The Cookie Collective

Cookie Collective
133 Creek Road
London SE8 3BU
Tel: 0203 176 2920
www.cookielaw.org

Five Models for Cookie Law Consent

Contents

Summary

First of all, this document is deliberately long and detailed. I accept that many will find it difficult to justify wading through it all; however I make no apologies for that. In fact despite its length, many things have been left out that could have been included.

This has not been written for those primarily interested in either the letter of the law, or the technical and user experience design considerations involved in complying, though such readers will hopefully find much of interest.

The main audience for this document is website owners who need to understand the core choices open to them for getting consent for the use of cookies. These are fundamentally about the amount of control you decide to give to your visitors, and when they are able to exercise that control.

Each model has some validity, though some will be less acceptable than others in certain EU countries. Compliance however is rarely a black and white issue. If someone says to us 'just make me compliant' what they often mean is 'I don't want to get a fine'. The solution to these statements does not have to be the same.

As much as anything, it is about balancing risks: the risk of enforcement; the risk of a loss of information; the risk of alienating your visitors; the risk of competitive disadvantage.

When choosing which model of consent to apply to any website, you need to consider what the law says, guidance from the applicable regulator(s), what your visitors expect, what your competitors or peers are doing and what your brand values are.

All of these can have an influence on the model you choose and how you execute that in the user experience.

I hope that this article helps you to decide which model suits you best.

About Us

The Cookie Collective is a specialist provider of software, advice and services to help website owners comply with the cookie laws.

We have developed a range of software products, marketed under the Optanon™ brand at www.cookielaw.org and www.optanon.com. We are also responsible for the Cookiepedia open knowledge base all about cookies: www.cookiepedia.co.uk.

Our ePrivacy Centre (http://www.cookielaw.org/optanon-eprivacy) is a Software-as-a-Service solution to enable website owners to meet the requirements of the cookie law across the EU, whether they have one website, or hundreds.

Combining flexibility with control, the Optanon™ ePrivacy Centre, not only supports all the models described here, it allows website owners to switch from one to the other at any time and within a few minutes. The user interface is also of course fully customisable for individual branding.

Most importantly, when you buy a licence from us, you are also buying our commitment to keep updating it to suit the ever changing regulatory landscape.

We offer a free 30 day trial for any website owners. To find out more please visit our website, or call us on **0203 176 2920**.

Introduction

The EU cookie laws have been in place since 2011, although there was very little awareness before 2012 when the first websites in the UK started displaying consent notices. In most countries we have seen what can largely be called a ‘softly-softly’ approach to enforcement, however activity has picked up since mid-2014. In the UK the Information Commissioners Office (ICO) has kept it a low profile issue, but Spain and The Netherlands have seen enforcement action and fines for non-compliance.

In September 2014, the French regulator the CNIL led a series of ‘cookie sweep days’ to assess the state of play amongst big websites in particular. Then in October they began exercising powers to run remote compliance audits, with follow up enforcement action to follow closely behind where needed.

In the UK, many sites have implemented some kind of ‘cookie banner’ and introduced ‘cookie policies’. A lot of sites went for an approach of doing as little as possible and then waiting to see what happened next. When that turned out to be not very much many site owners left whatever they put in place alone, and stopped thinking about the problem.

However, as websites inevitably get re-designed and user experience journeys get re-evaluated, many site owners are now looking at their initial solution to the problem, as well as what their competitors have done, and begun to ask themselves ‘could we do it better?’

When looking around for up to date guidance however, there is very little information available about the choices to be made. Much advice is now several years old and amounts to nothing more than a re-statement of the requirements of the law, ‘*tell people about cookies, and get their consent*’.

This is hardly very helpful, so we have decided to put together something more comprehensive, and lay out some clear, practical choices for site owners and designers to make.

Although a cursory glance seems to suggest that everybody is doing basically the same thing, if you look at little closer, and look at the requirements laid down by different regulators, there are in fact no less than five broad models for cookie law consent.

These can be characterised as:

1. Information Only
2. Implied Consent (opt-out)
3. Soft Opt-in
4. Explicit Consent
5. Mixed Consent

In addition to these 5 there is also the option of giving extra consideration to Do Not Track (**DNT**) requests.

For the remainder of this article, I am going to look at each of these in more detail.

What We Talk About When We Talk About Cookies

It is important to realise that the ePrivacy Directive from which the cookie laws derive, applies to more than just HTTP cookies. Other technologies that perform similar functions to cookies, such as web beacons, etags, Flash locally stored objects (often referred to as Flash cookies), and HTML5 local storage, are also subject to the same legal requirements.

