpful to be clear about when you can process data for the purposes of complying.

Elizabeth Averill:

Denise, did you want to add to that?

Denise Tayloe:

Yeah, I just wanted to say that when you're dealing with services that are primarily directed to children, we're already assuming that they're children at that state, so we probably shouldn't launch into age estimation. So, it's going to depend on what is your audience type on different services. But my last point here is going back to Jim's infographic, is I think that we need to provide parents with a way to assert their child's age and make that available through, not just tokens that the kid can choose to use, because he's not going to use his 10-year-old token when he's trying to fib that they're 15, but to also allow, like in the fintech world, lookup service that would allow the parent to say, "No, this phone number, this data is absolutely associated to a minor and that signal needs to be respected." I think you're going to see some of that in the global privacy control type format, not a broadcast that this is a child, but more an approach of this device needs an age-appropriate experience and you need to default to child until you can get a better answer. Thank you.

Elizabeth Averill:

Great. And Iain, did you want to add to that?

Iain Corby:

Yeah, just briefly to add, I think there is a distinction here between when we just accept the parent's word for the child's age and when services need to get an independent verification of that age. We do know, this was mentioned earlier, that often, parents are complicit in helping their kids to access services which are age-limited when they shouldn't be accessing those services. So, sometimes you will need to do an independent age verification rather than simply relying on a parental attestation. So, it's sort of one step up from self-declaration, but it's not an independent view of the age of that user. Thanks.

Denise Tayloe:

And certainly, we shouldn't assume that all parents are liars or fibbers, but they will. And technically, they don't understand that when they do it in one place that they feel they should have more latitude, because they're being denied the chance to actually provide consent to let their child onto a service they think is reasonable, they're not understanding that that then creates a global capability for the child to use in many, many places. So, perhaps when we're talking about creating federation or re-usability, we need to take a step farther than what some of the processes are to ensure that parents really understand the global nature of these credentials.

Elizabeth Averill:

And Denise, I think you might've raised this point earlier, that when you're asking children to provide information as part of an age-assurance step, there are concerns about their ability to consent to sharing biometric data and other information.

Denise Tayloe:

Yeah. I mean, I think the representative earlier in the morning, stated the elephant in the room, kids in the United States, 17 and under, can't sign contracts, and terms of service are contracts. So just at that level, we've got issues. But certainly, you've got how is a kid to know that one service requesting its biometrics is an actual real service doing that and not a link they've clicked through their Facebook account, or some other social media account that they have, to a site that they're spoofing information from them. They're not really equipped to make consent decisions. We've already agreed to that. And also, this sort of conflating, this isn't about parental consent, but parental vouching, when you're talking about age verification.

So like I said, the controversial thing about laws for impersonation, and I know people don't want to hear this, but it would not be legal for a parent to take their child to the DMV with a fake birth certificate, so why is it legal for parents to fake their child's age for a digital driver's license on the internet? So, I think that we have to take a step back and think about parents can't just be left on the sideline here and assumed to be ignorant to all of this, and they do have some responsibility to at least not lie on behalf of their children if it's going to have a global impact.

Elizabeth Averill:

I want to thank this amazing panel of experts for sharing your expertise and experience with us. I feel really fortunate that we've had the opportunity to learn from you today. And I'm happy to announce that we're now going to take a lunch break, so I hope everyone gets a chance to walk around and eat something. We'll be returning at 1:30. Thank you.

MUSIC:

(Instrumental music).

Manmeet Dhindsa:

Hi, welcome back. My name is Manmeet Dhindsa and I'm an attorney in the FTC's Division of Privacy and Identity Protection. We're going to begin the third panel of the day, titled Navigating the Regulatory Maze of Age Verification. As we've discussed on a number of occasions today, there has been a flurry of activity in the United States, particularly on the state level, in which it's becoming increasingly commonplace to see laws requiring the use of age verification.

We have a panel of four experts today who are going to help us understand the purpose of these laws and how the different legal frameworks can work together to best protect children online. I'll provide a brief introduction of our panelists, but I would direct you to the FTC website to see their full bios.

First, we have Katie Hass, director of the Consumer Protection Division in Utah's Department of Commerce. As we'll discuss more in a bit, Utah in particular has been very active in passing laws that require the use of age verification. Second, we have Jennifer Huddleston, senior fellow of technology policy at the Cato Institute. Next, we have Sara Kloek, vice president of education and youth policy at the Software and Information Industry Association. And last but not least, we have Clare Morell, fellow in the Bioethics, Technology, and Human Flourishing Program at the Ethics and Public Policy Center. Thank you all so much for being with us today. We have a lot to discuss, so let's jump in.

So I want to start by building on the first panel and kind of digging a bit deeper on the purpose of these laws. And in particular, I want to delve into the question of whether we need laws requiring the use of age verification in order to protect kids in today's online world.

And Katie, I want to start with you. As I mentioned, Utah has enacted various laws that require the use of age verification for different purposes. So for example, one of Utah's laws would prohibit kids from accessing pornography sites. Another would empower parental involvement with kids' engagement with social media platforms, and yet another would allow parents to take actions to protect their children from certain types of apps or in-app purchases. So my question for you is, why are these laws necessary, and do you think these types of regulations are effective at protecting kids online?

Katie Hass:

So yes, they're necessary, and yes, I do think they're effective. I mean, a lot of these are relatively new, so we don't have the impact studies to show how they've helped kids. But clearly, as Chairman Ferguson said earlier today and everybody else, we want our kids to flourish. Utah, I think, per capita has the most children, and we care greatly, as many states do, about our children and making sure that we create an environment in which they flourish.

As many people have talked about today, we live in a technical world. There's technology everywhere. I have four kids. Most of them need their phones to log into various apps at school, to do homework assignments, to engage with communities, their sports teams, et cetera. And so because of that, it's also handing a device to your child that is the Wild West. And while I wouldn't let my kid go to a bar and I wouldn't let my kid or want my kid to go to a liquor store and be able to purchase alcohol, I think it's a valid point that I do let my kid out into the world.

And so there is some sort of obligation of these companies that are out there on the web to make sure that they're protecting youth who are coming onto their platforms, in some cases, a full ban with pornography, and in other places, tailoring an experience that is appropriate for children. And in some cases, depending on the age of that child, it's appropriate for parents to be engaging with those kids on those platforms that have an understanding of what they're doing.

And so all of the laws that Utah is passing, at their heart, are to protect children from harms that we are seeing, whether it's social media with their algorithms, whether it's now AI platforms and chatbots, whether it's downloading harmful material on the app store, and most especially also pornography and making sure that the porn companies are doing what they need to do to age verify before they allow young people... Well, to basically prevent young people from being on those platforms.

And so all of these laws, I think, are reasonably tailored to get at the heart of the issue, which is to protect our children and to make sure that the experiences that they're having online are appropriately tailored for children. Did I answer both your questions?

Manmeet Dhindsa:

You did. Thank you very much.

Katie Hass:

Okay.

Manmeet Dhindsa:

And so I want to turn to Jennifer on this point as well. So Jennifer, in some of your writings, you note that regulations that require the use of age verification can actually pose more problems for kids than they solve. So I imagine that you're not necessarily going to agree with Katie's response here. So what problems do you think these types of laws can pose?

Jennifer Huddleston:

Thank you. And thank you to the FTC for arranging this workshop on what is a very important and very timely topic. One of the key risks when we're thinking about these laws is that they're a policy solution that often results in a one-size-fits-all solution when every child and every platform is going to be unique. And this may mean that you're preventing young people from accessing what could be beneficial elements of the internet, things for educational services or other things that could further their career or their educational needs, or even just opportunities to connect on a particular interest, particularly for young people who may feel isolated.

But I think one of the key concerns, and we heard about this a bit in the previous panel, are concerns around data privacy and the data privacy of young users who these laws are intended to protect. Oftentimes, these laws will require additional data collection, which could create a kind of honeypot for bad actors to know where all the young people's information is. And we know from looking at what's happened when these policies have been introduced in other places, like in the UK where we've seen, after the Online Safety Act, pretty significant hacks of biometric information or other age verification information such as on the platform Discord, that these can and do happen.

And so many parents or young people may be concerned about the privacy risk involved in that information being out there. We also know that young people, just like adults, can be subject to identity theft and that sometimes they don't even know that's occurred until they're much older. And again, additional data collection of what can be very sensitive information that is required to verify identity or age under these procedures can certainly increase that risk as things go along.

So when we're looking at this from a policy perspective, and we can recognize that every child and every family may be slightly different. To Katie's example of letting her kids out into the real world, a lot of parents probably have different ages that they're comfortable allowing a child to do different elements in the real world, even within the same family. And the same can be true in the online world. And that's why parents, not policymakers, are often the best decision-makers when it comes to when it's appropriate for their child to have certain online experiences. And in approaching that from a policy point of view, we should look at ways to educate and empower both parents and young people to ensure that they're able to have more positive experiences online, and so that they know what to do, how to go to that trusted adult or to that platform should they encounter a negative experience online.

Manmeet Dhindsa:

Thanks, Jennifer. If I could ask you a quick follow-up on that. So you mentioned that we should think about how we can empower parents and young people. Do you have specific recommendations or ideas on how we can empower parents and kids in this space?

Jennifer Huddleston:

So I think there are a variety of things. There are things that civil society groups and platforms themselves are already doing to make sure that parents know what parental controls are available on their platforms, and so that civil society groups are modeling for parents how to have what may be unfamiliar conversations with their young people around how to approach new technologies or around what their family's values are regarding different technologies.

We've also seen, in some states, pushes to update digital literacy curriculums. Many schools or many states already have some degree of digital or media or technology requirements in their schools. And so updating that based on our 21st century needs to cover things like social media, or even now as we approach artificial intelligence, so that young people are aware of how to use these tools in responsible and beneficial ways.

Manmeet Dhindsa:

Thanks, Jennifer. That's really helpful. And I want to turn to Clare because I think that this discussion about parents is particularly relevant to you and the recent book that you wrote that discusses the role of parents in protecting kids in today's technological society. So Jennifer raised the idea that parents can be this tool in protecting kids online, so these laws that require the use of age verification are not necessary to protect today's kids. Do you agree with that position?

Clare Morell:

I do not. I actually think that age verification laws are incredibly empowering to parents. And most of these state social media laws that have been passed that require age verification are actually all tied to parental consent, because the point is that parents are actually not in the driver's seat right now of their children getting access to social media. Because right now, a child can enter a birthdate, there's no verification process, so they can falsify their age and then click a box agreeing to all these terms of services with these giant companies, and a parent is not involved in that process whatsoever. And so I do think a lot of these age verification laws, at the root and the heart of these laws, is both to protect children, but also to empower parents to have more oversight to be able to protect their kids. And I just want to briefly address, one of the other reasons these age verification laws are so needed and that they are empowering of parents is they provide collective solutions to problems that are too large and complex for individual parents to address on their own. So we all know, it's well-documented now, that social media does not just harm the individual, but it actually changes the group's social dynamic. And any parent can see this firsthand, that social media just changes the way kids interact. It changes the entire social environment.

And so even just these laws like we've seen in Australia... That are not a safe environment for children that have addictive properties and effects. And the same way... It's applying that same logic to the online world. And it empowers parents because then they're not fighting these individual battles. And even for those parents who opt out of social media for their kids, they're still not able to protect them on their own from those group level dynamics.

And it's the same with pornography. The porn industry has said, "Well, parents can just install filters." Well, in the smartphone social media era, if I install every filter possible on all my child's devices but then they get on the school bus and any other kid can lean over and pull up pornography websites on their smartphone, well, that's a collective harm. That's... Verification laws are talking about the internet as a whole, but it's recognizing there are parts of the internet... Content they expose children to that should be age-restricted. And those age restrictions are actually meant to empower individual parents, particularly when there are these collective aspects... [inaudible 04:37:32] to what parents want, but actually empowering parents and helping them out, the government providing critical backup.

And I will just last end with this, that in the real... They're not getting into bars, they're not going into strip clubs, they're not purchasing alcohol and tobacco. And I think we have to recognize that there are portions of the online world, as there are parts of the physical world, that are just not safe environments for children to be in. And so we have to be able to then distinguish between adults and children online, and that is really the heart behind the age verification laws.

Manmeet Dhindsa:

Thanks, Clare. That's really helpful. Sara, I think you wanted to jump in?

Sara Kloek:

Yeah, I wanted to talk a little bit about... And I know we're focusing on age verification here, but what we're talking about, what happens after we verify the age? Are we restricting access to content, and content being, there's pornography and then there's news sites and then there's sports websites. What sort of stuff, after we verify the age, are we making sure is appropriate for kids and not appropriate for kids? And how are we deciding that, I think is important to think about as we're having this conversation, because yes or no, age verification is good or age... And I know we've talked about pornography, but there's other things that kids are accessing online, whether it's the Washington Nationals sports team scores or... You can tell I'm not a huge sports fan by what I just said. But I think it's important to recognize that context, that we need to make sure that we're talking about what Mark Smith was saying, from CIPL, earlier, risk-based approach. When is age verification necessary? When is it not necessary? Are we having age verification on a connected fridge? Where are we having that sort of level of age verification to make sure kids are still being able to access information online and access information without having all sorts of data collected about them?

Manmeet Dhindsa:

Thanks, Sara. So you mentioned pornography sites as a place where we've seen these age verification mechanisms implemented, and I think it sounds like you support the use of age verification on those properties. Are there other places in your opinion that we should implement age verification mechanisms?

Sara Kloek:

I'm going to give an “it depends” answer, because I think, like people have been saying, parents or caregivers are really the ones that can decide what's appropriate for their kids at a time. I know you can send your kids with a permission slip to see a R-rated movie if they're old enough. Is that the right approach for what we're thinking about accessing things on the internet? Maybe, maybe... and groups, that's the direction that some want to go and we need to make sure we avoid the unintended consequences of, for example, collecting more data from people that don't... report. I think there's a lot of questions to be answered, so I'm just going to give the “it depends” answer.

