[Slide 1]
Hello, and welcome to my talk about COPPA and the Age Appropriate Design Code, first presented at the I A P P Data Protection Intensive Event in London, in March 2020.

[Slide 2]
My name is Carl Gottlieb and I'm a consulting Data Protection Officer, working inside technology companies, to lead their privacy programmes. I'm not a lawyer and nothing I say in this presentation constitutes legal advice or professional advice. It's just one man's observations and opinions which you'll hopefully find useful.

[Slide 3]
Oddly, I'm going to start with describing what I'm not going to talk about in this presentation. I'm not going to discuss the ethics of following the rules, or the morality of doing the right thing by the data subject. I'm not going to talk about where these rules come from or why they're important in the first place. And I'm not going to criticise anyone's privacy practices. I know full well that privacy is not the top agenda for most companies, and I'm not going to pretend anything different or preach to you.

[Slide 4]
This presentation is all about what's going on with child privacy across the Atlantic, with COPPA, the CCPA, the GDPR and the UK's Age Appropriate Design Code. I'll describe some of the challenges that globally focused companies face in weaving these compliance regimes together and I'll discuss some of the approaches you can take. I'll also show you some examples of current practices with handling child data online. I only have an hour, so this is a fast paced talk with plenty of areas for you to dive into much deeper at a later time.

[Slide 5]
In this presentation I'm going to make the wild assumption that complete compliance is your plan. I know this isn't how the real world works, but for me to assume anything less than complete compliance means I'll have to explain which pieces of the puzzle you can cherry pick and which you can try and put off for a rainy day. And I don't really have time for that. And this is aimed at a general audience of privacy professionals, so going into specific scenarios of implementations just wouldn't be practical. Cherry picking compliance is what everyone does, for many reasons, but it's a highly problematic approach. It's like a house of cards, one wrong move and the whole thing can come falling down.
If you're not 100% all in, things can get very complicated. A Privacy Policy is a good example. It needs to say what you do, it needs to be correct, and you need to do what is says. For example, if you want to use cookies to track website visitors without consent, you either claim in your Privacy Policy that you rely on consent which isn't valid, or you say you rely on Legitimate Interests which the law doesn't permit. Either way, you're screwed, and it's the privacy policy catching you out. So for this presentation I'm assuming that you're going to do it all correctly.

[Slide 6]
The most obvious place to start is with COPPA, the Childrens' Online Privacy Protection Act from 1998. It's a slightly strange privacy rule in that at first glance it looks really simple - you need consent to process children's data. But when you read all the FTC documentation, guides and FAQ documents it looks super nuanced and complicated. But then after a while it starts to look simple again.
At its heart is the focus on collecting data from a child, which it defines as being under 13. This isn't about data processing in general, but about directly collecting the data from the child. So collecting child data from a parent wouldn't be in scope. Like many other privacy rules, it has origins in the Fair Information Practice Principles (FIPPs), so it includes the usual things like rights of access, erasure, minimisation, consent, notice and security. As I

said, it defines a child as being under 13 years old, and it applies to personal information collected from general audience services with actual knowledge of the child's age, and it applies to child directed services where the age of the child might not be known but it is expected that they're under 13. And COPPA uses the term Personal Information in a common US standard with fairly rich identifiable information, but this was updated in 2013 to be much wider and include identifiers that you commonly see online.

[Slide 7]
Many people think that COPPA requires parental consent to collect data from a child, but this isn't true. I've created a chart here to show how the data collection activity you have determines the type of parental notice or consent you need to have in place. The more data you collect and then share, the stronger the need for parental involvement.
At the top you can see that if you have a website that only asks for authentication details from the child, such as an email address to log in, you don't need to inform the parent. The same goes for if the child wants a password reminder emailed to them, or they give you their email address for a one-time request such as entering a competition. But you have to minimise how you use that email address so that it can't be used for anything else.
If the child wants to receive ongoing emails from the website, such as by subscribing to a newsletter, then you need to notify the parent. So at some point you'll need to be capturing the parent's email address too. Up till now, we've only been doing the bare minimum data collection and processing, but the reality is that many services want to do more, such as requesting the name, photo and biographical information about the child, and this is the stage where you'll need to block the user and gain parental consent. Where the risk is low and no external data sharing is going to take place, COPPA allows for a consent method known as Email Plus which is where you email the parent asking for consent, and the parent clicks a link in the email to authorise it. Clearly you don't know who is on the other end of that email clicking that authorisation link so it's a fairly low bar to meet. And if you intend to collect any more information that you will share with third parties, such as to provide the user with behavioural advertising or to track their identity with third party analytics, you'll need stronger parental consent. Email Plus won't cut it here. So you're looking at things like a video call with a parent or a credit card transaction to prove the consenting party is an adult.

[Slide 8]
Some of you may be aware that recently YouTube has been having fun with America's Federal Trade Commission, the FTC. This all started with a complaint that YouTube was in violation of the Children's Online Privacy Protection Act, COPPA, and ultimately ended in September 2019 with a vast settlement of $170 million dollars - $136 million to the FTC and $34 million to the New York Attorney General.