However cookies are the most common form of local data storage and tracking, which is why this became known as the Cookie Law. Within this document, as in our other advice and publications, when we talk about cookies, we are using the term to encompass these and any other technologies covered by the same regulations.

The Strictly Necessary Exemption

There are many different ways that cookies are categorised by websites, usually by reference to their purpose. We believe that the UK International Chamber of Commerce (**ICC**) categories are the most helpful for consumers, and have long supported them. They are: **Strictly Necessary**, **Performance**, **Functionality** and **Targeting/Advertising**. Most sites now use these or a close variation of them.

However, only one of these categories carries any specific significance in the law, which is the **Strictly Necessary** category.

Cookies that are strictly necessary for the functioning of the site are exempted from the requirements for consent under the law. This means that any such cookies are outside of the consent models discussed below. Strictly Necessary cookies need no controls to be applied, and can be set as needed.

However, it is important to note that the definition of what is a strictly necessary cookie is very narrow – and cannot be applied more broadly to suit the business needs of a particular site. For more detail of when cookies can and cannot be categorised in this way, the Article 29 Working Party has issued an opinion document which serves as a general guide: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp194_en.pdf

Although consent is not required for such cookies, it is considered general good practice to identify them so that people can distinguish them from other types of cookies if they want to.

However when discussing the different models and introducing user controls over cookies – Strictly Necessary cookies can be excluded from such decisions.

The Granularity of Cookie Controls

In all of the models we present, except the first one, visitors are given controls over whether cookies are set or not. However this does not have to mean a binary all or nothing choice.

The law itself is silent on the issue of how much control users should be given. In theory you could choose to give users choice about each individual cookie. However from a practical, as well as a user experience perspective, most site owners would find such an approach unworkable.

In our view the most sensible approach is to apply controls at the broad purpose category level, such as the ICC categories mentioned above.

Model 1: Information Only

In summary: By visiting the site, you accept our use of cookies.

Basically this model tells the user that cookies are in use, and their choices are to accept the fact or navigate away.

This is technically the simplest approach, and the most widely adopted. It requires the least amount of effort and change to a site. It can be done well, but it is also very easy to get wrong.

The amount of information provided by sites using this approach can vary a great deal. Some have nothing more than a short statement in some kind of banner, with variations on the above sentence, and a mechanism to remove it from view.

Others will go further by perhaps linking to an internal cookie information page that says a bit more about the cookies in use.

Good Practice Tips:

With this approach generally the more information you provide, the better. Avoid generic statements about cookies that have been copied from other sources. Instead, list out the types of cookies your site uses, and explain their purpose. This shows you have at least given some consideration of the privacy implications for users.

Be open and clear about why you choose to use cookies that are not absolutely necessary for site functionality. Some visitors may not like it – but providing less information, or misleading information makes it appear you either don’t know or care less than they do. It is better to provide more information than most people will care about, than not provide enough for the one person who may decide to make an issue of it.

Good practice also includes advising people of the ability to delete or prevent cookies in their browser or by other means. However, don’t try to go into all the detail yourself.

Instead it is better to link to external sources which are more likely to be kept up to date, such as: http://cookiepedia.co.uk/how-to-manage-cookies

Mistakes to Avoid:

Don't give visitors an illusion of choice. A lot of poorly designed information banners ask people to 'Accept' the use of cookies with a button or tick box which makes the banner go away. However, if you have already set cookies, and there is no equivalent option to delete or refuse, this is misleading for users, and more likely to get you noticed, resulting in complaints or regulatory action. If all your button does is close the banner, label it that way.

We have talked to a number of webmasters with non-functional 'Accept' buttons who say that almost none of their visitors click the button. This is then interpreted as an indication that users don't care about cookies and privacy. However it is just as easy to suggest that they recognise it as a false choice that they don't need to make, that they don't want to accept cookies and would refuse them if the choice was there. A site with this approach can also open itself to accusations of deliberately misleading users into thinking they wouldn't get cookies if they didn't click 'Accept' – which is much more likely to lead to stronger enforcement action from a regulator.

Another common bad practice is to have a banner that disappears of its own accord, before the user has a chance to read it. This is particularly bad if there is no obvious way they can retrieve it, or there is no other information in a cookie policy page. From a user's perspective this can feel like another attempt to mislead.

There is nothing inherently wrong with the disappearing banner. However if you are going to set it on a timer, like a fixed number of seconds, make sure they cannot fail to notice it on first arrival, and that there is enough time for slower/inexperienced users, to read and interact if need be.

A slightly better form of the disappearing banner is to leave it in place until the user takes another affirmative action, such as navigate to a second page.