Manmeet Dhindsa:

Jennifer, it looks like you wanted to jump in.

Jennifer Huddleston:

Yeah. I think this goes to part of the reason that there's so many concerns about the potential speech impact of these laws, as well as the impact that these laws could have well beyond just kind of simple social media platforms or things like that. I know Clare said no one's trying to age gate the entire internet, but when you look at how broad some of the definitions are, I think they apply much more generally than a lot of people would think about. When you think about a platform, for example, like YouTube and how that's become a very complicated issue when we look at the Australian law, for example, and whether or not it's better or worse for a young person to be able to access YouTube without an account versus with an account.

When we think about things like the UK's Online Safety Act, which has very broad terms when it comes to what it considers harmful content for minors. And what we saw is in response to that, many platforms out of fear of potential non-compliance with the regulation or fear of the consequences of the regulation, were age gating information with regards to the wars in Gaza and Ukraine, were age gating other information that many people would find that there could be beneficial usage for, whether it's for newsworthiness or for other knowledge.

When it comes to many of these issues, we're going to very quickly move from pornography to what are a lot of gray areas where individual parents may disagree or where different states may disagree with one another, let alone the fact that when it comes to what is the actual underlying concern for young people on the internet. For some, it's exposure to certain types of content. For others, it's the amount of time their young person may be spending online. For others, it may be very particular when it comes to something that they feel goes against their family values or things like that. For others, it may be traditional kind of concerns about who's contacting their child. And again, that one size fits all approach isn't necessarily going to solve what is a much more nuanced problem for each individual family.

Manmeet Dhindsa:

Thanks all. There's clearly quite a bit to unpack in this space. So maybe we can move on to the next bucket of issues that I want to discuss, which is digging a bit deeper into the regulatory fragmentation as I think Mark put it in the first panel. So as we look at the different legal frameworks in the states and even internationally, it doesn't seem like there's broad consensus on the best approach to verifying the age of users, whether that's when to use this technology, the types of harms to address, or the type of technology to be used.

So I want to use this panel to dig into the question of what the ideal law in this space looks like. And I want to start with you, Sarah. So at SIIA, you represent a variety of different industry players. So you have an interesting perspective, I think, on some of the on the ground work that's occurring to implement age verification in online platforms.

We've seen various state laws pop up that require different entities to implement age verification. For example, some laws require individual websites or apps to verify the age of its users while other recent laws require app stores to conduct that age verification. In your opinion, who's best positioned to conduct age verification? And don't feel limited to the two entities that I just mentioned. I know there are a ton of other players in this ecosystem that have touchpoints with kids. So it begs a question of whether this age verification should be also done at a higher level, maybe the device level or the operating system levels. So what are your thoughts on that?

Sara Kloek:

Thanks for that. And I forgot to mention earlier that I wanted to give a shout out to all the parents that are joining today and watching this workshop that are listening to people talk about how to navigate the online world, especially for those folks that have kids at home in the DC area or are snowed in with their kids. So my kid is at a play date, so you aren't going to hear any Frozen lyrics, but I'm sure you can play Let It Go afterwards if you're missing out on that. But I know everyone's doing some hard work today.

So I think to your question, it's a big question that I think a lot of the panelists on the next panel have deep opinions about. We have some deep divides within our membership on the right approach for this as well. I know that some support one method, others support other methods, some may have opportunities to provide age verification in the future. So we've been having these discussions internally, but really we have come to the big conclusion thus far that everyone in the ecosystem is going to have a role to play when it comes to protecting kids online and also when it comes to age verification, age assurance, age estimation.

We also think that includes the Federal Trade Commission and Congress. I think it's really important for Congress to pass a comprehensive federal privacy law because we should not be passing age verification laws without thoughts on how to protect the privacy of users. I think that's really important. Some states that are passing age verification requirements do have privacy laws, so that is one positive aspect, but we do think federally it's really important to have a comprehensive privacy law to ensure those privacy protections and guardrails are up for... that we can institute guardrails for tech companies and other companies that operate online to make sure that there is an ecosystem to protect the users when using age assurance mechanisms.

I think that this workshop and other discussions on this topic, perhaps even a 6B investigation on age verification methods could be helpful to learn more about the process and accuracy and unintended consequences or bias. I think that it's really important to make sure that there's a general consensus about the right path and what needs to be done in a way that respects the rights of Americans, including kids' constitutional rights. And I think it's also important that whatever age verification process, whoever is doing it is done using a risk-based approach.

Manmeet Dhindsa:

So if I understood you correctly, it sounds like you think that there's maybe a benefit to having age verification done on different levels. So maybe it's happening at a higher level and then done later at the individual app or website. If I understood you correctly, and correct me if I'm wrong, but if I understood you correctly, do you have any opinions on what to do when the different entities come away with different results? For example, if at the device level it says that the user is 15 and at the individual app level, it says that the user is 12?

Sara Kloek:

I think that's a really good question and a hard one to answer because we've seen that, I think, and my fellow panelists can correct me if I'm wrong. I know that there was some rulemaking happening in New York that was saying that for kids under the age of 13, the accuracy rate needed to be, I think it was like 0.1%, but earlier a panelist was talking about, I think it was the hand method was a 2% accuracy or inaccurate rate. So I think that that's a hard question to ask. Are we going to then want to get social security numbers from kids? I don't really want to be sharing my kid's social security number out there with anyone really to make sure that what her age is. And I think that parents should have a right to not share information if they don't want to. And I think that's a tenet of COPPA of being able to say, "No, I don't want to share that information."

Manmeet Dhindsa:

So that raises another follow-up question that I have for you, and I'll direct it to you, but if anyone else on the panel also has opinions, please feel free to jump in. So one of the things that was discussed, I believe, in the second panel is that a number of these age verification or age assurance methods are subject to circumvention by kids. So for example, they can just hold the phone up to a parent's face and pretend it's for something else. Do you think that there is a particular place in the ecosystem where we can avoid... If we conduct age verification at that particular point in the ecosystem, can we avoid some of these circumvention issues?

Sara Kloek:

I think that kids are really smart and I think that it's really hard to figure out how to... I think... I don't know, kids are smart and it's an opportunity for a discussion amongst the family. But also the FTC has long said that yes, kids are going to lie at the age gate, so they recognize that kids might not be telling the truth at an age gate. And if kids aren't going to tell the truth or figure out how to circumvent something, what should companies be doing? Is the requirement then to collect more information on kids to make sure that kids aren't lying? Is the requirement going to be to get the parents to provide more information? I just think getting more information from people to address a problem that is just developmentally how kids might be acting at that certain point, whether it's not listening to parents or trying to navigate the world on their own, figuring out how to legislate something that kids might do anyways is hard.

Manmeet Dhindsa:

Yeah. Those are really helpful considerations to think about in this space. So Clare, I want to turn to you. If you could create your ideal law here, what types of things would you recommend that policymakers keep in mind when drafting laws that require the use of age verification?

Clare Morell:

Yes. And bear with me because this will probably be the longest answer I give for this whole panel because I think there's a lot of factors that go into crafting the ideal age verification law. My brain likes to work in buckets, so I'm just going to give you my five buckets upfront, and then I'll just briefly explain each. So I think the first is just how you define the platform. The second is what age you've set for the age gate. The third is the methods, the process of age verification, how that is specified. The fourth is privacy protections for the data collected. And then the fifth is enforcement, what is actually enforcing the law. So just each of those briefly in turn, and I actually hope some of this will help respond to some of the points that Sarah was raising about some of the concerns around collecting information from children and things like that.

So the first is just, I think, and this was brought up earlier, we're not trying to have these be age gates for the entire internet. And I think especially, I know Jennifer, you mentioned maybe a UK definition being expansive. Just talking about the US context, all the laws that I've been looking at are quite narrow. In fact, a lot of them for social media platforms have size caps. And so there's specifically targeting certain platforms. And so anyway, how you define that, how you define what it classifies as a pornography website or a social media platform, or now there are measures to look at age gating AI companions. So how you define that will be very important. And again, I am really for narrow kind of targeted definitions that make clear this portion of the internet we have deemed to be unsafe for children, and so we're going to age gate that. It's not meant to be age gating the internet as a whole. So I think specificity, narrow targeted definitions is important.

The second is the age that you set. So I think actually I'm in favor of higher ages. So 18 for pornography, I would say 16 or 18 for social media. And part of that is we actually have procedures and processes for determining an individual's age around the age of 16 or 18 or 21. Most states have issued driver's license by age 16. And so I don't think... A lot of us are not comfortable with the idea of companies collecting lots of biometric information on children to try to determine who's a child or not. But we're just depending on where you set the age, that can be a clearer marker or not. I think it's more difficult to determine what's a 12-year-old from a 13-year-old. And so I'm in favor of actually higher age limits because I think there's actually less data then that needs to be collected to determine that level of age gating. So that is an important consideration to keep in mind throughout all these laws is where are we setting this age and how easy or not is it then to determine a user's age at that particular set point?

The third is the methods. And I think I am really for a both and approach, methods that are both effective at keeping kids off, but that also really protect and maximize user privacy and choice and freedom. And so even the Australia example makes really explicit that the platforms cannot rely only on document-based options. So if an individual like an adult doesn't want to upload a government ID, there have to be other options available to them. So I think trying to actually write the laws in such a way that maximizes user choice in how they're going to verify age.

And the second piece of that method is we are talking about these circumvention issues when there's this kind of one age check on the front end. I like laws like Florida's and Australia that have a kind of ongoing responsibility on the part of platforms and they expect them to be employing kind of multiple means for age restricting. So an age check on the front end, but then there's also an ongoing responsibility on the part of platforms. Given the data they collect and the behavioral analytics they run, they understand how people interact online and they're able then to determine, okay, this person, maybe they flash the selfie to their parent, but they're really interacting like a minor user. And let me just say, the companies, they know who the users are. I mean, their business model is targeted advertising. And so over time by someone using an account, they can develop a very good sense of their age. And if they were to wrongly flag an account as a minor and try to remove it, then there is a whole process of then you can verify that you're actually an adult.

But I do think a combination approach where there is some check on the front end, so minors just can't create accounts, but then also some ongoing responsibility on the part of the platform to be employing multiple methods because I think it should be clear in how the law's written that a platform is not necessarily completely off the hook just because they ran an age checkpoint if that age checkpoint is really ineffective. And so I think those kinds of approaches in a law that kind of combines a front end check with some type of level of ongoing responsibility ends up making sure that they're effective because the company is then incentivized to make sure it's effective. So I've been really following how Australia has been doing this and I think they've been thinking through that really well.

The fourth thing is privacy protections. I think an age verification law needs to have really strong privacy protections for all the reasons that we've mentioned. We do not want to compromise, especially the user privacy of adults. And so I think ensuring in the way a lot of these laws are written, that they can only use the information collected for age verification purposes, so they can't just hold onto it and use it for other business reasons. And then they need to immediately delete it and not retain it in order to ensure that, like Jennifer was warning about, that there's not some big honeypot out there of all this information collected that someone could hack into, but that information is not retained by the provider or by the platform.

And I'm really in favor of a lot of methods that actually kind of have a two-step process, or they call it a double-blind method where no information really about me is transmitted to the platform, but only whether or not I'm above or below a certain age. And so I think innovations like AgeKey for any of you who have followed that, just encourage me that really the US can be the leader in age verification. We have some of the most amazing companies in the world with incredible technological prowess that can really put their minds together and come up with good solutions.

And so the private industry developed this open age initiative where you can just create an age key. You verify your age once on a website and then your age key, this kind of cryptographic signal like a pass key, I think a lot of us have pass keys for our various accounts is then stored on the device, but then can only be transmitted to a website when that specific user who owns the age key unlocks it with a biometric, whether it's your fingerprint on your device or your face scan on your phone. And so even a family's shared device could actually have multiple age keys associated with different users. And so this is really innovative and it's very much user controlled and very privacy protecting because the only information then that a site or platform gets if you're accessing something age restricted is just a signal, yes, they're below or above the age of 18.

And so I would just say, I think the message I would just want everyone to take away is the technological means are there. Honestly, the private industry is coming up with great solutions, but we need laws and regulations on top of that infrastructure that they're creating to ensure this kind of compliance across the board. And right now, this AgeKey initiative is voluntary, but we would want websites to accept the age key and we would want websites to have to verify age if we want that portion of the internet to be age restricted. And so laws are really necessary, but the technological means exist to do verification in a way that's both privacy protecting and effective at keeping minors off.

And so the last part of the laws is just enforcement. I think that's really critical, whether it's the Federal Trade Commission or state level consumer protection divisions like Katie from Utah, we definitely need to make sure that however the laws are written, that they are enforceable so that they actually have teeth so that the companies will comply or otherwise face significant penalties or fines. And so there needs to be a significant enough of a kind of penalty or threat or fine really to compel the industry to do the right thing. And I'll end there because that was like a long entourage, but those are my kind of five key components and how I've personally thought through what the ideal age verification law would look like.

Manmeet Dhindsa:

Thanks, Clare. There was so much interesting information in there, and I wish we had more than 40 minutes to really dig into a lot of those pieces, but I did want to follow up at least on one point that you raised, which is that you think that companies should have this ongoing obligation to determine the user's age. And you said that the age assurance mechanism should be conducted at the outset and then as the user continues to interact, they're continually checking to see the age of the user. At what point in that process can we say, "Okay, company, you've done enough of these ongoing checks to understand the age of the user to a sufficient degree"?

Clare Morell:

No, it's a really good question that you raised. And what I would say too, as I should have clarified, I don't think the expectation is 100% perfection That's not realistic for companies, but I think how Australia is doing it is it's more about ensuring that there are practices and processes in place that sufficiently the company has taken steps whereby they will be identifying and removing any minor accounts that are detected. So I think it's more that there's some kind of proof that the company has a process in place and not that the process is 100% perfect. And that's the other thing I did want to say is I think we also have to have a very realistic expectation of what these laws can accomplish.