But what is unique about this case is how news of its ramifications was most notable not in the privacy sphere, but amongst YouTubers themselves - the content creators that upload their videos onto YouTube. If you want to know about COPPA, or at the very least get some vitriolic opinions on it, just ask a YouTuber.

The reason for this is that as part of the settlement with the FTC, YouTube made some major changes to the way videos are tagged and in turn the functionality surrounding those videos, with many a YouTuber expecting this to be the end of their livelihood.

To explore these changes and why they're necessary, we need to look at where this case came from and why YouTube had a problem in the first place.

YouTube is like many other websites in the sense that its business model is to gather a huge volume of visitors and monetise them by displaying ads. And of course those ads are going to be personalised ads, based on your behaviour all around the Internet to get you clicking

on relevant stuff. And that's how YouTubers, the video creators, make their money. They specify that ads should be displayed around their videos, and so they get more revenue the more their videos are watched.

And like with most websites, personalised ads are a lot more effective than non-personalised ads, so personalised ads is the method of choice to get everyone the most revenue. That's why YouTube is one big hotbed of tracking, with most interactions tracked and used to ultimately drive revenue. Of course, that relies on personal data, or let's use the more relevant term here, Personal Information, since we're going to be talking about COPPA.

At a high level, COPPA has two sides to it.

1. Are you collecting personal information? and
2. Are you doing this knowingly from a child?

So here you've got two get-out-of-jail-free cards you can potentially use for COPPA.

The first being to ensure your service doesn't collect personal information beyond the bare minimum you need to operate the service. That way it doesn't matter what age your users are.

The second is that your service, like most services out there, does indeed collect personal information in excess of what you truly need. For instance you're using cookies and tracking to display personalised ads next to videos. But if you can show that your users are not children, then COPPA can't apply. This was the YouTube approach.

And so just like most other large tech company websites, the terms of service for YouTube stipulates that you have to be at least 13 to register, since that is the minimum age from which YouTube can collect your Personal Information without parental consent under COPPA.

This is where the game starts - trying to stay out of scope of COPPA by stipulating you have to be 13 or over to use the service. COPPA is all about collecting Personal Information from under 13s so if you can show you don't knowingly have any under 13 users then all good right? Well this approach has some merits, but if you're going to take this line, you have to stick to it rigidly. And allegedly YouTube didn't.

[Slide 9]
As an example, here is what I see when I open YouTube.

[Slide 10]
And here are the YouTube channels I'm apparently subscribed to. Clearly someone else has been using my YouTube account. No doubt a small person with access to my devices and is a fan of Mickey Mouse and nursery rhymes.

[Slide 11]
YouTube fell down by allegedly knowing full well that its users were under 13. This was seemingly evident in two places, the first being YouTube's promotional pitches to companies such as Hasbro and Mattel where they made comments such as, "YouTube is unanimously voted as the favourite website of kids 2-12" and "In fact, it's the #1 website regularly visited by kids." Clearly some people at YouTube were very proud of their appeal to kids. Secondly, a huge amount of video content within YouTube is undeniably targeted at young children, from Mickey Mouse cartoons, to nursery rhymes to basic counting songs.

Whilst YouTube may have had an age gate on their registration flow, preventing users from registering when they self selected an age under 13, they allegedly knew that millions of young children still did access their content, whether it be from lying about their age or from unsupervised use of an older person's shared device.

[Slide 12]
The FTC settlement required YouTube make some fundamental changes to get compliant with COPPA, and these changes got rolled out in January.

[Slide 13]
Back to those two get-out-of-jail-free cards on offer, clearly YouTube could no longer rely on the "we don't have child users" card. And this left a fundamental problem for them - how do we stay in compliance if our registered users aren't our actual users? The logged in user might be known to us as 40 years old, but the actual user might only be 4.

The alternative card on the table was to stop collecting personal information from users, and that way it wouldn't matter what age the real user was. But this would be hugely damaging to revenue, requiring all ads to be non-personalised.

It's pretty rare you get a service where the operator doesn't control its content and has literally no idea of the age of its users.

So with no simple solution on the table, YouTube was forced into a "zero trust" compliance model.

The logic goes something like this:

1. We don't want to collect personal information from children.
2. We don't know which users are children.
3. But we do know which videos are directed at children
4. Let's assume that all videos for children are being watched by children.
5. Stop collecting personal information around only these videos.

[Slide 14]
This change requires all videos or their channels to be labelled as "made for kids" or "not made for kids". And YouTube double checks this. Their algorithm proactively checks for labelling "errors or abuse" and amends the videos accordingly.

[Slide 15]
And with the "made for kids" label in place, a whole host of features are disabled, because they all rely on collecting personal information from the user who is assumed to be a child.

[Slide 16]
These disabled features include, comments, donations, live chat, the notification bell, the miniplayer and most importantly personalised advertising.

[Slide 17]
Disabling personalised ads is the big one. Instead, any ads will be context based which are clicked on less, and ultimately produce less revenue for the content creator.

