COOKIE MONSTER

HOW EUROPE SHOULD FIX THE WEB'S POP-UP ECONOMY

By Anda Bologa

This report is presented by

E+Europe

Anda Bologa is a seasoned artificial intelligence and digital policy expert and one of the '35 under 35' tech leaders recognized by the Barcelona Centre for International Affairs. During her tenure at the European Union Delegation to the United Nations, she was responsible for high-level negotiations on artificial intelligence resolutions and the United Nations Global Digital Compact. Her background includes legal work on technology and business at Wardyński & Partners, handling digital platform cases at the European Court of Justice during her time with the Legal Service of the European Commission,. A Fulbright scholar, Anda holds Master of Laws degrees in Information Technology Law and Intellectual Property Law from Fordham University and in International Arbitration from Bucharest University, along with a Master's degree from the College of Europe. She is an editor of Revue Européenne du Droit.

William Echikson is a Non-resident Senior Fellow with the Tech Policy Program and editor of the online tech policy journal Bandwidth at the Center for European Policy Analysis (CEPA).

Before joining (CEPA), he worked at Google for six and a half years, running corporate communications for Europe, the Middle East, and Africa. He began his career as a foreign correspondent in Europe for a series of US publications, including the Christian Science Monitor, the New Yorker, Wall Street Journal, Fortune, and BusinessWeek. He is the author of four books, including works on the collapse of communism in Central Europe and the history of the Bordeaux wine region. He also directed and wrote documentaries for BBC and PBS.

An American and Belgian citizen, Mr. Echikson graduated from Yale College with a Magna Cum Laude degree in history. He spent a year at the Harvard University Russian Research Center. He is based in Brussels.

(This publication benefitted from insights gathered from interviews with a range of stakeholders. The author values their contributions while noting that neither they, nor E+Europe, endorse the paper's conclusions. The views expressed in this paper are solely those of the author. Any inaccuracies or oversights are the author's sole responsibility. Nothing herein should be construed as legal advice, or endorsement. The research received financial support from Amazon.)

Table of Contents

Executive Summary **2**
Digital Omnibus and Cookies 4
Recommendations **5**
Timeline: How Europe Built the Pop-Up Economy 7
Introduction **12**
Chapter One: The Cookie Revolution **15**
Chapter Two: GDPR Locks in Cookie Monster **19**
Publishers Lament 23
Chapter Three: Consent Becomes Cost **25**
Calculator Cookies 29
Chapter Four: A Fix That Still Needs Fixing **31**
Lunch Cookies 35
Chapter Five: Fixing Cookie Monster **37**
Bibliography **40**

Executive Summary

Europe set out to give people control over online tracking. It imposed restrictions on cookies, the small file a website saves on your phone or computer. That privacy-protecting goal remains legitimate. Users should not be followed invisibly, pushed into choices they do not understand, or exposed to opaque data practices.

But European web users going online now are greeted by a pop up asking, “do you share data?” The current consent system has grown into something few people understand, find useful, or defend: too many banners, too little meaningful choice, and too much legal uncertainty for businesses that rely on the open web.

A series of often conflicting privacy regulations, from ePrivacy to GDPR, combined with fragmented enforcement, have pushed routine, low-risk functions into a relentless permission machine. Some cookies do basic jobs, such as checking whether an ad appeared, stopping fake clicks from bots, or helping a website work better. Others follow people across websites and build detailed profiles about them. EU privacy rules often push both into the same consent box, despite their different privacy risks.

The result is a European “pop-up economy.” It protects privacy imperfectly, irritates users daily, and imposes costs on the firms least able to absorb them.

Consumers face consent fatigue. They are asked the same question again and again, often through notices they do not read and cannot meaningfully assess. Many click without thinking to get to desired content.

Small merchants see banners act as “do-not-entry” signs. More than 40 interviews conducted for this paper suggest users often abandon smaller websites after repeated prompts. Some merchants report losing up to 10% of potential customers when the banner appears.

Publishers face a direct revenue hit. Advertising funds journalism, and that advertising is worth more when publishers can show basic facts: how many people

saw an ad, whether the views came from real users rather than bots, whether the same person was shown the ad too many times, and whether the campaign actually ran. As refusal rates rise and regulators demand granular choices, publishers are pushed toward business-killing paywalls and “consent-or-pay” models.

Startups and SMEs face a scale problem. Large platforms can absorb legal advice, consent-management tools, engineering changes, and cross-border compliance risk. Smaller firms cannot. Fragmented enforcement turns one European rulebook into many practical interpretations.

Innovation suffers when firms hesitate to launch, expand, or test new services because the legal guidance is unclear. The cost is measured not only in compliance spending, but also in products never built, markets never entered, and customers never reached.

The European Commission’s proposed Digital Omnibus recognizes the frustration with cookie banners. It proposes “simplification.” That is welcome. But its current approach risks preserving the same consent machinery or even worsening it.

- **Keeps Scope:** The proposal fails to draw a clear line between high-risk tracking and low-risk operational uses such as limited measurement, and fraud prevention. It fails to reduce the "consent surface."
- **Risks Centralization:** It moves consent choices to browsers or operating systems risks turning tech giants owning browsers and operating systems into permanent web gatekeepers.

A better way forward is possible. Europe can protect privacy, reduce pointless prompts, and give businesses clearer rules, if it follows the below recommendations.

Recommendations

1. **Shrink the consent surface**

Allow websites to use basic, low-risk tools without asking every time for consent. This should cover essential but innocuous data collection such as counting visitors, checking that ads were shown, stopping bots and fraud, limiting repeated ads, and keeping the site secure. Strict rules still could be applied, requiring that data be used for that specific purpose, kept only briefly, and never reused to build profiles or track people elsewhere.