And I think sometimes people are kind of holding this to the standard of perfection, but that's not the case in the real world. I think we know that if a kid is really motivated to get access to alcohol or tobacco, they'll find a way, they'll find a friend, or they'll coax a parent into doing it. And I think it's the same in the virtual world that there will be children who circumvent these laws and that they probably won't get caught by the company, but it is important that it's setting a norm. And I think we can't lose sight of that fundamental purpose of the law, that it really helps set norms because the vast majority are not going to be trying to get VPNs or use their parents' biometric scan. But to that point, there will be some. And so trying to have some type of process or procedure in place by the companies where they can then identify minors who may have circumvented their age verification check, I think is important. And I think that's something that, again, a government agency like the FTC is really the best positioned to be determining. And so I think a lot of these laws could actually leave it to the division of consumer protection in their state, like Utah or the FTC, to actually be the one on the frontline regulating and updating those regulations. Because another important point I didn't mention is these laws, we need to make sure also that the methods being kind of used are evidence- based and adaptable, that as new developments come about in age verification technology, that the law is not kind of static and stuck in time, but is actually adopting those better methods as they become available like zero knowledge proofs or this age key that I mentioned.

And so I think it is important to have a regulatory body or agency that's the one that's kind of writing those guidelines for companies because this space might be changing and evolving so quickly so the companies know what the expectation is for them in order to ensure they're complying with the law. I don't know if that answered your question, but that's my thinking on it.

Manmeet Dhindsa:

No, that was really helpful. Thank you. And Katie, did I see that you wanted to jump in there?

Katie Hass:

Yeah, I just wanted to add, to Clare's point, first of all, I think there's just so much out there with the IEEE standards that we can use and utilize in creating these laws. We relied on them and their standards in making sure that companies had an option to use the privacy methods that they felt were appropriate or the methods to determine whether or not somebody was a minor that was going to be most comfortable for their clients. So rather than prescribing, you have to use a specific method, we gave companies wide latitude to just get to an accurate standard. And whether that was through a funnel or whether that was through using an ID, that was up to the company themselves.

The goal was to get them to a certain level of accuracy in the process. And I think laws like that don't stifle innovation that allow for companies to know their clients, know who their clients are, what they're going to be comfortable with. Because you're right, people in Utah, for example, might not be really comfortable with sharing an ID and they might want to go with some other type of method to get there, whereas somebody in another state might not mind sharing their ID. And so we wanted to make sure that we weren't stifling any innovation, that we were allowing technology to grow. And I really appreciated Clare's thoughts on that.

I also just wanted to comment that I feel like we're ignoring the fact that most of these companies are data collection companies, they're data mining. They already know a ton about us and who our children are, who our families are, our connections, our friends, the worlds that we live in. And so I think it's kind of rich that we're sitting here concerned that we're asking them to age verify when they have the technology, many of them to know in a rough estimate who kids are. And if you look at social media, for example, and the way kids are posting about themselves, you can get a relative age bucket about that and tell that somebody is maybe 14 and not 22 when they originally lied to come in. And so I think somebody on an earlier panel mentioned AI to kind of do that ongoing monitoring that we may need in order to trust but verify as people are coming into these platforms and when something seems suspicious, putting some of that burden back onto the companies to figure that out.

And finally, I just wanted to throw in there that some of those parental controls, which I'm glad to see, many of the social media companies in particular starting to add parental controls, are the result of laws being passed and lawsuits against the companies, that they're trying to shift and move and get ahead of what they already knew was coming. So I don't think that we would have these protections that we are, and I'm going to use protections loosely there, but some of those parental controls that we're starting to see on these platforms, but for the lawsuits from the states, from the FTC's inquisitions into them under the 6B authority, and also the laws that are being passed across the country. They see the writing on the wall, which is making them now move towards giving parents more insight into their children's accounts, et cetera. So I do think that to say that laws aren't necessary and that these companies are doing it on their own is not accurate. I think the companies are doing it because these laws and because of the lawsuits that the states have filed.

Manmeet Dhindsa:

Okay, thanks. Sara?

Sara Kloek:

Yeah. I would say that just because companies have data that they can use to identify users doesn't mean they should. If it goes against their privacy policy and they go and identify who is related to somebody or what is being done, the familial relationships doesn't mean a company should be doing it. And again, especially if they say they aren't. So just because they're collecting the data, we should not be saying all companies should be doing it.

Katie Hass:

Sara, I'm actually going to give you that one, and that's why in Utah we made sure that we also included a lot of data minimization and other things around our age assurance. We didn't just say they have to do age assurance. We said, no, no, no, no. And with that, here are all the things you have to do to restrict the use of that information. So I do agree with you on that. I'm just saying they also though, through the natural course of their business, are gathering information that could send signals that somebody is also not of the age that they originally said that they are. So I see what you're saying though, and I'll concede that one.

Sara Kloek:

Thanks.

Manmeet Dhindsa:

All right. Well, thank you. Thank you all for jumping in on that bucket. I'm looking at the time. I want to jump to the next bucket of issues that I wanted to raise today. So far we've focused a lot on state laws, but there are other legal frameworks that come into play here such as COPPA, as Chairman Ferguson mentioned this morning, as well as the First Amendment. So I now want to turn to a discussion about how we can get these different legal frameworks to work together so we can maximize protection of kids online. And since we are at the FTC, I think it makes sense to start by focusing on one of the Commission's primary tools to protect kids online, and that's the Children's Online Privacy Protection Act or COPPA. And in case anyone online is not familiar with COPPA, I'll start by giving a very, very high level description of the law.

So under COPPA and the FTC's implementing regulation, the COPPA Rule, operators of commercial websites and online services directed to children or those with actual knowledge they're collecting personal information from a child are required to provide certain privacy and security protections to personal information collected from children under 13. So for example, some of these requirements include requiring operators to provide parents notice with information about their information practices and obtain verifiable parental consent before these operators collect, use, or disclose personal information collected from a child.

So with that background, I want to turn to how COPPA relates to the recent laws we've seen that require the use of age verification. And I want to start with Katie. Katie, what do you see as the interplay between COPPA and recent state laws requiring use of age verification? And in particular, I would be interested in understanding if there are any places where you think COPPA interferes with what states are trying to do with age verification.

Katie Hass:

That's a great question. So I want to back up and just say our law in Utah was 18 and under for social media in particular. And then when you're under 18, there were certain things that parents get to have control over. There are default settings and you need a parent to override it. So as far as COPPA's concerned in the 13 to 18 bucket, which it doesn't really apply, I want to say we actually relied on COPPA though for parental consent. We said if you're complying with COPPA or you use the same methodologies that you would for COPPA for somebody under 13, then that is enough to say that that is a parent or guardian in relation to the child. So we are really relying on the FTC when it comes to that parental or guardian relationship and establishing that for right now.

And therefore allowing, I think one of the things I would ask for is some fluidity so that... Rulemaking takes time. And so to go back to Clare's point and maybe even Sara's point, making sure that we're allowing for innovation and quick innovation in this space. If we move into more of a key access type situation, maybe those key accesses are tied to a parent and then the parent is able to give the access that the child needs and able at the same time to verify. Those kind of technologies I think are going to grow exponentially and quickly. And the ability for COPPA and specifically parental consent to pivot and to move with that rapidly would be really important for states like Utah who would simply rather refer to the FTC, and frankly, I think companies would rather have, and not that I want to advocate for companies, but the checkerboard's hard. And so knowing that you're complying with COPPA means, or at least a standard set in COPPA, therefore you can comply with the state's laws is an elegant solution to kind of say that.

And then one of the things we hear from the companies all the time when we talk about age verification or age assurance is, well, we have to comply with COPPA and that requires us to do a neutral age gate. I understand that. A neutral age gate to start is fine, but most of the state laws are not going to settle with just a self-declaration as even a starting point for age assurance. It's going to have to be something more than that. And so making sure that companies feel like they can both comply with COPPA because it's neutral at the outset, but then build upon that and making sure there's no conflict there, no preemption or anything like that that allows the states to add the plus levels that they need in order to ensure that their laws can be enforced, I think is really important.

So down the road as you're looking at this, when people start asking for preemption, it would just be good to keep that in mind, that the neutral age gate only is not sufficient and I think you guys already know that. So those are kind of like my initial thoughts on this, but I really feel like as Clare was talking earlier about the key access, and I know the state of Utah is looking at state endorsed digital identity and moving towards some sort of both ID, but then also something that links parents accounts with kids so that parents can get that consent going. I think those are the types of solutions that we're really going to need going forward, both for COPPA, but also for all of these laws that are being passed.

Manmeet Dhindsa:

Thanks, Katie. So I want to follow up on a few points that you raised there. So first, you mentioned that self-declaration is not enough as a neutral age gate. So do you have an opinion on what type of age assurance mechanism would be enough in your view?

Katie Hass:

I think it really goes back to the risks associated with the account. If you're talking about under 13 and we want... If the kid self-declares that they're under 13, great, then I think you can trust that. I think if somebody is saying that they're over and we're more concerned about that they really might be under and there are risks associated with the type

Katie Hass:

The platform that they're going on to, then as the risk increase, so the methodology accuracy needs to increase with it. There, again, I think it needs to be left open to the type of methodology that's appropriate. What is needed to get on a pornography website is probably higher than what is needed to get onto a social media platform, which is probably higher than what is needed to get onto the Lego brick building app.

So, I think it really needs to be tailored to the right approach of where the concern is and where the harm is, and that's why laws should be somewhat specific to that.

Manmeet Dhindsa:

Thanks. And I think that aligns with, I think, our earlier discussion on the panel, and definitely what we've heard in other panels as well. One other thing that you raised is that the FTC should make sure to allow for quick innovation. What does that mean to you?

Katie Hass:

That means if you're choosing methodologies along the way and saying this new methodology is now approved, that that should be able to be done in a very rapid manner. That testing, of course, we want it to be thorough and accurate, but if you're approving methodologies because you feel like it meets a standard that the FTC is set for accuracy, we also don't want to see that those methodologies get bogged down in a regulatory bureaucracy that prohibits them from coming out in the quickest time possible.

So, it's that balance between taking the time to accurately test it and then list it, but also making sure that we are not stifling innovation or somehow preventing that innovation or not allowing some companies to start relying on it maybe.

Manmeet Dhindsa:

Thanks for that additional detail. So, Sara, do you see that companies are struggling to understand how to comply with COPPA and then simultaneously comply with these state laws requiring use of age verification? And if that's the case, are there places where you think COPPA could be changed to better facilitate compliance with both federal and state frameworks?

Sara Kloek:

So, this is another complicated answer that there's a couple points I want to make. First, when it comes to products that are directed at kids, that are already directed at kids and already need to comply with COPPA, companies are working to do that right. I think where we get into a little bit of a chicken and an egg situation is where if and when age assurance would be needed for kids that are under 13 or under the age of the state law, would you need to get parental consent to collect information for age assurance before the age is determined, or would you need to determine the age and then get parental consent for whatever the platform is going to do? It's a hard one to figure out. And I think the answer would be it depends on the age verification, age assurance, age estimation method, and what information is collected. I think one of the biggest challenges that we haven't really talked about is operators that are not directed at kids that may now have actual knowledge that a kid is on their platform.

I know that earlier, Chairman Ferguson talked about how it's really important to make sure that child protections are in place, and I agree with that. I think we need to make sure that kids are protected online. I think it's also important that we need to make sure that COPPA compliance is done in a way that makes sense. And if there's a kid on the platform, what needs to happen depending on the application or website or something like that.

I thought of a scenario here about a grocery store website. Let's say there's a 12-year-old that the family had a discussion and the 12-year-old needs to learn about budgeting and family management. So, the parents are like, "Okay, kid, you're going to go and do the family shopping for the week." Does the store website now have actual knowledge that they receive a flag that that kid is using the website?

Does the store need to set up a COPPA compliance regime? What does that mean? Do they need to get age verification, parental consent for the kid using the grocery store app or should the store say that the kid can't use it because they're not old enough yet? Should the parent let the kid use their account? What are we protecting the kid for?

I think these are silly questions that it seems like I'm asking, but I think it's a really important question to ask of what happens when there's a app that isn't traditionally directed at kids, but a kid is using it, what needs to happen where, and how do companies need to make sure they're protecting kids, but how do we make sure that it's done in a reasonable way where kids can still access information that their families think is valuable for them to access?

Manmeet Dhindsa:

So, I mean, do you think that these types of general audience properties that you just mentioned would be interested in using these age assurance technologies? And I'm thinking about this from a little bit of just as they're balancing their exposures to different things. So, for example, if they use age verification and then through it, they obtain actual knowledge that a user is a child. Do you think that these properties would be interested in using these types of technologies?

Sara Kloek:

I mean, that's a question for the grocery store, but I think that it's a question that needs to be talked about a little bit more is the general audience apps that may now be receiving or getting actual knowledge that a kid is on their platform. What are the expectations for, yes, we should protect the kids on the platform, but are we going to ban them from the platform? Are we going to say that they can't access that? Is the app supposed to set up a parental consent mechanism? I think it's a tough question that hasn't... Before these laws were passed at the state level, there was no way that they were going to get actual knowledge unless they set up a COPPA compliance regime just because they're a general audience app that isn't directed at kids. But the family decided they wanted the kid to learn about budgeting and family management, and that's the app that they decided to do. So, I think it's just an interesting question, what are we going to do about companies that now have actual knowledge that are general audience apps.

Manmeet Dhindsa:

Clare?

Clare Morell:

I just had a brief point about COPPA that I think is relevant to a lot of these state social media laws. So, I actually don't view COPPA as contradicting these state laws. I think that it actually can be a help to the state. So, I know Katie mentioned that they've done this in Utah, but I've seen a lot of states with these age verification parental consent laws actually rely on the verifiable parental consent standard in COPPA to actually help guide how these laws would be enforced.