[Slide 18]
Unsurprisingly, many YouTubers were extremely upset by this, as well as the uncertainty surrounding it. One of the biggest questions was how to decide whether a video is made for kids or not. Is a Minecraft video made for kids? Is an unboxing video of a Mickey Mouse

DVD made for kids? Both the FTC and Youtube have since issued guidance in the form of videos and extensive Q&A articles to help clarify. But crucially, the focus from YouTube is about helping you - the content creator, comply with COPPA, NOT YouTube comply with COPPA. YouTube has pushed the compliance focus outside of their four walls, saying they are doing their bit, and opening the door for the FTC to go after content creators for breaking the rules.

We're only a month into these changes, so we're yet to see many metrics of the real negative impact on revenue. Generally commentary from YouTubers since the change has been minimal, so it's hard to know what is going on behind the scenes. With that in mind I reached out to one channel owner that produces their own animated videos for small children to see what impact they had seen. In terms of their size, this channel is pretty big, with 280,000 subscribers and 220 million views.
I'll be honest, I didn't expect a reply, but I did get one, and it was a long one, explaining that they had seen a

[Slide 19]
75% drop in revenue. Their exact words were, "YouTube's changes are catastrophic for us." Being perfectly honest, I'm not overly sympathetic to those YouTubers that produce trashy content such as box openings or episodes of Mickey Mouse played in reverse. But when I hear that some of the very best children's educational content is being put in jeopardy, I am truly sad. Most YouTubers won't have understood the COPPA rule or how YouTube was working within it. They would have assumed that YouTube is a massive organisation that would obviously be fully compliant with the law, and in turn, so would they. At the same time though, these YouTubers were profiting off the collection and profiling of children, in direct contravention of COPPA, which is clearly wrong and needed to stop.

YouTube is still the only game in town for monetising your videos, so people have no choice but to accept the changes and make less money. YouTube has promised to bring back some of the disabled features in the future but in a format that doesn't rely on collecting personal information. Ultimately though, personalised ads will continue to be unavailable when the content is child directed.

YouTube has taken much criticism for agreeing this settlement and the changes they've implemented. But personally I don't see what alternative they had. As a parent I've found the changes to be positive in some respects but I would like to see YouTube do more to support the high quality YouTube creators that produce rich content that is safe for kids. The YouTube Kids app is a good alternative, but it lacks many of the very best children's channels and is very clunky to use.

[Slide 20]
Alongside COPPA, we have to consider the other notable laws that relate to online child privacy such as the California Consumer Privacy Act, the CCPA. This came into effect in January this year and I'm just going to focus on its child related elements.

The CCPA's main focus for children is the application of its Do Not Sell provisions as they relate to child data, or "minors" as they describe them. Similar to COPPA, we're looking at parental consent for under 13s, but the CCPA includes a new bracket for 13 to 15 year olds, who can provide consent themselves, but it has to be opt in, unlike anyone aged 16 and older for whom data selling must have an opt out.

The big point to note here is the under 16 category needing extra protection, rather than the traditional age of 13 being the magic age that a child becomes an adult.

And similar to COPPA again, the rules are based on "actual knowledge" of a child's age. It's worth noting that parental consent requirements within the CCPA are separate to those of COPPA in that the CCPA rules are about data selling, whereas COPPA is about direct collection from the child. So consent for one doesn't provide consent for the other. And this is important since the classic Email Plus consent that COPPA offers isn't available as a method of parental consent for the CCPA.

As many of you will know, the CCPA is infamous for its amendments and moving goalposts.

[slide 21]
And this is seemingly going to continue with its new iteration, the California Privacy Rights Act which will land next year and extend the CCPA.

[slide 22]
From a child privacy perspective, the CPRA doesn't change much at all other than tripling the fines for improper handling of child data.

[Slide 23]
But we do have some COPPA amendments being discussed. There are two main ones that have bipartisan support, with COPPA 2.0 proposed from the US Senate, banning personalised ads to under 13s and individual consent requirements for 13-15 year olds. The PROTECT Kids Act from the House keeps it simple and looks to extend COPPA to apply to under 16s rather than under 13s.
Overall, we're seeing a US trend towards 16 being the magic number, and anything below that is probably going to need prior consent.

[Slide 24]
I've included a quick slide on China, as it's an area that not many of us get exposure to, until the day your CEO tells you they've created an office there and are launching the product next week. What's really interesting is that China has a surprisingly tight regime for child privacy. Their Regulation on Network Protection of Children's Personal Information came out on the 1st of October 2019. Overall it's pretty similar to COPPA, following Fair Information Practice Principles again, but this time there are fewer exemptions to parental consent, and that prior parental consent is needed for under 14s. This is certainly one to look at if China is on your radar.

[Slide 25]
Across Europe we have a mixed bag too, which stems from Article 8 of the GDPR and the ability for each EU country to decide its own age for child consent to an Information Society Service, or ISS.
This picture evolves every so often, for example, with Slovenia changing from 16 to 15 for its age.
We'll talk about this slide again later on.

[Slide 26]
Turning to the UK now, we have the Age Appropriate Design Code, or the "Kids Code" as some have described it. Who knows what nickname it'll be given when it goes live, but for now I'll refer to it as the Code.