The forced click or acceptance, i.e. not allowing use of the site without clicking on a cookie notice, is probably the biggest mistake for the information only model. It is rarely seen – but there are examples. Some sites have even offered a refuse option which re-directs the user elsewhere, like to a search engine page. This solution is not only high risk – as people may just leave the site, but it also is highly deceptive. If visitors click through to get to the content, they have not given any valid consent, and if they leave they will assume that they have not had any cookies from the site – which is not the case in the information only model.

Risk Factors:

This model, especially if information is kept to a minimum and hidden very quickly, is the least likely to be considered compliant or fair by any regulator. However it can be useful as a quick fix or stop gap while an alternative is being decided on or developed. In the short term many visitors may find it acceptable, but if more of your competitors go further, you can risk losing trust, and visitors.

It can give the impression that you don't trust your visitors to make choices that benefit you, that you don't care about the law, or you don't understand your legal obligations, it can also make a site look unprofessional.

Giving people a banner with neither choice nor adequate levels of information can also be viewed by visitors as showing you have something to hide that they might object to, even if you don't.

The reality is that you are unlikely to get into immediate trouble with a regulator, but if your site is investigated, you are more likely to be required to change it to meet the legal requirements. It is a good idea therefore, even if you choose to go with this approach until you get a complaint, to have an alternative ready to roll out quickly.

Model 2: Implied Consent

In summary: We are using and have set cookies, but you can switch them off if you want.

The key differentiator to the Information Only model is that the site provides the ability to directly opt-out or refuse cookies, even though they are set by default on first arrival.

Good Practice Tips:

When offering opt-out controls there is balance to be struck between usability and the effort required to opt-out.

Creating options for different levels of opt-out is good. Best practice suggests grouping or categorising cookies by purpose, and giving control at that level – perhaps over 3-5 different categories. This gives real choice to users, without it being too difficult to choose, or too many clicks to make.

It is also a good idea to explain the consequences of opting out, especially if it may negatively impact user experience.

If people do opt-out it is also perfectly acceptable to incentivise them to opt back in again later. You may find you have to block certain functionality when users opt-out of some types of cookies. When you do this, make it clear that this is the result of their choice – and you can then enable them to opt back in, if they want to use the blocked functionality.

If you want to there is nothing to stop you preventing access to premium/valuable content or services after users have opted-out, even if these do not rely on cookies to work, but at the same time don’t try to trick users into opting back in.

It is also a good idea to make sure that the controls to opt-out or in again are always readily accessible to the user, such as a page element or link that is on every page, and is clearly identified.

We also believe that implied consent can also be done without the need for banners or pop-ups that automatically appear when users first arrive on a site, although this may not be consistently true in all jurisdictions.

As users become familiar with the concept of being able to control and opt-out of cookies, those with an interest in doing so will automatically seek out access to information and control mechanisms. As long as links or buttons are easily identifiable, always available, and offer real choice, there is less need to interrupt the user experience which many visitors find annoying.

However, if you do want to introduce a banner message, an approach that works well with implied consent is the banner that automatically disappears after a period of time. It works to tell users you are complying whilst not requiring an action to get rid of the message. As long as there is another always available link to the opt-out controls, this can be an additional assurance that you have given clear notice to visitors.

Mistakes to Avoid:

Probably the biggest mistake we see is confusion between the Information Only and Implied Consent models. As noted above a lot of sites try to use the language of implied consent in an information only notice, but implied consent notices can also be easily confused with explicit consent.

An opt-out mechanism will inevitably require you to make some technical changes to your site, in the event that users choose to stop cookies being set. It is vital therefore that you put those changes in place and test them. If you are using a pre-built script or service, make sure you read the documentation, and where necessary involve your web developer. This includes making sure you understand the requirements for using such a script before you purchase or licence it. Giving your visitors the appearance of choice when their choices don't actually do anything can easily be seen as being deliberately misleading, which is clearly something to avoid.

Try to avoid forcing people off-site or requiring them to install third party tools to exercise their opt-out. Some opt-out mechanisms do this – for example requiring installation of browser plug-ins for Google Analytics. It can seem like an easy option, but has some significant drawbacks.

It is not only annoying for users, but it puts control into the hands of a third party rather that you. It also means that you may not be able to incentivise users to opt back in again at a later date – which could be critical for some businesses. If relying on the installation of third party software – some users may not be able to do this (for example users at work may have had this disabled by their IT department) – and therefore cannot exercise their rights properly. Plus of course, if you direct people off-site, there is a significant risk they won't come back again.