And I think it's worth noting that the social media companies chose rather than comply with COPPA to just set the age at 13, they didn't want to go through the rigamarole of obtaining verifiable parental consent. And so, that could also just be a likely outcome of some of these laws where a company just decides that they don't actually want to go through that whole process. And so, they just then set the age accordingly.

And I will just say, looking through the FTC's COPPA rule where they outline these are acceptable methods of verifiable parental consent, it's very clear. And some of the laws I've saw, even Louisiana's, actually just lists out the same methods that have been approved by the COPPA rule. That's the methods that they've listed out in their law for age verification and parental consent. So, I really view the states and the FTC being very complimentary in this, that the work that the FTC has done in laying a lot of the ground rules and groundwork in COPPA can now be helpfully, I think, relied upon by states as they're crafting these laws.

And I appreciate that FTC says this can be updated. If new methods of parental consent become available, we will add them to our list of acceptable methods. And likewise, we will also make clear if there are certain methods that we're going to deny in terms of not being an acceptable verifiable parental consent. So, I just think it speaks to the important role that the FTC plays in this space in general and actually how that can be a benefit and help to states.

Manmeet Dhindsa:

Thanks, Clare. I want to turn to Jennifer. So, Jennifer, as we discussed in the first panel today, we've seen that a number of states age verification laws have faced First Amendment challenges. So, my question for you is, how can we create laws that require the use of age verification without running into these constitutional challenges?

Jennifer Huddleston:

So, I don't know that we have seen such a proposal yet, in part because one of the things I've noticed that we've done on this panel that we often do when we talk about these issues is that we've talked about only how these laws would impact young people. But the reality is age verification laws don't only impact those that are under the age set in the law, whether it's 13, 16, or 18, they impact all users of the internet.

So, if we're talking about age verification, whether it's for social media websites or for access to pornography or access to more broader internet services, those are also going to impact adult users and their speech rights as well. And this is going to be particularly true when we're thinking about certain groups like whistleblowers or others who may need anonymous speech, who may be concerned that age verification will require identity verification to some degree.

It may also impact certain communities more than others or certain people who may be less comfortable with going through those steps than others because they have a greater privacy sensitivity. We also have to think about when we're thinking about these laws, what do they mean for those exceptional cases?

Not only for perhaps the young person who's super advanced in trying to do college level calculus or is enrolled in a university when they're under 18 and needs to be able to access the academic message boards or certain information that might be otherwise covered by these laws, but what also does it mean for, say, the young person in foster care who doesn't necessarily have their parent able to give consent, or even for a far more common situation, for a divorced couple, are we going to have to see that who has the right to give parental consent for social media has to come down in the divorce decree or what happens when mom says yes or dad says no.

Let alone for those truly heartbreaking cases where perhaps the internet is a literal lifeline for a young person who's in an abusive situation where the parent may, in fact, be the problem and the internet may be how that young person's able to get help.

So, one of the real issues when we're talking about age verification is not only the impact that it has on those what may be at times perceived as exceptional cases when it comes to young people's speech rights, but particularly what it means for all of our speech rights online and all of our ability to access information, and someone alluded to this earlier in the panel.

Some of that question comes from the, what happens if you fail the age verification? How are you able to overcome it? How burdensome is it? What does that mean for your ability to access information? And what we've seen is that time and time again, because there are these other options available in the market for adults to protect their young people, that the impact on adult user's speech rights has really come into play here with perhaps the narrow exception of what the court found this previous term in the FSC v Paxton case as it relates to pornography.

Manmeet Dhindsa:

Thanks, Jennifer. Do you have any opinions or insight into how we can balance the adult versus kids issue that you raise? And I think, in particular, kids are often seen as a more vulnerable population on the internet. So, how should we think about balancing those two different populations as you suggest?

Jennifer Huddleston:

Again, I think this goes back to we've had this debate over the internet for quite some time. While it certainly has seemed to increase recently, we can go back to 30 years ago in the early '90s when we were having this about the early phases of the internet. This is a really difficult and complex issue, and it's been had over other forms of media as well, including video games, movies, books. There's certainly cases throughout history where parents have been very concerned about their young people's media consumption.

And I think that oftentimes, the better solution is to make sure that we're having those conversations with young people by trusted adults to encourage them to know what they're consuming, that parents are talking to their young people about the media that they're consuming, but also that we're not putting those burdens on adults so that they can't engage in what's often very important, political speech, the importance of anonymous speech, thinking again about anniversaries that are coming up with it being the 250th anniversary of the Declaration of Independence. Our founders certainly knew a thing or two about the importance of anonymous speech when they were writing around the time of the revolution.

And so, really looking at those key First Amendment values to make sure that we're recognizing that this is a form of expression. And that is something that makes it distinct from some of the examples we've heard around, for example, a bar or driving a car. It's something that we've seen have a different analysis and a different standard when it comes to that.

Manmeet Dhindsa:

Thanks, Jennifer. Clare, do you have any thoughts on this?

Clare Morell:

Sorry, I was just trying to unmute. Yes, I do. And I would just say just to set the stage a little bit, there are a lot of social media laws that have been passed requiring age verification and parental consent that actually have been upheld or allowed to go forward by circuit courts.

So yes, there have been some that have been enjoined by district courts, but both the Fifth Circuit and the 11th Circuit applied intermediate scrutiny. For the Fifth Circuit, it was to Mississippi's state social media law, in Florida, it was Florida's law for the 11th Circuit. And what they were both saying was that intermediate scrutiny applies because these laws are not content-based.

And so, I think when we're talking about the First Amendment, I think it really is important for states to not be defining these laws or justifying them for content-based harms or content-based reasons, but really focusing on this concept of account creation and minors entering into contracts with these companies.

And so that is, I think, important just as lawmakers are thinking through these laws and not wanting to get tied up with constitutional challenges under the First Amendment is just really how they're constructing and designing the laws to not make them content-based so they don't trigger strict scrutiny, but really making it about the regulation of these contracts where minors are entering, I mean, they agree to this whole host of terms of services with these companies and making it about that or making it about the medium.

So, I think, again, to Jennifer's point, this is not saying that kids can't express themselves online. There are plenty of places for them to do that. There's plenty of places for them to access news and information. This isn't talking about blog posts or Google searches or Wikipedia. This is just saying all these laws have been targeting either pornography websites for their obscene content, which we know there's a compelling government interest to protect children from, or social media platforms that have been shown to be highly addictive by their design features.

And so, I think that's the other important piece in just a note to legislators is in your legislative record or the history, what you're using to justify these laws, focusing on the design feature elements of the social media platforms, not about the content that kids are being exposed to, but how they use features like infinite scroll, notification, push notifications, autoplay, these aggressive recommendation algorithms, and that we're taking issue with that and how that has created an unsafe medium for children to be part of, but there is plenty of other ways for them to express themselves on the internet.

And so, I would just say those are some considerations to keep in mind, but also to recognize our circuit courts have not yet determined that any of these laws are unconstitutional. On the contrary, they have allowed for them to go forward for now. And so, I would just say it's important also, Jennifer mentioned, the precedent set by Paxton. And I'll just say while Paxton applied to pornography websites into a pornography age verification law, I think there are principles in that opinion that can just be helpful to keep in mind with other age verification laws.

And one of the foremost parts of the opinion just being very clear that adults have no First Amendment right to avoid age verification. Yes, adults have protected First Amendment rights to free speech, free expression, privacy, but it's not saying that because of those reasons, they have a right to avoid age verification. And I think increasingly with the technological means available that shows this can be done in a way that's not burdensome to adult speech, that's not chilling of adult speech.

I mean, processes that are double anonymous and take less than five seconds if you're just sharing your age key with a new website, I think just prove that point that these laws do not need to burden the First Amendment rights of adults and that the Supreme Court has been explicit that adults have no right to avoid age verification, are just some things I would encourage us to keep in mind in this conversation around the constitutionality of age verification laws.

Manmeet Dhindsa:

Thanks so much, Clare. So, I'm looking at the clock and we have a little bit over five minutes. So, I want to turn to the last question I have for you all, and this is about the role that the FTC can play in this space. So, as you've heard throughout the workshop, the Commission is interested in learning more about this technology and promoting innovation in this space, given its potential to protect kids online. But as we've heard today, there are also things we should keep in mind to ensure the technology is used responsibly and effectively. So, my question for you all, and I'll start with Jennifer, is what do you think the FTC should be thinking about in this space? And in particular, as Chairman Ferguson said, we would be particularly interested in understanding what we can do in the COPPA context. So, Jennifer, I'll turn to you first.

Jennifer Huddleston:

So, I think first, focusing on COPPA where there has been a clear congressional delegation to the FTC around what is and isn't appropriate with regards to the Children's Online Privacy Protection Act. One of the things that I think is very important to remember is that we would want to see a clear delegation from Congress should there be more expansive authority in this area, particularly given the nature of the privacy and speech sensitivity.

If an agency such as the FTC just were to suddenly expand its scope, that could certainly have concerns both in terms of the administrative procedure, as well as in terms of the First Amendment and privacy issues that have been raised today. So, I think it's really looking at what is clearly within the definition of the authority that the agency already has in the absence of any further congressional delegation, and that in light of any, if we were to see any federal law recognizing that such is likely to have further considerations around whether or not it has First Amendment or privacy concerns.

Manmeet Dhindsa:

Great, thanks. I'm just going to call on people unless they raise their hand. So, I'll turn to Katie.

Katie Hass:

Well, I appreciate Jennifer's comments, but regardless of where your authority may lie, the states are definitely looking for you to lead out in the determining the technologies that are sound, accurate, that we can build upon the principles, even if it's applying to 13 and under, or if new laws raise it to 16 or whatever it might be, you guys carry a lot of the laboring oar in helping vet a lot of the products and the technologies that are out there, you have resources that the states don't currently have, although we wish we did.

And so, we really look to you to flag issues of concern, to vet the technologies that are out there, to help us make accurate determinations on what are good methodologies, where the innovation is going. And I would just say we love partnering with the FTC. Manmeet and I have had the opportunity to partner on a case, and we welcome more opportunities to partner with the FTC, whether it's on this or other similar issues, because we learn a lot as we do it, but we definitely rely on a lot of your expertise in this area. So, continuing to do this for the states so that we can use it into our laws is incredibly helpful.

Manmeet Dhindsa:

Sara?

Sara Kloek:

I think that I spoke earlier about the FTC using their authority to understand age verification mechanisms better in a way that would both figure out how they can align with providing robust privacy protections, but also protecting the constitutional rights of all Americans, adults, and children.

I think when it comes to COPPA, understanding the limitations of the existing statute and understanding and figuring out where some parents might choose to allow their kids to have their information shared and some parents don't want their information shared and recognizing that all families are different and that parents and families should be able to be heard in what they want there and be able to exercise their own rights.

Manmeet Dhindsa:

And Sara, when you mentioned understanding the limitations of the COPPA statute, are you thinking of something in particular?

Sara Kloek:

I'm thinking Congress said what the FTC can do with COPPA. There's a law there. I know Congress is considering a new COPPA proposal and have been for many years, but right now, what the COPPA law says is what it says and making sure that whatever regulations or rulemaking is thought to go forward is within the bounds of the existing law.

Manmeet Dhindsa:

Thanks, Sara. Clare?

Clare Morell:

I'll be brief because I know we're basically at time. I would just completely second what Katie said. I think the FTC's expertise and your manpower is just so important for really charting a course for what age verification procedures and technologies are going to be privacy protecting for adults and effective for keeping kids out of the portions of the internet we're trying to protect them from.

And so, I think any guidance that you all can provide to states as they're passing these laws is really critical. And I think Australia's eSafety bureau has done this well, implementing their social media law and just they even ran age assurance technology trial and then put out a report of those results. Those types of information and resources that the FTC can provide, I think are invaluable. And then, the second thing I would just say is under your existing authority of the Federal Trade Commission to investigate and enforce unfair and deceptive trade practices, the ability to do that, I think, some of these social media laws are in response to just harmful practices by these companies to children. And so, I think there's also just enforcement of your existing authorities that can hold some of these parts of the internet that have been harmful to children, or now AI chatbots accountable if they're deceptively advertising their products and targeting minors and saying they're safe for minors, but then they're knowingly harmful to minors.

And I think just imagining down the road that then part of a hypothetical settlement agreement could be the FTC then requiring some pretty robust age verification measures and protections in place by those companies. So, I think those would be the two things I would say is just resources on how to do this well, and then using your existing authority to also make sure that none of these places we're talking about, pornography websites, AI chatbots, social media platforms, are unfairly or deceptively advertising their products to children and harming them.

And so, using your existing authority to enforce the law there and then potentially requiring more from them when it comes to age protections out of any settlement agreements.

Manmeet Dhindsa:

Thanks, Clare. And thank you to you all again so much for joining us today. I personally found this to be an incredibly interesting and informative discussion, and I imagine the audience did as well. So, thank you very, very much for joining us. To everyone online, we'll now take a 15-minute break and return at 3:00 p.m. for the last panel of the day. Thanks again.

James Trilling:

Good afternoon, everyone. Welcome to our fourth and final panel of the day where we will discuss deploying responsible age verification at scale. My name is Jim Trilling and I am an attorney in the FTC's Bureau of Consumer Protection. I will be moderating this panel along with my colleague, Diana Chang. We have a stellar group of panelists who work in different parts of the online ecosystem. We look forward to getting their perspectives on issues that have already been mentioned today as well as additional topics. As a reminder, short bios of the panelists are accessible on the FTC webpage for today's workshop. That said, I am going to ask the panelists to introduce themselves in the context of describing how their organizations fit into the age verification landscape. I will start with Emily Cashman Kirstein.