[Slide 27]

The Code is what's known as a Statutory Code of Practice. Conforming with the code helps you comply with the GDPR. The Code is a specific thing that the UK is required to have as dictated by the 2018 Data Protection Act, and it solely relates to an ISS that is likely to be accessed by a child. It's underlying focus is about putting the best interests of the child first, to, "protect children within the internet, not from it."
Section 127 of the Data Protection Act states that the Information Commissioner "must take the code into account when considering whether an online service has complied with its data protection obligations under the GDPR or PECR." This means that you must be able to demonstrate conformance with the Code if asked. Ultimately the Code is both a tool to help you comply and the stick to beat you with if you don't. In time there will be certification schemes that align to the code that will help show your conformance, but just like GDPR certification schemes, they don't exist yet, and I wouldn't hold my breath for one to arrive any time soon.

[Slide 28]
Let's look at who the Code applies to. The headline is that it's for "Relevant information society services which are likely to be accessed by children", and remember that an ISS is "any service normally provided for remuneration, at a distance, by electronic means and at the individual request of services." At a high level, that means that we're talking about transactional commercial websites where some kind of offering is being provided.
This means that some websites will likely be out of scope, such as those from public authorities, law enforcement and broadcasters. Static brochureware websites that promote a service elsewhere, such as an appointment booking website for a GP surgery might not be in scope either, since the service in question isn't provided at a distance, which an ISS does. The Data Protection Act also explicitly excludes preventative or counselling services from the Code too.

One big question is whether your service is "Likely to be accessed" by children. The ICO encourages common sense here and advises that it should be more probable than not that a child would access your service. And if you decide that you're out of scope then you should document your decision.

[Slide 29]
The Code focuses on child data processing by establishments in the UK, and any global organisations that target UK children but are not based in the EEA. After the UK's Brexit transition period, the Code will also apply to EEA establishments too.

[Slide 30]
The Code itself has had a few rounds of drafts, consultations and amendments but we now have a final version that is being laid before parliament. This sits for 40 days for the usual discussion and debate, and if there are no objections, which hopefully there aren't, then we wait a further 21 days, and then the Code comes into force. At this point we start a 12 month transition period where the code is in force but not enforced.
So realistically this means you have 11 months to tell your company about the Code, and then one month to actually implement stuff once they see the deadline approaching.
All that means that by Autumn 2021 the Code should be in full effect.

[Slide 31]
The best way I can describe the Code is to "Give Children More Privacy by Design."

[Slide 32]
Focusing on the first part of that, defining what a child is can be a huge problem, but in general the ICO is talking about under 18s here.

[Slide 33]
And the second part is about Privacy by Design. Clearly everyone should already be getting Privacy by Design, after all the GDPR demands it. But it's such a vague term that it's often skipped over and not really taken seriously. The ICO stresses that children deserve special treatment and this code tries to state what that privacy by design looks like. The Code is the "how" of ISS child privacy compliance, albeit it at a high level.
So really, the Code is simply describing what you should already be doing for GDPR compliance, but in a way that is much more tangible.

[Slide 34]
The code has 15 sections. Each are interlinked and flexible, and you need to adhere to them all where relevant.

[Slide 35]
Back to the age thing, I said that the ICO considers a child to be under 18, but how does that marry up with the UK choosing 13 as its age of consent for an ISS? You can see that the UK territories are all aligned on 13.
But to complicate things further, the ICO stresses that organisations treat children as individuals

[Slide 36]
recognising that age isn't always the most important factor. Each child has different maturities, sensitivities and abilities. We also need to consider disabilities and equality legislation too.

[Slide 37]
Whilst the Code has 15 sections, there are some underlying principles that we can pick out and focus on. As you can imagine, with the clue being in the name, age appropriateness is a big one, making sure everything is relevant to the individual child. So just like COPPA, we need to consider the degree of data processing and its associated risk, along with our level of confidence of a child's age, then plan accordingly.

[Slide 38]
So when we communicate with children it needs to be focused on helping them, whether it be with child friendly notices and explanations or letting the child decide how childish they want that communication to be. Back to the theme of the best interests of the child, the Code requires driving child behaviour away from anti-privacy activity, such as freely allowing a child to expose their geolocation.

[Slide 39]
The Code requires privacy by default to be on a high setting for children, but how high and for what children? Minimisation as a principle is important, not just around the data processing but around the service in general. A child might have signed up to an education site, but that doesn't mean they should be enrolled in public sharing of their profile information by default. The Code wants a service to provide the bare essential things that the child needs - focusing on the best interests of the child. And when you do allow them to change a privacy setting, this should come in a non-permanent form, so that the change is only temporary for this active session. The Code recognises the YouTube scenario of shared devices and so encourages services to provide multiple profiles per device, allowing for different age users to have different privacy settings.

[Slide 40]
The Code has a whole section on nudges, which is the name for that practice of pushing the user towards one option or another. Here you can see a nudge that I like, but you might not.

Generally, nudges push people away from pro-privacy choices, such as getting you to consent to something you might not want, like cookies. And as you'd expect, the ICO stresses that nudges remove the freely given nature of any consent. Instead, the Code looks for pro-privacy nudges to drive children to make a safe choice about what to do with their data.