You don't need to worry about deleting cookies already set if users opt-out. Technically this is more challenging to do, especially with third party cookies. Opting out means stopping reading existing cookies, however if you use the right mechanisms to stop setting of new cookies, this will also prevent reading if existing cookies, which is consistent with the implied consent model.

Risk Factors

Implied consent is potentially the least user-interruptive model for compliance, if done in the right way. It can give real choice without getting in the way of the user journey for those that are genuinely not interested in exercising their choice.

Not only is it considerably lower risk from a regulators' point of view than the information only approach, it also shows respect for visitors who want to exercise control.

Research suggests that web users are employing more privacy defensive technologies than ever before. Though this makes life easier in the short term for site owners, it has bigger long term consequences in terms of loss of control over your own websites. It is much better to try to prevent that by using a little more effort to give users choices that you can remain ultimately in control of.

Model 3: Soft Opt-in

In summary: We will use cookies if you continue to use the site.

Soft opt-in can look a lot like Information Only, however the crucial difference is that cookies are blocked on first arrival to the site (the landing page). Any further user interaction, such as clicking on a link to a second page, is then taken as consent, and cookies are then set normally on the second page.

There is however an exception in that if the first user action is to follow a link to more information about cookies to be set, this cannot be seen as consent, so the cookie information page should not itself set cookies until a second action is taken.

It can be technically quite challenging to get this model right; however it is mandated by some regulators, notably the CNIL in France, as the minimum compliance level.

Good Practice Tips:

Getting the content and format of the message correct is critical to this model. It needs to be clear to users that they have a chance to not accept cookies before they continue, so this does mean an initial notice must be prominent on the landing page, and it must stay in place until the user takes further action. The content of the message should also spell out the choices clearly.

A layered approach to messaging works well with this model, especially if you can present sufficient information for the user to make a choice, without navigating to a new page, which will simplify the technical implementation.

It is also a good idea to test thoroughly when implementing this approach, as it is easy to get wrong and end up misleading visitors.

Mistakes to Avoid:

You don’t need to stop people entering the site unless they have clicked to accept cookies, so don’t go overboard with something like a page-takeover approach or ‘cookie wall’ as mentioned below, it is unnecessary and can end up losing you visitors, as many will click away.

The time limited, automatically disappearing banner is definitely one to avoid in this model. Even if there is another mechanism that remains, it could be seen as confusing to visitors.

Similarly even if users have continued and allowed cookies to be set, there needs to be an always available control to opt back out again. Omitting that option would not be seen as acceptable by regulators requiring this model.

Risk Factors:

Generally from a regulatory perspective this approach is pretty low risk, as long as you get it right. The main issues are errors in implementation that result in behaviour that is different from messaging. If this results in unexpected cookies it could easily be interpreted as deliberately misleading.

Making opting out significantly more difficult than opting in could also be viewed negatively by both visitors and regulators.

Another issue with the soft opt-in model is how long the consent can be valid for. In France there is a requirement that consent not be stored for more than 13 months after which visitors would need to be given the choice again.

Model 4: Explicit Consent

In summary: Please click to accept cookies on this site

With this model you have to block cookies until users perform a specific action that signifies their acceptance of cookies. The action should only signify that acceptance. Essentially this means they have to tick a box or click a button or a link that says ‘I accept cookies’ or something very similar.

Technically it need not be any more difficult that an implied consent model, and could actually be much simpler to achieve.

However the greatest difficulty can be in getting people to click on the accept link, without completely disrupting the user experience. In The Netherlands where explicit consent was initially adopted, many sites erected what became known as ‘cookie walls’. These forced users to accept cookies before they could actually get to the site. After which the rules were softened up the rules a little, although some types of cookies still require explicit user consent.

Good Practice Tips:

Getting this model right is mostly about considering the overall impact on the user experience of not having cookies set by default.

One option is to go for the cookie wall and force users to accept cookies before they can access the site. If you are going to do that, you must make sure they can access information about what they are accepting before they do so. This means a cookie policy page that also has cookies blocked. Without this it can be argued that consent was not informed – which makes it invalid under the law.

The cookie wall can work well for recognised brands with very strong or unique content, but even then expect drop offs in user numbers. For less compelling sites, this should generally be avoided.

For most modern sites running without cookies means there will be page elements and functionality that will need to be blocked unless or until users accept their use. This is actually an opportunity to get consent and a good approach here is to replace valued content with in-line cookie accept controls. This highlights the value of the exchange for the user whilst allowing them to access the parts of the site that don't require cookies.

Some kind of general, persistent 'nag' notice is also a likely feature of a site with this model. Getting this right is about balancing the need for getting an opt-in with the user experience. If the banner can easily be ignored it will be, but if it gets too much in the way, you can also risk losing visitors.