Emily Cashman Kirstein:

Thanks, James, and thank you to the FTC folks who organize this and the fellow panelists. It's great to be with you all today. So I'm Emily Cashman Kirstein. I lead child safety public policy at Google. And Google's been really deeply entrenched in this work for several years now. We've consistently taken a principled approach based on various ways that... Excuse me, age assurance affects our wide-ranging user base and the products and services that they come to Google for. So these principles include being risk-based, something we've heard a lot about today, privacy preserving and part of the wider ecosystem. So how does that work in practice? So we do ask users to declare their age at account creation, but on top of that, we have fully rolled out an age inference model in the United States and we'll be rolling out globally. That helps to determine if the user's an adult or not.

Earlier, we heard about the promise of using machine learning and AI for this purpose, and that's exactly what we're doing at Google. Our age inference model takes information we know about our users without collecting additional data and works to confirm whether or not that user is an adult or not under or over the age of 18. And those are based on factors like how long has a user had their account. If this person has their account for 20 years, probably an adult. Is that person, depending on their privacy settings, is that person searching for tax assistance on search? Are they looking for how-to plumbing videos on YouTube? Again, probably an adult. And if the model gets it wrong, as models are inherently imperfect, we offer as users the ability to prove that they're an adult using a variety of methods, something we heard from other speakers today, maximizing those options.

So they could upload an ID, use a credit card, take a selfie, use email verification. We agree that there should be a lot of options there for, again, that wide-ranging user base that we have in different products and services at Google. So from our point of view, it's really incredibly important to get this right, because we want to ensure that users have that right experience on products and services. If the user's an adult, we don't want to block them from access to critical information or put restrictions on their account to hamper their ability to use those services as an adult or ask for something people have talked about. We don't want to ask for privacy intrusive information if the risk doesn't warrant it. But on the other hand, of course, if a user's a minor, we want to ensure that they can take advantage of all of the default protections and settings that we offer as part of that experience, because that's how we've developed them for their unique needs.

And just to get into how we think about that approach from the wider ecosystem, right? We have skin in the game, we're working to be part of the solution, but we're not the only part. And we continue to identify ways to be a good partner. So for example, we've launched APIs for app developers and websites to receive age information through a zero knowledge proof pipeline. We've heard folks talk about that today too. In practice, this connects things like a digital government ID or an age estimation credentialer to an app or a website looking to receive the age signal. This is a signal that's requested by a developer and sent through that double-blind pipeline to ensure privacy. So I promise I'm wrapping up. Just want to say that we're very aware that this technology changes rapidly. This is something that I'm sure we're going to talk about a lot this afternoon. And because of that, it's so important to keep an open mind, see where the elements like privacy capabilities, precision improvements will take us going forward and finding new opportunities to find solutions. This is very much an ongoing conversation for us. Thanks.

James Trilling:

Thanks, Emily. Let's now go to Nick Rossi.

Nick Rossi:

Thanks, Jim. I'm Nick Rossi. I serve as Apple's Director of Federal Government Affairs here in Washington, DC. I actually spent 25 years in federal service before joining Apple, including as the staff director for the Senate Commerce Committee with jurisdiction over privacy and technology policy and oversight of the FTC. So I've been living with a lot of these issues for many years. And I'm very excited now to be at Apple because at Apple, we believe in advancing technologies and policies that protect children from online harms and protecting their privacy. We want users of all ages, really, to be able to have a great and safe experience with our products and services. And that's why a core part of our design is focused on keeping our users safe, and that's especially true for kids. So for example, I think one of the things we'll talk about today is our role in curation of the App Store.

And with the App Store, we've created a safe and trusted platform for users to discover millions of apps, but at the same time, we've also created a suite of tools and features to help keep kids safe. That includes tools that allow parents to approve or disapprove of any app download or in app purchase, to set app specific time limits or to control who can start a conversation with their kids. And we make these tools not only for our own services, but we make them available to developers to use within their apps as well. And then in this specific context, as we'll discuss, I'm sure in more depth later, we've rolled out within this last year a privacy protective age assurance solution that gives kids and parents the ability to share kids' age ranges with developers for the purpose of providing them with safe and age-appropriate features and content, but only with the approval of parents. So it's an important piece of this and one that we want to make sure is part of the dialogue. So thank you and really look forward to the conversation.

James Trilling:

Thanks, Nick. Antigone Davis, you're next.

Antigone Davis:

Thank you first of all for having us and also for all of the panelists who are participating. I've worked with a number of you before on this issue, and it's nice to see you all here, and thank you to the people who are interested in listening. My name's Antigone Davis. I am the head of our safety efforts at Meta, particularly our safety policy efforts, to give you a sense of our work in this area, which I think is really important from the developer perspective. At Meta, we've taken a comprehensive approach to ensuring that teens have an age-appropriate experience online on our apps. Since 2022, we've required teens on Instagram to prove their age if they try to change their age to over 18 through a video selfie or ID check. In addition, now if a teen attempts to change their age from 13 or 14, 15 to indicate they're 16 or 17, we also will ask them to do that ID check.

That's because we've launched something called Teen Accounts, which have built in protections that are offered specifically for 13 to 15-year-olds and also for 16, 17-year-olds, and they change across those age boundaries. In 2024, when we introduced Teen Accounts, we were really trying to reimagine the experience for teens online and to create an experience where we had built in safeguards that provided parents with peace of mind, but also gave parents the ability to have and put in place supervision and controls so teens wouldn't be able to change these safeguards without their parents being involved. We've been also using artificial intelligence along the way to help us determine if someone is a teen or an adult and continuing to expand in that area. We've launched a partnership with something called OpenAge Initiative. This is an initiative where they create basically a privacy preserving age key that can be shared with many different apps, including ours.

We think this is a very promising piece of technology, although it still puts parents in the position of having to do this or teens in the position of how to do this across numerous apps, which is why we are also really pushing for a piece of legislation that would essentially put in place an approach at which at the app store, you would be able to basically collect both the parental approval and age from the minor. We think this is a way to really address an ecosystem-wide approach. So while we'll continue to take our proactive steps in this area and ensure teens on our platform are placed into age-appropriate experiences and that we're doing what we can to understand age, we think there's an industry-wide challenge here to try to address. And we think the most effective way to do this is to really to basically have a simple process.

When a parent gets their teen a smartphone, the parent can easily go into their Apple account, their Google account, or their other account and confirm they're the parent or guardian, give the teens age that can be passed to us in a privacy preserving way with permission from that family to ensure that we can provide those age-appropriate experiences. I think what's really important here to know is that parents are likely with their teens when they're purchasing that smartphone. So it really sets up this moment in time that's very easy to collect this. Once you pass that phone to the teen, it becomes much more challenging. And we think that this will help parents because teens on average are using about 40 different apps. And so having to have them do this app by app, by app, by app by app really creates a challenge for them.

And what we've seen is that parents are highly supportive of this. In fact, 88% of parents support this approach to... Helping to address this issue at a multi-layered way. So in closing, I would just say we're very, very committed to providing age-appropriate experiences and to finding a way to know the age in a privacy-preserving way. And we think there is a simpler solution that starts at that layer of the ecosystem and then continues with what we do to assure age.

James Trilling:

Great. Thank you, Antigone. Graham Dufault, would you like to go next?

Rick:

That sounds good. Thank you, Jim. And thank you so much for having me participate in this really important discussion. I'm Graham Dufault. I'm general counsel of ACT, The App Association. We're a trade group. We represent small business app developers, connected device makers. I always think of our member companies as working to solve problems by leveraging software, leveraging smart devices. And they're across all different industry verticals from digital health and education to cybersecurity and agriculture. And so we fit into the age verification landscape in a few different ways. And the first thing I'll have to acknowledge is that a pretty large percentage of our member companies, maybe the majority don't use, they don't have reason to use age assurance services, including verification. And it's a really important piece of this puzzle to understand and to understand where age verification fits into the framework of making sure kids are safe online because so much of the app ecosystem and so much of the digital ecosystem in general is business to business. It's software for hire.

It's the type of business where age assurance, let alone full-on verification, has never really been a good fit because it presents a risk without having that commensurate need to address an age-related risk or need to provide an age-related benefit. And so one of the examples, I always think of a member companies that don't fall into that category of companies that really need access to age verification services. SwineTech, it's a provider of a tool that helps pig farmers manage their farms. It's software, it's distributed on Android, and it doesn't present those sort of age-related risks that you think of when you think of imposing on your users the process of going through age verification. So that's one example. Another example I'm pulling out of a hat is a member company in Cincinnati that runs Cincinnati's startup website, and they've also created Toyota's augmented reality app, where you can look around at the car.

These are services that don't create this sort of foreseeable age-related risks that would justify undertaking the measures involved with age assurance, including age verification. And the legislation that Antigone described, part of the reason we have some issues with an approach like that is because it would require all of our member companies, everybody with an app to receive a signal and handle that age-related information. So even so, we do have a lot of member companies that they're actively navigating Kappa. We've been around since 1998, so we have a lot of scars that are associated with how Kappa has evolved. In the mid 2010s, we were very active in trying to make sure that all of our member companies understood Kappa's obligations and knew how to apply them in the mobile space. We partnered with Moms With Apps, created a resource called Know What's Inside, where you can scan your app, figure out where your software development kit might be transferring information if you didn't already know how it worked.

And that was really critical at the time. People wanted to understand where ad networks were sending information and get a clear picture so they knew how to disclose things for purposes of Kappa. And we also just helped with best practices and mocking up short form privacy notices and things like that. And so we've been really involved with all this stuff. Some examples of member companies that deal with kid focused content are like FlipAClip and Thinkamingo, which provides a creative writing app. And so we do have a pretty significant percentage of the membership that are in kid-focused content. Second, we do have member companies that provide age-restricted services in one way or the other. We have online wagering app that needs to understand age and makes use of age assurance. And then last but not least, we do have a couple of member companies that provide identity management and age-asserted services themselves like SheerID and PRIVO.

So we have a range of interests in this space, but one thing we really can't lose sight of is that significant proportion of the app economy for which age assurance is not something that they necessarily need to be doing.

James Trilling:

Thanks, Graham. Amy Lawrence, you go next.

Amy Lawrence:

Yes. Thank you for having me and for allowing SuperAwesome to present our position on this. We're a little bit further away from the direct user than some of the other panelists, and so I hope that we can provide some color from the user and the kid-focused economy side of things. I'm chief privacy officer and head of legal at SuperAwesome, which is an advertising and technology company. So we're an intermediary in the advertising ecosystem or an ad network. In that position, we work with companies who are focused on reaching kids and teens in a responsible manner. Our focus is the under 18 market, and we want to use all available technologies to do that. So on one hand, we're working with brands that have products that are directed to kids and teens, think movies, toys, snacks. On the other hand, we have companies that have sites and services with an audience of kids and teens.

We use contextual advertising and our internal expertise to ensure that the ads we serve don't collect unnecessary data from minors and that the ads are age appropriate for their audience. And so in this space, knowledge of a user's age or an age signal can be very helpful in how the publisher provider or the advertiser interact with that service. And what we're also seeing is that as companies maybe shift from a general audience perspective to more of a mixed audience perspective or want to take in a teen audience or even advertise to teens. Based on the landscape that we have right now, age information or an age signal would be really helpful in their compliance efforts and working through some of the issues that come through state laws.

James Trilling:

Thank you, Amy. And to round things out, Robin Tombs.

Robin Tombs:

Thanks very much, James. My name's Robin Tombs, and I'm co-founder and CEO of Yoti, which is an 11-year-old business, and we provide age and ID verification and assurance services. We do a lot of age checks. We've done over one billion facial age estimations over the last seven years and about 1.1 billion age checks in total, the vast majority being facial age. We do those in lots and lots of sectors, so particularly social media, gaming, gambling, adult sites, often known as porn, vaping, e-commerce, supermarket self-checkouts, and a few other areas like gambling machines. We have quite a lot of big clients. We have TikTok and Meta, Sony PlayStation, Amazon Games, Epic Games, several others in the gaming sector. And we also have Philip Morris, British American Tobacco, Pinterest, and all sorts of others. So we have quite a lot of understanding of how to help businesses comply in the age sector.

And we've seen how that's changed over the last few years as technology has improved and more sectors and more regulations have been introduced and all of the challenges that has brought. We also have a digital identity, which is the Yoti app and over 20 million people over the last nine years have downloaded that app. You can either use it just to do age. So for instance, you could put your face into the app and not put any ID information and then share an over 18 or an over 21 depending on how old you are or what we estimate you to be over. And you can obviously add an ID doc so that you can prove your ID and that app is certified in the UK market by the UK government to a certain trust framework. Thank you.

Diana Chang:

Thank you all for those great introductions and previews of the ways that your organizations and the entities that you represent have been approaching this space. We have great representation today on this panel from different players in the online ecosystem. So we wanted to start by zooming out a little bit. Our panel is about deploying age verification or understanding age online at scale. So when we're thinking about scale, what types of online sites and services should be seeking to understand the ages of users online? I wanted to start with Graham. I think you touched on this a bit in talking about general audience apps, so I want to kick it over to you first.

Graham:

Sure. It's a great question. That's a critical one because you got to figure out where does it make the most sense for age verification. And there are a few different types of online services. It's services where you're providing access to age-restricted content, and Robin ticked off a few of these. So porn, online betting and gaming, content that presents distinct age-related risks like social media platforms, that's user-generated content, where you're facilitating peer-to-peer contact between strangers potentially. E-commerce platforms, in some cases, where they're facilitating the sale... Or they're selling age-restricted goods or services like tobacco. And so those are some examples and types of categories. But even if you're a company that's in one of those categories, it's not always true. And that might not be wise to get back to age assurance for every single one of your users or in every single scenario. And one example is where you have an e-commerce platform that generally provides non-age-restricted items.

95% of sales are non-age-restricted, but they do want to facilitate some age-restricted items. And so in a scenario like that, it would make a little bit more sense for you to provide an option for your user to go through the age assurance process, go through age verification before entering in one of those transactions instead of stopping them before they even enter the store. And if it's just the average person coming to the store, 95% chance they're going to buy something that's non-age restricted and saying, "Hey, you need to verify your age before you come in here." And so that's why I think you have to look at this from a risk standpoint because age assurance itself, it does present some privacy, some security risks. And that's why companies like my members are looking to experts like Yoti, experts like SuperAwesome or PRIVO to do it on their behalf because they're well-equipped to do it for them.