[Slide 41]
Interestingly the Code specifically calls out online tools, or privacy controls, that allow the child to exercise their privacy rights. Here the ICO wants services to actively help the child access these controls, rather than hide them away, and also recognise urgent cases and helping the child stay informed about the progress of their request. Putting the child first is again the cornerstone of this one.

So the Code isn't exactly revolutionary in what it says. It's common sense for most privacy professionals in their understanding of the GDPR. But we've not seen this level of detail before, and had nothing that we can present back to our business that says we actually need to do this stuff.

[Slide 42]
I've covered the latest developments in the rules, but if you're operating across all of these, you'll no doubt be starting to see some of the complexities of how they overlap, interact and contradict each other. We need to find pragmatic ways to make this all work.

[Slide 43]
The first challenge is what defines a "child". In the EU we have the ages of consent for an ISS, but clearly that's only relevant if we're talking about consent as a lawful basis. And in the US, we're starting to see the 13 age getting nuanced, with 16 becoming the new number, depending on when and what we're talking about.

[Slide 44]
The ICO takes this one stage further, remembering about treating children as individuals. The Code suggests these five age stages, which have associated abilities and needs. Clearly if you're offering a service to those that can barely read, it needs to look very different to a service designed for those that are legally able to join the army, even though both groups might technically be classed as children. Remember that even two year olds are watching YouTube independently and able to control touch screen devices with ease.

[Slide 45]
And there's even more nuance to consider. First of all, what do we even call these children? Is it a "child", a "kid", a "data subject under the required age of consent for an ISS depending on your location?" And how do we even articulate this for simple age brackets? Should it be based on consent, or should it be based on contract law and what each country says is the minimum age that a child can enter a contract? Some child privacy laws even vary depending on the child's involvement in higher education.
And because we have all these per-location age differences, the location of the child matters greatly. But location changes and is imprecise, especially when we don't want to invade the child's privacy by collecting an accurate geolocation. A common challenge is knowing whether to base location on the user's IP address location upon first registration and keep this static, or should you apply the age rules based on their current location, or what about letting the user select the preferred location for themselves? For example, does an Austrian 14 year old need parental consent for a service when they're holidaying in Germany where 16 is the minimum?

Many of these problems are not new, but the YouTube COPPA example shows that we should be reevaluating our assumptions. Ordinarily, the legal attitude would be to create a defensible position. But YouTube had that, and it cost them 170 million dollars. So rather than a defensible position, a "child friendly" position would be better, and simplicity is the key to getting to that. For example moving to a global standard of 16 removes many of the problems on this slide.

[Slide 46]
Yes, this is the most sinister title for a slide ever, but the question of whether you're targeting children is a critical one. Depending on your service, it can be very hard to know if the person you think is your user is actually your user. It's often not the case for YouTube. So the content itself can be the best indicator, along with the goals of the business, as to whether you are appealing to children. If your business wants more child users then you're likely trying to appeal to them. Often your assumptions will be wrong, so test them with your colleagues, consult with children and parents and make sure you document it all. The ICO states in the code that for DPIAs, "We will expect larger organisations to do some form of consultation in most cases."

One mistake I have seen is a business claiming that it didn't target children, but accidentally their PR team had placed their app in the "Designed for Families" section of the Google Play Store. No business is perfectly aligned, and you might have colleagues that are accidentally out of sync with your approach.

[Slide 47]
Age Gating - it's something you're thinking about after you read the Code and COPPA, but believe me, it's something you want to avoid doing at all costs, and so does your business. But just like COPPA rules on consent, there's a progression of age knowledge all the way from no idea to that painful verified parental consent.

The less knowledge you have, the easier it is for your service to register users and easier for the user to get using your service. But clearly you create a challenge of not knowing how much these child privacy rules apply or how you can tailor anything to the different needs of the child.

At the other extreme we know we have consent from the parent, which helps us comply, but we still don't really know which user is the actual user, and the burden of a hard wall age gate is massive. Services see a huge drop off rate when even the most basic Email Plus parental consent is requested, and verified forms of consent are significantly worse.

But you have an obligation to do this right, and a middle ground I would promote is self selection of age. It's a nice sweet spot, allowing you to tailor information to your audience, stays away from parental consent and helps both you and the user. It also keeps you away from knowing date of birth, so if you intend to graduate children to adult users when they reach a magical age, such as 16, then you have to accept there will be a lag and potential complaints with you thinking that a user is younger than they actually are.

[Slide 48]
If you do want to go down the parental consent route, then COPPA has these requirements, with the potential for Email Plus - the click a link on the email method, if there's no data sharing. And the CCPA has similar accepted methods for data selling which obviously would preclude the COPPA Email Plus method.

It's worth saying that there are third party systems out there for verifying identity and age. Some are actually very good. But even the very best of these has a massive amount of

friction imposed on the child or parent, usually requiring them to register for the service before hand and prove their identity.

When attention spans for children and even adults on web pages are measured in milliseconds, any friction from parental consent and verification is bad and you're going to lose users.