We would recommend an ongoing commitment to A/B testing of message content, design and page location, to get the right balance here.

Mistakes to Avoid:

The biggest error is in giving users the impression of an opt-in model when cookies are being set by default.

Another one is using an opt-in model when it isn't necessary for the particular jurisdiction. We tend to see this with small websites that have picked up free scripts and haven't sought any advice.

The other big issue is assuming that opt-in always means stopping people entering the site until they have accepted cookies. It's really not necessary and can be very damaging to both traffic and engagement on sites that aren't so compelling that almost all visitors will just click to get to the content.

Risk Factors

Opt-in is going to be low risk from a compliance perspective, as long as it is done correctly and cookies don't slip through before consent is given. The biggest risks are really business related. If you don't get the balance right, you can either lose a lot of traffic, or you have a lot of unmeasurable traffic. Neither of these options is ideal, but clearly the latter is preferable.

It is also important that once users have opted in, there should remain somewhere on the site the ability to opt-out again, effectively withdrawing their consent. Not providing this could be deemed unfair by regulators, and therefore not fully compliant.

Model 5: Mixed Consent

In summary: We have set some cookies already, and would like to set some more.

As the name suggests, this is really a hybrid approach where different models are applied to different types of cookies according to their purpose.

An example would be relying on Implied Consent for web analytics and Soft Opt-in for third party advertising.

This is a fairly sophisticated approach, and not one seen very often. However detailed following of guidance from some cookie law regulators would result in this model being applied more widely.

We expect to see greater use of this model over time. However most of the lessons from the different models mentioned above can be applied as appropriate if considering this approach.

The Do Not Track Question

A final consideration to add to the mix is to decide whether and how to respond to browser Do Not Track (**DNT**) signals.

Although DNT is a standard feature of most modern browsers, it is almost completely ignored by websites.

DNT is a ‘preference expression’ setting in a browser, which is designed to provide websites with an indication of the end users wishes. However, in itself it doesn’t do anything to protect user privacy or limit the setting of cookies or any other similar technology.

There have been years of negotiations to set a global standard for what DNT should mean, and how web sites should respond. This has been carried out by the World Wide Web Consortium (W3C) – which sets most web standards, including the very basics like how HTML works. However the opposing interest groups, mainly privacy advocates and online advertising businesses appear unable to agree any compromise. So the standard has stalled around lack of agreement over what ‘tracking’ means in the context of the request.

Therefore although there is no legal requirement to honour the DNT signal, websites can choose to do so, and some do. This includes interpreting it as an opt-out from certain types of cookies. The main question however is which ones?

A narrow interpretation of the request would support opting users out of Targeting/Advertising cookies, especially as these are mostly third party and capable of profiling users across websites. This is the kind of activity that people most associate with the term ‘tracking’. However, some people believe tracking includes recording page visits within a site for analytics purposes – so that would include stopping the setting of Performance cookies.

The cookie laws allow for websites relying on ‘browser settings’ to signify their consent for cookies. This has largely been interpreted as meaning using direct controls to block cookies, although these limited in scope. However, it could also be interpreted to mean that if a browser sends a DNT signal, then the user is signalling a lack or withdrawal of consent to some types of cookies. That would then require a response from the site

It is important to realise that usage of DNT is significant amongst general web visitors. Although reliable statistics are hard to come by, in many parts of Europe, it is estimated between 10% and 15% of users are using DNT to request not to be tracked. We anticipate that it will become harder for site owners to ignore such requests from significant number of users over the coming years.

If you do choose to respond in some way to a DNT request, it is a good idea to also communicate this to visitors. This could be done both in response to a signal, as in ‘we have switched off cookies because of your DNT request’, or you could add some explanation into your privacy/cookie policy, such as ‘if you turn on Do Not Track, we still stop setting these types of cookies’.

Conclusions

The EU Cookie Law is in fact 28 different laws – one for each EU member state. Although they are substantially similar, being based on a common EU Directive, there are some subtle (and not so subtle) differences between them. Different regulators in each country also take very different views of enforcement, and have provided different levels of guidance.

This means there is really no one size fits all approach to the cookie law. In fact, multi-lingual websites targeting different users in different EU countries may need to apply several models in one website.

It is our view that these five models represent the different basic choices available to website owners when deciding how to apply a solution to the cookie law anywhere in the EU. Making the right choice is about balancing the interests of your brand, your customers, and your regulator. It is also important to make that choice **before** any other decisions about the message, design and user experience.

Once you have done that, it becomes much easier to manage the implementation.

The Cookie Collective
October 2014