So whichever online service is thinking about making use of it... You have to think about it as a trade-off. How high is the risk I'm addressing here when I am looking to provide age assurance or a contract for it? Is the risk more along the lines of risk to mental health or potentially to reputation or is there a safety risk? And then if the risk is especially high, you might want to know with a greater degree of certainty of what the age of the user is. So you're looking more at verification rather than inference or estimation. So you have to look at it through a risk lens when you're thinking about scaling up, where is it going to make the most sense. And giving businesses a little bit of autonomy to decide what level of assurance is best for the job is critical. I don't think one size fits all is super great in this instance. So bottom line is, look at it through a risk lens, and those are the types of services that you can automatically look at as a category that's most likely where age assurance and verification might make sense.

Diana Chang:

Thanks, Graham. Emily, I suspect that you may have a different perspective on this. Would you like to weigh in?

Emily Cashman Kirstein:

Of course. Yeah. I would say going back to what we mentioned earlier, this really does get to our ecosystem-based approach. We really believe that every layer has a role to play here. And in practice, just to reiterate what I mentioned earlier, we do age assurance on our own products and services and offer mechanisms for developers and websites to request age signals from age credentialers like a government ID or a facial estimation provider through that double- blind privacy preserving API. But on top of that, we appreciate that there are really important conversations happening about, "What else can be done at scale to the purpose of the question? How can these processes be improved? At what level should that happen?" And for us, it's really, again, all of the above. We've heard from a number of speakers today, and we agree that app developers know their users best.

They know their product better than anyone, and whether or not they have an experience that creates risk to minors or not. There's an important responsibility there that shouldn't be shifted to other parts of the ecosystem. And that said, we understand the role of app stores and how they could be part of that whole of ecosystem approach. For example, by being required to share age signals as long as developers use that information responsibly to better protect kids. And the details matter there. We think this needs to be done in a tailored way. Only the apps that need this sensitive information should have access to it.

There's some proposals that would mandate that all apps have to receive this information whether or not they want it, and that's not privacy preserving and it's not risk-based. And I think I'd add separately that in many cases, users can and do access the app experience through a website, not through an app, accessed through an app store. So it's important to think through that angle when we're talking about scale as well.

Diana Chang:

Thanks, Emily. Amy, SuperAwesome sits further downstream. Can you weigh in on this?

Amy Lawrence:

Yeah. I think that there are companies along a spectrum here. So I think that SwineTech at one end is maybe a place where younger users are unlikely to show up, but there is a wide gulf of gray area of companies that have a product or service that might be attractive to kids. Maybe they don't think so, maybe some kids think so, maybe it's a niche product or service. And so I think that there needs to be a way to take into account where kids are likely to be present, that it's not just a black or white, that there is this scalable conversation about whether or not kids are likely to be present in your specific app or service. Kids' online lives right now are deeply multi-platform. So they're on mobile, they're on web, they're on social platforms, they're gaming, they're on streaming services on TV.

And the risks that they face are fundamental to how many of those modern platforms operate. And so things like social content... I'm sorry, social contact, content exposure, tracking is ubiquitous. And so the risk-based approach of whether or not you're doing some of those things does come into play with that. And if kids are likely to be there and you're facilitating certain data collection or contact with strangers, those are all things to take into consideration. And what that gets me to is that saying, "We're not a kid's product," is no longer a sufficient reason to avoid understanding age. And that's sort of where we've been with Kappa for a very long time is that if you want to avoid children, that that is probably the most risk averse way that you can handle it is just ignore that there might be kids on your site or service.

But if minors are likely to show up, age aware design should be the safety and compliance baseline instead of just avoiding the idea that there might be kids showing up. And in practice, kids are on a big chunk of the internet. And one of the earlier panelists that today said that one third of internet users are under 18. And so it's difficult to think that there are very many places where there aren't minors online. And if you read privacy policies, a lot of sites and services already claim to collect and use data to personalize content or to better understand their audience. Age information is really just one facet of that. And in particular, I'll just note that any service using advertising or an ad tech ecosystem should be thinking about

Amy Lawrence:

... the age of all their users, including platforms, SDKs, any monetization tools, because a legal landscape is already there such that the age of the user directly affects what you can do in terms of targeting, profiling, tracking, even if it's just measurement of your ads and the personalization of that as well. There's a lot of gray areas that some are darker, some are lighter, but I don't disagree that we should take a risk-based approach. I just think that there's probably more risk out there than a lot of companies would like to really think through.

Diana Chang:

Thanks, Amy. Antigone.

Antigone Davis:

Yeah, I just wanted to pick up a little bit on Amy's last point because I agree with it. I think as I was listening to the examples that were given, I can't speak to the SwineTech app, which has now gotten a lot of play on the panel, but if the SwineTech app, for example, is an app that can be downloaded in the app store by anybody who has access to that app store, then there are necessary steps that need to be taken to ensure that certain minors are not on that app already by law, by COPPA. But in addition, if a minor can download that, even if they're over the age of 13, and that app provides... Let's say that it provides a way for that farmer to link out to research that's in the broader internet, you've now unintentionally within that app makes sense for an adult using the app given a minor who may download that app access to the internet, which is why I think when we think about how we solve here, it becomes incredibly important to think about where parents and families are.

Parents are providing their teens with a smartphone. A smartphone provides them with access to an app store. It also has an operating system, and that provides them with access to in the neighborhood of one million to two million apps. They use 40 on average. Parents need to have a way to approve the download of those apps because I think as Amy indicated, the risks are probably quite across all the apps, and I think that risk-based approach may not be exactly the right answer here. You're really looking at a place in which we should be thinking about teens downloaded the app, their data is being processed. They've taken on terms of service. There's a contract there. Where does that parent fit in that role?

I think when they have access to a million to two million apps, we need to be providing a mechanism to ensuring there's a way for someone to prevent teens from having access to that app. So that app developer likes SwineTech, that it has an app that's built for adult farmer doesn't have to spend time trying to deal with those other issues. So apps like ours where we know a minor's going to use it, we're going to put in place certain safeguards.

I think the last piece that I would say here is that we already have the ability to do this at present. We see it in the context of in-app purchases where parents have to approve in-app purchases. We have the systems in place to do this in a privacy-preserving way. Providing an age signal of over 18, under 18, between 13 and 15 is pretty privacy-preserving. You're not providing a date of birth. You're not providing an ID to the extent that you have to verify. You would be verifying in one place in the app store instead of across multiple apps with varying security protocols.

Diana Chang:

Thanks, Antigone. Nick, I see your hand up. Let's take your response as the last one for this question.

Nick Rossi:

Thanks, Diana. I appreciate it. I'll be quick because I know we're still on the early questions. But I just wanted to underscore that even though we call, for example, our tool Ask to Buy, it's not limited to in-app purchases. It enables parents to approve or decline to approve the download of any app, even free apps. If you'd watched some of the advertising that's occurred in at least the D.C. market over the last year, you would think that that didn't exist, but it does exist. We've given parents the tool to approve every single app download.

James Trilling:

Thanks for those responses, everyone. We're going to shift gears a little bit and assume that we know where age verification will be employed and that a risk-based approach is in play in some way. Let's talk about some of the challenges in deploying age verification at scale. Throughout the day, numerous speakers have raised concerns about balancing privacy with age verification. Graham has alluded to that on this panel, as well as some of the other panelists. There also have been allusions to friction. Nick, how does Apple think about these countervailing considerations when it comes to deploying age verification?

Nick Rossi:

Thanks, Jim. I think one thing to think about here is... And I know balancing is part of the question, but we don't really like to think about it as balancing. We very much see the goal is keeping kids safe online and privacy is part of that. We know, for example, that one of the reasons that people love Apple is our commitment to user privacy in everything we do. So the principle that we apply here, and I know a lot of folks have talked about it through the course of the day, is data minimization. It's really a key part of our foundation, and it's the idea that we're only collecting the minimum amount of data that's required to deliver what users need.

In this area, we certainly recognize that we have a role to play and that the App Store has a role to play, but we also think, as you've heard, it depends on the different levels and everybody having a role to play in that process. In this area, we think that in order to keep kids safe, we've got to be mitigating the risks of scams, frauds, or other harms that could result from children's personal information being collected or retained or shared unnecessarily with folks who don't need it. What you've heard today is that things like a child's birthdate, that's an indelible attribute. It's something that's highly valuable to commerce and it's something that can be leveraged for tracking or used for targeting advertising.

So the way that we've approached this with our declared age range signal is that we've rolled it out in a way that enables parents, first off, to choose whether to share their child's age range. We've also done it in a way that never discloses the actual birthdate, but also that trusts parents who have established that they are adults to provide their kids' ages without having to turn over to us sensitive documentation like a birth certificate or social security numbers. From our perspective, that keeps parents in control of their kids' sensitive personal information while minimizing the amount of information that's shared with third parties. So that's one of the ways we've tried to approach this.

James Trilling:

Are there other panelists who'd like to respond to that same question? Antigone?

Antigone Davis:

I actually just wanted to respond to the point about having the ability for parents to download to provide parental approval. While a tool like that, it may exist, what isn't the case is that that tool is automatically turned on for a parent when they set up a minor's account, which is why we see ourselves trying to figure out a way to create that. So you could imagine if that was automatically turned on any smartphone, on any iPhone that Apple sells and has an account for a minor that would automatically turn on, that parental approval was done, it couldn't be turned off if it was a minor, you wouldn't be having this conversation, but that isn't the case. So parents still struggle to figure out how to approve app, how to share, how to approve the app download.

James Trilling:

Robin?

Robin Tombs:

Yeah, I think also there's a lot of businesses online which are not using apps. Obviously, the adult sector is one where all of the sites are effectively browser-based sites. So yes, obviously, you can be quite efficient with an app signal, but there's lots of sites which would basically need to choose to also plug into that system. I think there are risks, but we know that children self-assert sometimes incorrect ages, but there's also evidence for parents to do that.

Obviously, a parent may feel in a strong position to be able to say, "Well, I think my child is old enough to play a particular game," even if that game is an 18-plus or a 15-plus, but the child is maybe 11 or 12 because their friends are playing it as well. So you've got a risk that some parents will basically allow that age to be wrong, and then you've got conflicting signals in the system that, well, allegedly this user is 18 or 19, I'm going to allow them to effectively play that game. The other person may be 50, but they're claiming also to be younger, and maybe the signals in the chat are that one of those players seems to be a lot younger. So it's not sadly as simple as thinking that there might be one place where you can do age and then hopefully rely on it for many months and years.

James Trilling:

Nick, did you have something to add?

Nick Rossi:

Yeah, I just wanted to clarify that when a parent sets up a child account on their iPhone, that Ask to Buy is on by default for children under 13 or 13 and under. When you invite someone into your family group, you're encouraged to turn it on for kids, all minors, for 18 and under.

James Trilling:

Graham?

Graham:

Yeah. On this point of the privacy issues involved with age assurance and scaling it up, I just don't want us to lose sight of the fact that conducting the age assurance, conducting verification in particular does itself still pose privacy risk, and even to the extent that there are technologies that allow you to pass along a signal, that doesn't include an entire identity. You are still injecting risk and you are still taking on risk and you have to match that against whatever the risk that's presented by the app itself or the content that's provided or the services that you're giving access to. The description of what might happen with a kid downloading SwineTech, that's the kind of risk that's pretty speculative, right? You have to match that potential risk against the risk that you are necessarily taking on if you require SwineTech to now, again, under the COPPA regime and receive age signals and stuff like that.

James Trilling:

Emily?

Emily Cashman Kirstein:

Thanks. Just to get back to what we were talking about and the importance of the risk-based approach here, I think part of this is in understanding that the age assurance really has to align with degree of risk to a given online experience. A lot of folks have been heartened to hear that people are really thinking about this with the importance it deserves as part of this entire panel. There's always going to be a privacy trade-off. This inherent privacy trade-off underpins everything we've been talking about today, and the degree of that trade-off really does have to be commensurate with the risk. So when we're getting into brass tacks here from our perspective, no one should have to upload a government ID to use a weather app, but you probably should, for example, to view pornography.

Just to be clear and to follow up from one of the previous panels, there is no pornography on YouTube. It's against our policies for all our users, not just minors. So just want to be clear about that and that we also have additional age restrictions to ensure that minors don't encounter things even like racing music videos or other age-restricted content that, again, may be appropriate for an adult, but not for a minor. We aren't opposed to hard verification as a concept. We just believe that it has to be limited to the highest risk activities on the internet. I think that just goes back to, from our point of view, it's really not a one size fits all. That doesn't fit the bill here. There's ways, really different ways that people use the online world, the way they use it is nuanced, and so should the protections. This isn't a satisfying answer to any degree, but in our view, it's really what needs to be done. This has to be thoughtful and the details matter.

Diana Chang:

Thank you for those responses. So we've examined through the last couple of questions some of the risk-based approaches, some of the privacy harms that can result when it comes to talking about age verification. I want to change the lens a little bit. Our discussions sometimes seem to presume that implementing age verification or understanding age can impose costs to online sites and services, perhaps without providing any benefits. We did hear earlier in the day that one benefit to businesses could be that age verification could facilitate compliance with the laws, including those that might require implicitly services to know the ages of their users. So I wanted to talk about that a little bit and whether or not what we're talking about, understanding age online, can provide any benefits to online services and sites. Amy, I think you've talked about this a bit. Do you want to start off?