[Slide 49]
Sony's PlayStation network is a good example of age gating. You can see here that I'm trying to register, and pretend that I'm one year old. It's blocking me, telling me I don't meet the requirements. So straight away there are some age requirements going on.

[Slide 50]
So I'll change the year of my birth and see what happens. I'm now 9 years old and it lets me proceed but says we need parental permission. And contrary to advice given in the Code, we have a back button here to allow me to go back and try again.

[Slide 51]
I've now changed my year of birth from 2010 to 1970, and now it allows me register without any parental consent. An easy bypass, and a tried and tested approach that any child knows.
Ideally, if you're going to have a gate like this, accept the user's first age selection and don't let them go back.

[Slide 52]
Probably the hardest problem in all this is the lawful basis for processing the child data. In the GDPR world we have six options, but only three of them are available for an ISS, Consent, Contract and Legitimate Interests.

[Slide 53]
And if we bring the US into the equation where lawful basis isn't really a thing, this comes down to something either needing consent or just being allowed, with lots of shades of grey in the middle.

[Slide 54]
Back to the GDPR lawful bases for the moment, we're immediately faced with some problems.
Consent is an awful word in that for 99% of people it just means agreeing to something, but in the EU we know it must meet a specifically high bar. Whereas in other countries the definition varies wildly. And sometimes, consent is absolutely required, such as for some eprivacy things like online behavioural advertising where we don't get a choice in the matter.

But for an ISS, it would be inappropriate to use consent for all the processing activities, especially since many children are too young to provide consent. But core processing shouldn't be based on consent anyway. It's a required function, so consent wouldn't be appropriate.

Contract seems an obvious alternative for the core processing, but then you run into the issue of whether a child can even enter into a contract. Some places have strict rules on age requirements for a child to enter into a contract, some do not. And to use Contract as a lawful basis we absolutely need necessity for the data processing, which in many parts of an ISS we don't have.

And lastly we have Legitimate Interests, which as a term rarely exists outside of the EU, but as a concept of doing what's fair and reasonable, we have a high burden to do what is in the best interests of the child. But this is what the Code is all about, and since Legitimate Interests doesn't have any age restrictions attached to it, this is likely a more appropriate model to be operating under.

[Slide 55]
As I said, sometimes you will need consent, and depending on the age of the child and what you're doing it might need to be parental consent. COPPA needs it for data sharing, as does the CCPA in its current name of data selling and the CPRA which calls out data sharing alongside data selling. And of course things like cookies and direct marketing components from the ePrivacy Directive probably need consent too, with behavioural ads being the perfect example. The ICO classes behavioural ads as being high risk and explicitly calls them out as needing consent.
So all of these pretty much amount to the same thing, which is nice to see. Anything non-essential that usually involves sharing data with third parties is going to need consent.
So back to our YouTube example, less is more. You can free yourself from the burden of consent by stopping doing all this stuff.

[Slide 56]
Just a quick note about consent in the GDPR versus the ePrivacy Directive. The GDPR has a minimum age for consent as a lawful basis (for an ISS) for processing personal data, for example in the UK it's 13. And the ePrivacy Directive reuses the GDPR's consent definition, with its "freely given, specific, informed and unambiguous indication" requirements. But crucially the ePrivacy Directive does not bring in the GDPR's ISS consent age minimum. So technically, this creates a narrow exemption where a child can provide ePrivacy consent on an ISS as long as no personal data is processed that would require the use of the GDPR consent lawful basis.

[Slide 57]
Now this sounds quite interesting, but unfortunately it's somewhat irrelevant since virtually all the cookie, tracking and behavioural ad stuff you want to do that requires ePrivacy consent, also involves processing of personal data and sharing it with third parties. So you'll still need the consent lawful basis.
But how?

[Slide 58]
With children and mixed audience websites it's a real problem.
It breaks down like this:
I don't know the age of my actual user, but I want to serve them behavioural ads or tracking. I know I need their consent, so I serve them a cookie banner. The user agrees and I start tracking them.
But if the user is an under 13 year old child, do I actually have consent from them? By definition, I can't. They're not old enough to give it. Am I now unlawfully processing and sharing their data?
And what about tracking pixels in emails I send to children? I won't have valid consent for that either.

With a mixed audience, is a partly "bad" consent the only answer? I'll show you some examples of this in a bit.

[Slide 59]
Just to make things even worse, there is a problem with ad funded websites and consent that nobody ever talks about, and it's that even non-personalised ads still need consent for

all the end user device processing they do. This comes from the tracking and analytics that the ad networks like Google need to do to measure the effectiveness of ads and prevent fraud. COPPA actually has an exemption for this, but the ePrivacy Directive doesn't and even Google calls this out as needing consent on their instructions for enabling non-personalised ads.

Whilst the simple answer might seem to be, "well okay, just ask for consent to show non-personalised ads", yes this solves the compliance problem but we now create a situation where users can just opt out of all ads, personalised or not. And your business would not be happy.