Amy Lawrence:

Yeah. Yeah, happy to. I think that you're right to call out risk reduction and compliance with other laws because I think that's the most obvious area where it is a real benefit that age assurance can lower your legal risk because it can be evidence that you've reasonably taken steps to comply with COPPA or other laws with age-specific obligations or restrictions, and especially, as we talked about earlier, the rapidly shifting landscape at the state level. Amelia Vance had a great slide this morning about the current patchwork of laws that have different restrictions. So for a company, being able to shift compliance from reactive to proactive is no small thing, and being able to avoid funds, litigation, bad PR, those are all benefits.

But what I'd say is maybe the best argument or the most effective for companies is that age assurance could also be good for business. It can be used as a trust lever and be something that you can rely on to help grow your site or service because when kids' environments are age appropriate, you reduce user churn, and that churn can be from kids that leave or disengage because they experience inappropriate content or contact from strangers that they didn't want or from parents that discover inappropriate content on a platform and then force deletion of that app or service from the kids' device. So there's that. Then it'll also open up monetization options because from an advertiser perspective, they usually want either assurance that inventory has no minors or assurance that inventory is family-friendly or safe for kids and compliant with youth regulations. Age assurance helps both by creating clearly defined inventory segments, or you provide an age signal that somebody is under 18, you can provide advertising that is appropriate to that.

It would also allow you to supply age-appropriate experiences by age bands. So in a lot of the app store age verification laws we've seen come out of the states so far, we have these age bands of under 13, 13 to 15, 16 to 17, and 18-plus. If you know that much about the audience or the user, then you can make sure that the tone of the ad, the tone of the content it appears on is all appropriate to that age group, and then that reduces the chance of backlash that an ad appeared next to an inappropriate content, that there was brand adjacency problems because an ad for alcohol appeared on content that is normally watched by kids or the reverse that kids' content ads appeared on pornography. You reduce all of those problems. So monetization model using contextual advertising has always been supported by COPPA. If safe spaces for young people are making money, if they are profitable, it encourages more and higher quality safe spaces for kids, and that improves the entire ecosystem around kids on the internet.

Diana Chang:

Thanks, Amy. Robin, I'm going to kick it over to you.

Robin Tombs:

Yeah, I guess we've got a very clear example. Luckily, we've worked with Yubo, which is a social discovery app and has about 85 million users, and they started using Yoti's age services back in 2020. They were the first company to use Yoti facial age estimation and they've done over 100 million checks, so effectively all of their audience. Interestingly, they do a survey of a large number of users. Those users, once everybody had been age checked, the users very clearly, a majority of them, something like 80%, said that they felt safer on Yubo.

They also said that they thought there were many, many less bots accounts because the bots couldn't pass our likness to do an age estimation, obviously, and that they effectively felt they were speaking to the right aged other people on Yubo, so that a 13-year-old who was claiming to be 17 in the past effectively was spotted and therefore a lot of the children felt they were actually having conversations with the right aged people and people weren't pretending that they were a different age. So that was a very clear example for Yubo that obviously there was a compliance benefit, but there was also a significant trust benefit for the brand.

Diana Chang:

Antigone, I see your hand up.

Antigone Davis:

Yeah. Thank you, Diana. You mentioned where does everybody benefit? I think Amy did a really good job of showing how all along the types of different apps you may have benefit. I think just going back, for example, to that if you think of a stack of different apps, types of apps, let's go back to the example of the weather app for a minute. Weather apps actually have terms of service. In fact, if you take The Weather Channel app, their privacy policy actually explicitly states that if they are collecting the personal data of a minor without parental consent, they want to remove that personal data. So having that information actually enables them to be in a better position to abide by their terms of service, to follow up on their terms of service, to execute other terms of service, whereas say a user-generated content site like Meta, where we want to be able to provide an age-differentiated experience based in relation to that content, we have a benefit from understanding that age to provide and ensure the safeguards that we want to offer to our users are in place.

So I think rather than saying, "Oh, it's this one set of apps or apps that offer..." I think one of the examples people give are apps that offer age-differing experiences. They should be required to do some kind of age verification. Well, if you set it up that way, you disincentivize people from actually creating age-differing experiences within their app because then all of a sudden they have to go out and do this added verification process. So not only is there a reason for almost every app that I can think of to have some sense of age awareness to effectuate their service in the way that they intend to, but it's also provides for companies like ours who do want to do age-differing experiences to provide those safeguards. It incentivizes companies that know that there may be young people coming to their app or that are interested in having young people use their app to be able to create those differing experiences and then effectuate them in the most effective way.

Diana Chang:

Graham, over to you.

Graham:

Yeah, I think the benefits are really interesting for all the different use cases that we're talking about here, including the weather example. I think, again, we just can't lose sight of the fact that there is a potential benefit here. There's also a risk. I think for that reason, it's probably the wrong approach to require every single online service to receive age information and then figure out, "Hey, how might I use this to benefit the service or in a creative way that might help users of the app, but also might be misused?" I think, as Antigone described, requiring a class of apps to receive a signal might create a disincentive to do that class of apps. I'm a little bit more worried about requiring all apps to receive a signal, and therefore disincentivizing them to put an app on the store in the first place that even might be a general audience app. So that is understood, and that is a risk that we're pretty worried about.

Antigone Davis:

Can I just respond quickly back? I wasn't saying that requiring age-differing apps. What I was saying is that if you have a system in which only those who offer age-differing experiences have to verify age, you incentivize the entire ecosystem of apps to not offer age-differing experiences because the minute that you do that, you take on this extra burden that you have to individually do. So slightly different than what you said.

Graham:

Okay.

James Trilling:

Nick, did you have a point you wanted to make before we maybe move on to another topic?

Nick Rossi:

Well, I can probably make it as part of the next conversation.

James Trilling:

Okay. Thanks. Most, if not all of our panelists represent organizations and entities that operate across multiple jurisdictions that, as we've heard today, might have divergent requirements when it comes to implementing age assurance or verification processes. What are some of the challenges that you have faced in implementing age assurance or verification processes across jurisdictions, and what solutions have you found to be effective? For this one, I'm going to start off with Robin.

Robin Tombs:

We see lots and lots of challenge. That's partly because different regulators in different countries have different thoughts on how to do age checks. Sometimes that's clearly written into the regulations and sometimes it's less clearly written into the regulations. So you have quite a lot of challenge there. Even if a company wants to comply, is it nice and clear how to comply and is everybody effectively therefore complying in a sensible manner, or are they worried that some people may interpret the rules to allow them to do X when somebody else feels that it's Y and maybe therefore it's less friction X? So I think that's been quite a big challenge in the adult sector.

For instance, Ofcom in the UK allows a site to do effectively one check. That doesn't mean that they should never do another check, but effectively that currently is the expected process, whereas in Italy and France, sites are expected to do those checks daily or even more often than that, if somebody signs off that account or closes their browser and comes back a few hours later. So that is a challenge for both users and sites if people are effectively being asked to do these checks potentially a disproportionate amount of the time. Regulation is a challenge.

I think also a lot of sites are not experts. Certainly, obviously, we're on a panel here with a lot of very technically expert brands, but there's lots and lots of sites which are not really sure, "How do I test to ensure that facial age estimation is accurate and not biased across skin tones and ages and sexes?" They don't necessarily have all of the capability to do that. Initially, there weren't really any independent testing houses because they probably didn't have the data to do those types of tests either. Now, that's changed in the last two, two and a half years with particularly the U.S. National Institute of Standards and Technology. They now do a huge amount of testing of vendors who offer things like facial age estimations.

So there are benefits now coming through that businesses without expertise can rely on independent testing to ensure that they pick vendors who are hopefully offering good services, but there's still lots and lots of challenge there in terms of it's quite straightforward to look at a facial age estimation technique and say, for instance, in Germany, but as long as you put a... It used to be a five-year buffer on 18, then you could allow somebody to use that. So as long as they basically looked over 23, they didn't have to use an alternative age check method like an ID document. They could pass the test without identification and that's reduced to a three-year buffer now. But on some other techniques, it's much harder to be as scientific. Is it definitely my credit card or is my child using my credit card without my knowledge to basically try and pass a test?

Unfortunately, there are a lot of challenges. I do think the technology as always is moving forward and allows more and more privacy-preserving techniques. We've heard about email assurance, which is really interesting. There are competitors to Yoti who now do age estimation on device. We do that on offline supermarket self-checkouts. We don't yet do it on the browser, but we do it in our app. So more and more privacy issues are being solved by people being able to say, "Look, actually, you can do the authentication of the pass key or the actual facial age estimation on device or you can get a reusable digital identity and then just share 18-plus and you don't actually have to go to each site and try and do an age check. You are just effectively choosing to share the 18-plus or the 16- plus." But there are plenty of challenges and there's a whole industry working to try and make that easier and easier over the years to come.

James Trilling:

Do we have other panelists who'd like to weigh in on the challenges of operating across jurisdictions when it comes to age verification? Amy?

Amy Lawrence:

I would add that there's operational complexity with partners and supply pads and communicating information through the ecosystem. So different publishers platforms implement age assurance differently based on their local laws. Some provide strong signals. Some set the COPPA flag in their bid requests. Some don't. They're not required to. Some do age segmentation or geofencing and others provide limited transparency. So what we've done historically and in practice is we rely on contextual analysis and content suitability frameworks.

I wanted to raise this just to make the point that an age signal really simplifies things for the recipients of that age signal to be able to say with assurance that, "Okay. We want to make sure that what we're sending here is appropriate for under 13s or is appropriate for a 16-year-old versus a 12-year-old," because the old school way of doing this is to actually look at the content of each site and look at what's happening there. Does it look attractive to kids under the COPPA test or another test depending on where the publisher is based? So, looking at that at scale using an age signal that can be passed through the ecosystem creates a lot of efficiencies and resolves a lot of inconsistencies.

Diana Chang:

Thanks for those responses. Amy, you just touched on this, and I think it gets us into our next topic. Everyone has touched on this at some point today, but we've got a great panel that I think can address the question. What is the right division of responsibility for age assurance across platforms, app stores, devices, and third-party providers? And I'd love to hear your opinions on why. If you can build it in, we just touched on this, but if you are receiving potentially age signals from a whole host of different entities throughout the ecosystem, how do organizations responsibly handle potentially conflicting signals? Antigone, I want to kick it over to you first.

Antigone Davis:

Yeah. I think what I would want to say first is that from Meta's perspective, we think everybody within this ecosystem has a role to play. So we think that on the operating system, the app store has this initial place that they can play a role in understanding age, understanding the parent-child relationship. They're uniquely situated to do that and then to provide a mechanism for sharing that age signal and that parental approval to download the app with apps that are a part of that app store. Then that age signal is passed to a developer like ourselves and we have a role to play in providing the right type of experience and ensuring that we've got a safeguarded experience in place like what you see for teen accounts across whatever app when you are working with a user, providing

Antigone Davis:

... The right safeguards based on that awareness. And then in addition, we will do the type of verifying of that signal that we already do, and developers should do that if they have additional information to make sure that we have that right age as well. And parents also have, in the context of a minor being on an app, a role to play. But we have to make it easy for them to do that. And that's why each of these parties has a role to play.

Diana Chang:

Emily?

Emily Cashman Kirstein:

Thanks, Diana. I think we're saying a lot of similar things, but one of the things I want to pull on is it's really tempting to look for easy one-size-fits-all solutions here, but the reality is it's just not that simple. So when we're looking at different policy proposals and what might make sense for the wide variety of ways that people are using the internet and apps, we've all said this, everyone's got a role to play. From our view that's app stores and developers. App stores could provide developers with the ability to call an age signal. Operating systems can enable parents to activate and elevate parental tools. And it also has to be risk based. And I know there's been some conversation here that I think it's important to dig into. We really do believe that only apps that are risky for minors or have those differentiated experiences that we've been talking about really should be required to utilize that age signal from an app store. Apps that are safe for everyone don't need or want that information.

And if you'll indulge me, I want to go back to some of the questions we were talking about, about those requirements and the claim that it's going to take away incentives for companies to do the work that many of us are doing here. And I think some of the proposals we're talking about actually create a framework where adopting safety features is the lowest liability. It creates a path of least resistance. It spares businesses with low risk from really tough mandates, really expensive mandates, but it forces high risk platforms to either put those protections in place or block kids entirely. And I'd add, the protections we're talking about are really common sense that developers can put in place and in many cases are already designed that way or already part of other legal frameworks.

Then on the privacy preserving side, I think we've talked about this and it does warrant repeating, that age information really should only be shared with apps that really need it. This is sensitive information and it should be with the permission of the user and the parent.

Diana Chang:

Thanks, Emily. And this is open to the whole panel, but I did want to turn it over to Amy. The second part of my question dealt with, what is the right approach to responsibly handle potentially conflicting signals? So as many panelists have said today, if everyone has a role to play and everyone's using a different method to collect age, what is the right way to deal with that information?

Amy Lawrence:

I think that that also depends on where you sit in the ecosystem. So if you have a direct connection with the user because their connection to the app store, their connection to the developers, then your positioning is a little bit different than someone who would only ever be a recipient and can't control what age verification method is used or how it is verified internally. And looking at the papers that Yoi would put out about its compliance with international standards or with its own stringent requirements. So I think that that positioning is very different from the apps and vendors downstream who might receive an age signal and don't have the ability to go back to a user or allow for a user to appeal what their age decision is.

When you're further downstream, you should set up a decision tree on how you're going to look at these things. It's really a governance problem of what do you do when you have these conflicting signals. And if there's a conflict, if there is a strong signal that exists, that is a good indicator that you should follow it. But we need some more guidance on what is a strong signal. Is that because they used a certain verification provider? Is that because it comes from two different sources? Is it because you just don't have any other of your own information? There are questions as a recipient of age signals that I think are easiest solved by choosing the most protective band. So that's where I think that you'll see this go is that if there's two conflicting signals that I've received from two different publishers and I have no reason to weigh one heavier than the other, then the defensible position is to use the most protective band that you have available.

Diana Chang:

Thanks, Amy. Nick, I'm going to kick it over to you.