A workaround to this is the "cookie-or-pay wall" approach, where you have a free service which includes ads and a paid one that doesn't. The conditionality nature of this would at first seem problematic under the GDPR's requirements for consent, but regulators are somewhat open to the idea. A ruling from the Austrian DPA in November 2018, and the Dutch DPA's analysis of cookie walls in March 2019 support the idea. In a Q&A page on their website, the Dutch DPA said that if someone refuses cookies, "Then you still need to give this person access to your website or app, for example after payment." They are very specifically stating that payment could be an alternative to consenting to cookies. And the latest drafts of the ePrivacy Regulation show some support for this approach too.

But for many websites, this cookie-or-pay wall approach isn't viable or desirable, so virtually everyone just sweeps this issue under the carpet and assumes that non-personalised ads don't need consent.

[Slide 60]
Rules within the EU ePrivacy space are always in a state of flux, or to be more precise, the laws aren't changing yet but the guidance is, with multiple regulators releasing their own guidance on what they think ePrivacy compliance looks like. They all seem to be converging on strict standards for consent, as you'd expect. But enforcement is still very patchy, and only a few organisations across the EU have been fined for improper ePrivacy consent within cookie banners and the like.

Nudges have been getting a lot of attention in the regulator guidance, and likewise within the Code, ensuring consent is freely given, as it should be. Even the CCPA includes fairness provisions with requirements to notify children that they can later opt out of any selling that they have opted into, along with a prohibition against asking people to re-consent again within 12 months. If we finally see the EU regulators enforce against fairness infringements then marketers will need to watch out. The industry of A/B testing user behaviour to optimise forms, content and consent acquisition will be hit hard.

And then we have the ePrivacy Regulation, which will one day happen, maybe not in my lifetime but I'm sure it will arrive. The latest draft includes legitimate interest as a basis for some end user device processing which is a radical departure from consent, but a highly sensible one. However this excludes children so it's not very useful for us here. And we have also seen a potential allowance for cookie-or-pay walls too.

[Slide 61]
Consent might be hard to get for a mixed audience, but most of the Code is about doing Privacy by Default, fairness and transparency. So the best place to start is by just complying with the basics, such as strict compliance with the ePrivacy Directive. As privacy professionals we know where the laws aren't enforced very well and what defensible legal positions we can take, but if you're sincere about it, you know full well that your business is taking a risk.

I have some examples here of websites that would likely appeal to children.

[Slide 62]
The first is Disney's website.

[Slide 63]
You open the page, and almost instantly you get their cookie banner with a clear nudge to get you to consent to all their tracking stuff.
Naturally I gave it the sub millisecond moment's thought and clicked yes.

[Slide 64]
And now when I go on Twitter, even many weeks later, every ad in my timeline is for Disneyland Paris.

And now I'll show you some toy shops, and just like Disney, these are all big favourites of my family so I was pretty hopeful they'd be doing things well.
[Slide 65]
Galt is a big British toy maker.

[Slide 66]
And you can see that it's definitely geared towards the younger child market. Now this time they have the old fashioned cookie banner at the bottom.

[Slide 67]
Here it is close up, with the usual, "we do stuff and you're already being tracked, deal with it."

[Slide 68]
And you can see this by using a tool like Ghostery to see what tracking is active.

[Slide 69]
And here it lights up like a christmas tree. Remember - this is a mixed audience website that will clearly have a child audience.

[Slide 70]
So at this point I'll go to close the webpage, and then

[Slide 71]
Oh!, They'd like to offer me a coupon!

[Slide 72]
And here's that close up. So again we have a nice big nudge, and of course we press the big colourful button to continue.

[Slide 73]
It asks for my email address, so I put it in. And because I'm a smart child I don't fill in the consent option and just press the big green button again to get my code.

[Slide 74]
Oh, that consent option was mandatory. Okay, I'll leave the questionable conditionality of that consent for someone else to ponder over and just consent so I can get my coupon.

[Slide 75]

I'll enable that.

[Slide 76]
And now I get my coupon. But I need to confirm my email address? For what? I'll just ignore that thank you very much.
So again we have some questionable consent and nudges going on.

[Slide 77]
Here's a major UK toy shop named Smyths Toys so we can see what their website looks like.

[Slide 78]
And straight away they want to know my location for some reason and they have another implied consent cookie banner.

[Slide 79]
Here it is up close.

[Slide 80]
And here's another big British toy shop named the Entertainer.

[Slide 81]
And their cookie banner is….. nowhere to be seen.

[Slide 82]
But they do want to know my location as well, and they want me to sign up to their newsletter.

[Slide 83]
And if I look hard enough there actually is a tiny link to their privacy policy which might explain something about that newsletter.

Now you'd assume that if there's no cookie banner they're obviously appealing to their child audience and not doing any tracking right?

[Slide 84]
Well no, all the usual retargeting, tracking and Facebook integration is there. It's all on by default, with no ability to turn it off.

So that's three major toy retailers and all of them are some way off the standards the Code is looking for.

For my final example I thought I'd go for the absolute extreme. And some of you can probably guess who I'm going to pick.

[Slide 85]
TikTok.

[Slide 86]
For those of you that aren't down with the kids, or you were just born before 2010, TikTok is a video sharing app for very short videos. There's pretty much no content moderation, so it's full of stupid, inane, dangerous and offensive rubbish that goes viral on a daily basis. Kids

love it. And as a Chinese owned app, it has a reputation for being less than perfect when it comes to privacy. It was even subject to a $5.7 million dollar settlement in 2019 for COPPA violations. So I wasn't expecting much here.

For ease of comparison I focused on their website. I ignored the cookie banner and looked at what nastiness was loaded by default.

[Slide 87]
Nothing. Literally nothing.
I assumed that something wasn't working in my browser so I tried again and there's just nothing there. It's as if they are following the rules of the ePrivacy Directive. Yes, actually obeying the cookie law.

[Slide 88]
And when I clicked on the cookie policy in the cookie banner to see was going on, they even have a nudge free set of options. I was pretty shocked.

[Slide 89]
Turning my attention to the privacy policy, they describe how users under 13 are given a different experience of the product.

[Slide 90]
And then you can see that only limited information is collected from these children. So clearly there's going to be some age gating going on.

[Slide 91]
And you can see here, that for these children some minimisation of the service will be happening along with a reduction in data collection. It's pretty evident what is going on. TikTok are sticking to a hard line of age restricted service, solely based on the registered user being under 13, and nothing else.

[Slide 92]
Compare that to the YouTube example, TikTok is still assuming that the registered user is the actual user. So TikTok isn't taking into consideration the child directed nature of the videos and the shared device scenario that YouTube has.

[Slide 93]
The bigger issue TikTok face is not a data protection one but a content moderation one, and already we've seen the press focus heavily on the kinds of videos that parents would be shocked to learn their children are watching.
[Slide 94]
This raises the question of risk. Is bad press actually that bad for TikTok?
If anything it'll give them more notoriety and popularity.

As data protection people we like to tell a good story that breaches and incidents and bad press coverage will hurt your reputation but this is rarely the case. Any dominant company, whether they be YouTube, TikTok or a big toy shop isn't going to lose any customers over somebody complaining about their lawful basis for processing child data.

Even regulatory enforcement is unlikely for most companies out there, until you get to the size where privacy activists and child welfare activists and probably some politicians too get involved and demand a regulator take action. Personally I expect that all the attention TikTok are getting for their content problems might lead to some questions on their privacy side.

In terms of real negative impact, what's more likely with privacy screw ups is tangible problems when working in the Business to Business world. A good example is in the education tech space where a whiff on non-compliance will get you thrown out of any bidding process and potentially cost you big money.

And lastly we need to think about the liability that individuals carry. There's a mistaken belief amongst DPOs that we're special and protected by law from any flack being thrown our way. And that's just plain wrong. If you're bad at your job, don't expect to stay in it. And the same goes for anyone after a data breach. Yes, a massive data breach didn't sink Equifax or Cathay Pacific or TalkTalk. But in my experience the one thing that does always happen after a big breach or fine is somebody gets fired. Maybe it's someone in security, maybe it's the CEO or maybe it's the marketing manager that assured the business that consent isn't needed for sending those million text messages. So whilst you might think that legally you don't have any liability attached to your role, ultimately we all have liability for following the law, doing a good job and not being to blame when things go wrong.

Just one final point on getting people motivated to do the right thing. People are always happy for the company to take risks when it's the company carrying the liability. So the next time someone in the business wants to take a big risk, make sure to tell them that for your documentation, you need to record the name of the person in the business that is signing off and owning this risk. It's interesting how often that risk acceptance suddenly disappears.

[Slide 95]
There's a lot to consider when your scope is global and children are involved, so here are my four top things to focus on.
Firstly, assume Global Convergence of privacy rules. Aiming for the minimum bar per territory is a nightmare and think ahead to future expansion plans for when you might move into territories you've not yet considered. Countries are rapidly converging on GDPR level Fair Information Practice Principles and age standards so focus on simplicity and consistency with the GDPR and COPPA as your backbone.

Align to the Business. Privacy by Design needs true exec buy-in as it's going to involve turning off features and functionality and stripping down the child's experience of your product. This goes against everything growth marketers and product people want. It'll potentially mean less users and less data, so try to convince marketing that since children can't really buy stuff, it's better to go where the money is, and create a safe environment for children and focus on purely adult data for targeting.

Avoid Consent. Consent is a bit of a nightmare and it's often not really what you need to get the job done. Minimising the service like YouTube has done, will help remove some of the need for consent, but places like ePrivacy often still require it. And with mixed audiences, accept that you'll end up getting bad consent for some activity like acceptance of cookie banners.

And lastly, take a phased approach. Do the obvious first - just following ePrivacy rules would be a good start and you don't really have any excuse for not getting that right. You'll need to create a plan that chunks the implementation of changes of practices, policies and support processes. This ensures they will all match up to what you say you do in your privacy policy. And plan for the end game. Don't trap yourself into relying on consent when later on you plan to rely on a different lawful basis.

[Slide 96]
You've got a lot to do, and less than 18 months to get it done. So I'll leave you to it!

Thank you for listening.