Nick Rossi:

Thanks, Diana. I think I want to echo a little bit some of what Emily said about risk and focusing on risk, but also on the importance of parental permission in this process. It may very well be the case that there are lots of apps that could make use of an age signal, but some of the requirements and some of the legislation would seem to require that age signal to be pushed out to every app regardless of how much they might need it and regardless of whether or not a parent has a role in saying yes or no about the sharing of their kids' age information.

So I think risk also comes back to this question about, for those who were able to watch much, if not all of the discussion earlier in the day today, you heard about everything from simple check-a-box exercises to parental vouching that's paired with confirmation that a parent is an adult, you heard about facial scanning and you heard about behavioral clues and all kinds of different possibilities here. It may be that the level will depend on the point in that stack where you are engaging. So what's appropriate for everybody entering the app marketplace, like everybody entering a shopping center may be different than those who are going to that one place that is selling alcohol or selling something that is more concerning. And it relates to how you handle potential conflicts as well, because the best answer is probably to look to where additional information may be available.

Our declared age range API does give developers a helpful addition to the set of resources that they can choose from, like third parties, tools like Yoti or the information that they possess directly. But we recognize that some developers, not all, but some collect or possess more information about their users than we do. And some may have a separate legal obligation to ensure a user is a certain age, like an app that sells alcohol, telling them that somebody's over 18 is not going to be sufficient to know whether or not they are 21. Developers are in a position as the ones who are creating and serving content within their apps to have important context about their app and their users. So in many cases, that's going to include data about the user's age. So looking to who may be in the best position to have the most specific information is probably part of the answer here.

James Trilling:

That's a good segue into, I think, a closely related question, which is how do entities that develop methods for understanding users' ages know whether their measures perform accurately? And relatedly, how do entities that seek to rely on a third party to understand users' ages, whether it be through an API, as Nick and Antigone and others have mentioned, or otherwise know whether that information that's coming from a third party is able to provide accurate information? Graham, why don't we start off with you for this one?

Graham:

Sure. It's a great question, and it reminds me of the questions that a lot of our member companies are asking around AI as they are sellers of AI services and developers with app layer AI tools, and also purchasers of AI services. They're hungry for standardization and looking at best practices against which they can match what they're developing and match what they're looking to purchase. So I always look first at the standard.

And I know what was mentioned earlier today, ISO/IEC 27566, sorry to make eyes glaze over by just throwing numbers out there, but that's the standard that you look to first for a bit of a taxonomy. And they have a couple things that you can look for, which are number one, classification accuracy, is it 97%, 98%? And second, you want to know what are your false negative rates versus your false positive? Is there a pretty even distribution between those two? You want outcome error parity to the extent you possibly can. You want the process to be testable so that you predefine your test points and that you support standard test protocols.

Then lastly, and this one's pretty important, is to measure and report on the completion rate. So how many people are dropping out of the process, adults and kids alike? And that helps you understand what the friction is like with the assurance process and whether or not people are liable to circumvent, as we all know, is likely to happen, or at least there's a temptation to do that. So you do the risk analysis and you need age assurance. And if so, where are you going to need to use it process-wise? If you do need it, look at the classification accuracy, look at false negatives versus false positives. There's uneven distribution, then that might be a potential issue. And try to understand whether or not there's a really big incentive and whether it's possible to totally circumvent the process. And then you'll have a good sense of whether or not the age assurance tool that you're looking at using is effective. So from our members' perspective, that's a taxonomy of what they look at.

James Trilling:

Thanks, Graham. Other perspectives on this one, Robin?

Robin Tombs:

Yeah, I think that whole issue of age signals, and if you include an age check as a signal, because age checks are not perfect and people shouldn't assume that they are. I think the key, as Amy said, was that you have a responsibility, or certainly the regulators should expect you to have a responsibility to look at things dynamically. You may well receive an age signal and take some comfort from that. But if you then see contradictory signals, those shouldn't be ignored. You should look at those and then think, "Okay, this is a dating site. We shouldn't have under 18s on the site." And if you begin to think you might well have somebody under 18 on that site, you should then basically request an age check to try and ensure that, as Amy says, that you're prudent. And that might be annoying for somebody who then proves that they are 18, but in some respects better to be safe than sorry there.

So I think there is an expectation that people have to understand none of these things are perfect, just like offline laws. And sometimes you need to take a second look. There's lots that you can do to benchmark that certain age checks are likely to be fairly accurate or they're less likely or less vulnerable to spoofing. But as people have said, kids are clever. Kids might be using shared devices, in many countries an older brother may be sharing a phone with a younger brother. So you've got to live in the real world and recognize that, okay, just because we've done one age signal on somebody or one age check, that doesn't necessarily mean months later that it's definitely the same person. It could be a different member of a family who's younger. So I think the regulations will mature and the regulations are likely to ensure that you can't basically say, "Look, I did something on that date and I didn't really feel I needed to do anything more at a later date." I think it's always going to be a dynamic situation where you have to assess for risk.

Diana Chang:

Thanks, Robin.

One theme that's come up today is this idea of empowering parents and families to help protect their kids online. I'd like to understand from the panelists whether you view age assurance, age verification as sort of one tool in the broader toolset that might include things like parental controls tools or other types of features or aspects. Can you talk a little bit more about this tool set and can you tell us a little bit more about how age assurance might fit into that toolkit? Amy, do you want to start us off?

Amy Lawrence:

Yeah, absolutely. So I definitely see age assurance as just one tool in the tool belt and that a multi-layered approach to this is the best way to deal with it. What we've seen in a lot of research around kids and teens and parents as well is that one of the most effective measures that families can take is to just have a conversation about what's happening in their kids' lives online. And so that is talking about household rules of where are the places that you spend your time, for how long do you take breaks? Where is the device? Is it in the living room or your bedroom? That conversation piece is maybe the strongest layer.

Then from there, we have settings that there are iOS settings, there are device level settings, there are settings in every platform that kids are spending time that would allow parents to be involved in setting the standard of what they want their child to be able to do, how they want to engage in that platform and how they don't. I know from experience that is no small task to discover and implement all of the settings across the board, but that is another piece of it. Then the third piece is really healthy habits of just making sure that you're normalizing the conversation of where kids spend their time and what they're doing online and what they're experiencing online to get that back and forth so that it's not a black box. Age assurance fits in because when it's implemented well, it can help enable safer defaults for younger kids in particular, with stricter privacy, limited social features, more moderation, more broadly can set up age-appropriate experiences. Think of age assurance less as a gate and more as a routing mechanism that if you are a certain age, you might go this direction if the service sets up an available experience for your age group. Then as ever, there are the parental consent flows that age assurance may help kick off, may help you avoid in some situations if the minor is over 13, but it may just be a different experience because other state age restrictions extend the 13 age limit from COPPA. So there might be a parental consent aspect as well.

But I just want to caution, and then I'll hand the baton over, age assurance alone doesn't solve safety. Knowing that somebody is a teen doesn't mean that they're not going to get bullied or pressured or scammed online, that the safety and age assurance are not the same thing. And it won't fix risky product design. So age assurance can enable a platform to do something about their product design, but it won't change it on its own. I'll leave it there.

Diana Chang:

Thanks, Amy. Graham, do you want to respond to the question?

Graham:

Sure. Yeah, it's a great question because that is how we see it. It's age verification. I agree with Amy. It's one of the tools in the tool belt. So there are a number of existing tools that parents have, and platforms are hard at work and our member companies are hard at work and developing meaningful controls for parents. So it includes web filters, restrictions on app downloads and purchases, activity trackers, screen time limits, privacy settings, communication restrictions that usually are sort of inside the app.

So one of the questions that comes up is how widespread is the use of these parental control features? And one report, it was very recent, the Family Online Safety Institute, I think it was 51% of parents reported using parental control features on tablets, and that was kind of the high watermark. So double-edged sword, it's a little higher than I had expected, but it also indicates that we have a little bit of work to do to make sure that parents understand how these tools work, and that policymakers understand what the practical burdens there are. I mean, shared devices, blended household, foster care situations, these are all scenarios are really hard to build a set of controls at the operating system level to capture. And you want to be mindful of it too when you're designing policy.

But I'll point to my own experience too and how it dovetails with that of our members. For younger kids, it is less complicated. On smart devices, when you set up a child's device, it is pretty straightforward. My son's eight years old, I set up his device. He's not allowed to access any apps with a browser. If you embed unrestricted browser access, you're 13 plus under the terms of service. And say those are off limits. He's not allowed to message anybody at restricted access to messaging for apps that are even within his range. He's not allowed to download any of those. His tablet will generate a text message to me, to my phone, and then I will accept or reject, always reject, always reject. And then facial recognition says, "Oh, you are who you say you are." And that's how it works for me right now.

So for that reason, I think it is really important to ensure that you are preserving parental agency, preserving the ability for parents like myself and like others to maintain that control and maintain that sort of flow, something that works for them and works for our member companies too because FlipaClip, for example, I approved my son to download FlipaClip and buy the service, and it's all subject to my control from my device. So it works for the member companies too, they work great inside that sort of ecosystem. So again, part of the reason that we're concerned with proposals that take a one-size-fits-all approach and that take that choice out of the hands of parents and turn it into a government mandate makes it a little bit harder. And it makes it harder to design these software ecosystems that are designed to be responsive to parent needs and an evolving app ecosystem, and instead just have it sort of a stilted statutory framework where consent signal goes to developer, developer receives consent signal. Right now, it's at the operating system level. It works a little bit more seamlessly.

So in all of these scenarios, not necessary for my son to be interfacing with an age verification service very often, I think when he gets older, I think age verification may start to come into play and necessarily he'll have a little bit more autonomy. So many of these more difficult policy questions come into play when you have tweens at issue. So my own experience there and then how that dovetails with the members, just for what it's worth to describe the broader ecosystem and the parental control that's in place that we're using now.

James Trilling:

Thanks, Graham. We're almost to the end of our time. So my final question I'm going to ask for panelists to treat as a lightning round and go very quickly in your responses. Can you address what you would like to see in the future when it comes to understanding users' ages online? And I'm going to start with Emily.

Emily Cashman Kirstein:

Thanks, James. Yeah, at the risk of sounding a bit like a broken record, we're looking for solutions, both technical ones and ones related to public policy that are risk-based, privacy preserving, and address all parts of the ecosystem. And the reason age assurance is so important and the reason why it's important we keep thinking about the future is because it really does, echoing what Amy and Graham were saying on the previous question, it echoes we want to get the right experience in front of the right user. That means benefiting from default settings, like on Google services, personalized ad blocking, take a break in bedtime reminders, blocking things like, again, racy music videos on YouTube. And that's even before you get to parental tools where parents can block and approve apps and websites. That's existing functionality right now.

Managing screen time on devices, setting up school time to completely lock down a phone during school hours if needed. I think making sure that as the technology evolves, that we are leaning into that evolution to be able to provide these experiences. And listen, this is one of the most complex, both technical and public policy issues I've ever worked on related to child protection or otherwise, and it's changing by the minute. So what's possible today really wasn't possible six months ago, and who knows what'll be possible six months from now. So my hope is that as this interconnected network of folks that we've had throughout the day, all looking for ways to do this safely and effectively, that we're all working together on solutions that are having a real impact for kids and parents and also are future-proofed as the technology evolves and improves. Thanks.

James Trilling:

Thank you. And thank you to all the panelists. And I'm sorry that we've run out of time. I'm going to kick it back over to Diana.

Diana Chang:

Thanks, Jim. Panelists, thank you so much for the discussion this afternoon. That brings us to the end of the panel. I'm now going to turn it over to the FTC's Director of the Bureau of Consumer Protection, Chris Mufarrige.

Chris Mufarrige:

Good afternoon. I'm Chris Mufarrige, the director of the Bureau of Consumer Protection. I'd like to thank everyone who made today's workshop possible, including the staff from the Division of Privacy and Identity Protection, who organized today's event and the thoughtful speakers and panelists we've heard from today. It is fitting that we are discussing children's privacy and age verification on Data Privacy Day. This is a perfect time for me to reiterate that there is no consumer protection work more important than protecting our children online. As Chairman Ferguson said this morning, the Trump-Vance FTC is dedicated to vigorous enforcement of the Children's Online Privacy Protection Act, both because of our commitment to enforcing the laws Congress has entrusted to us and because vindicating parents' ability to make decisions about their children's online activity is critical to a flourishing society.

Our COPPA-related enforcement actions speak for themselves. For example, we are taking action against the operators of the Sendit anonymous messaging App for allegedly unlawfully collecting personal information from children. We also settled actions alleging violations of COPPA with Disney, as Chairman Ferguson described this morning, as well as with robot toymaker, Apitor. And we are using Section 5 to protect kids online, such as our action against the operators of PornHub and other pornographic sites for allegedly deceiving consumers about efforts to crack down on child sexual abuse material and other non-consensual sexual content.

As we've heard today, age verification technologies will play an enormously important role in protecting kids online, but currently, using certain AV technologies may be intentioned with COPPA because some technologies require collecting personal information to verify a child's age before parental notice and consent is possible. COPPA, a statute designed to empower parents and protect children online, should not be an impediment to the most child protective technology to emerge in decades. The Commission is exploring potential solutions to this apparent inconsistency between COPPA and certain AV technologies. As we move towards wider adoption of AV technologies, we need to continue to learn more about them. The Commission has long played an important role in encouraging empirical work on technological issues. Indeed, next month, we'll be holding a workshop on injuries and benefits in the data-driven economy that will convene economists, academics, and other experts to examine how the agency can better understand and measure consumer injuries and benefits that may result from the collection, use, or disclosure of consumer data.

We need empirical work related to AV technologies. Some panelists have raised important issues about accuracy, ease of circumvention, and privacy, which suggests that certain AV technologies may be better in certain contexts. As the marketplace adopts AV technologies, the commission needs to understand which methods and which context mitigate these concerns. I encourage those of you listening today and the policy, business, and research communities as a whole to advance empirical work to support the adoption of AV technologies for the protection of children online. Thank you.