2. **Reinforce risk-based privacy protection**

Reserve consent for genuinely intrusive practices, especially cross-site behavioral profiling and sensitive data uses. Low-risk operational tools should be allowed to function under clear conditions without repeated prompts.

3. **Make enforcement consistent**

Reduce divergence among national authorities on issues such as analytics, cookie scope, legitimate interest, and consent-or-pay. The proposed move to bring the ePrivacy cookie rules into the GDPR framework represents a sensible step toward streamlining and harmonizing enforcement.

Today, cookie enforcement remains fragmented. National authorities can take different views on issues such as analytics, cookie scope, legitimate interest, consent-or-pay, and the boundary between ePrivacy and data protection law. Businesses need predictable rules across the Single Market, not 27 versions of the same obligation.

Under GDPR's one-stop-shop principle, cross-border cases are handled by a lead data protection authority. This should reduce the current patchwork.

4. **Clarify consent-or-pay for publishers**

Publishers need legal certainty on when they may offer readers a choice between paying for access or accepting advertising-funded access. Clear rules would protect readers while offering journalism a viable funding path.

5. **Avoid replacing banners with gatekeepers**

A centralized consent system must avoid default-setting power by dominant browsers or operating systems, avoiding locking in publishers, SMEs, and advertisers.

6. **Consider radical reform**

Imagine a scenario where ad transparency, not GDPR privacy, becomes the locus of "user choice"? Companies would gain the benefit of doubt to process data for ad-related—including personalized ad-related—without consent.

The Digital Service Act shows a potential solution. Under it, publishers who are "online platforms" need to provide detailed information about those ads. One clicks on "Ads" or "Sponsored" bringing up a window to "Manage Ads Preferences / Settings." An advertisement, not the website, becomes the focus of privacy concerns. With one bold stroke, such a reform would slay Cookie Monster.

Digital Omnibus and Cookies

The European Commission's proposed Digital Omnibus recognizes the frustration with cookie banners and attempts to simplify Europe's fragmented consent rules. Its main proposed reform, contained in Article 88a, would take the core cookie rules out of the e-Privacy rules and put them into the GDPR privacy rules. Consent would remain as the default rule for storing or accessing information. Limited exemptions are offered for low risk uses. GDPR rules would apply to enforcement.

Article 88b would allow users to express consent or refusal through automated technical signals, for example through browsers, operating systems, or digital wallets, rather than clicking through banners on every website. The aim is sensible: fewer pop-ups, more consistent choices, and less friction for users and businesses.

But the Omnibus treats symptoms rather than the cookie disease. It may reduce some visible banners, while maintaining the consent-first model that created them. If consent is simply moved from websites to browsers, operating systems, or wallets, Europe may replace one layer of friction with another. Large digital gatekeepers that already own the main browsers, and operating systems would gain additional powers at a time when policymakers want to reduce their power.

Introduction

In October 1994, a small banner appeared on HotWired, the online offshoot of Wired magazine. It asked: “Have you ever clicked your mouse right here? You will.” Almost half of users clicked.

That modest banner helped launch the commercial web. Advertising gave publishers a way to fund content without charging every reader. It gave merchants a way to find customers. It gave users access to news, services, and entertainment without paying for every page they opened. Today, digital advertising is part of Europe’s economic infrastructure. It has been estimated to add €526 billion to the EU economy each year and support around six million jobs. For media, the dependence is even clearer: advertising accounts for more than 81% of European newspaper and magazine digital revenues, and more than 90% for many digital-native publishers.

The costs of the current system are also now visible. One independent estimate puts the annual cost of cookie compliance at €2.3 billion, likely an undercount because it does not capture lost sales, abandoned visits, or services never launched. The European Commission itself estimates that reform could reduce administrative costs for businesses by more than €1.3 billion a year.

Europe was right to worry about online tracking. Users should not be followed invisibly, pushed into choices they do not understand, or exposed to opaque data practices. Civil society groups have long warned that online advertising can become intrusive and unfair, especially when it relies on dark patterns, sensitive data, or profiling that users cannot see or control. Those concerns deserve to be taken seriously.

But Europe has relied too heavily on consent as the answer. A tool meant to give users control has become a daily ritual few people understand and fewer still find useful. The web is covered in pop-ups. Users click to get rid of them. Businesses build systems to prove the click happened. Lawyers argue over button size, button order, wording, color, timing, and how often a user may be asked again.

The result is the cookie pop-up economy.

This paper is about how that happened, who pays for it, and how Europe can fix it. It is based on EU legal texts, economic studies, regulatory decisions, and a broad and inclusive outreach to privacy and consumer, industry, academic and regulatory voices, out of which more than 40 conversations with stakeholders materialised, including publishers, SMEs, advertisers, ad-tech companies, privacy lawyers, consumer-facing businesses, trade associations, and academic experts. Most of those conversations took place in the background, allowing participants to speak plainly about what the rules mean in practice.

The core problem is that different activities are pushed into the same consent box. Some cookies do basic jobs: keeping a user logged in, remembering a shopping cart, checking whether an ad appeared, stopping fake clicks from bots, or helping a website work better. Others follow people across websites and help build detailed profiles about their behavior. EU rules often push both into the same consent process, even though they pose different privacy risks.

The costs are uneven. Large platforms can absorb legal advice, engineering changes, consent-management tools, and enforcement risk. Smaller firms struggle. Publishers face a direct revenue squeeze because advertising loses value when they cannot measure audiences, prove that ads reached real people, prevent fraud, or limit repeated ads. Some have responded with “consent-or-pay” models, not because this was their preferred business model, but because the advertising-funded model has become hard to sustain.

The European Commission recognizes the frustration with cookie banners. Its new Digital Omnibus seeks to reduce friction and simplify parts of the privacy framework. That ambition is welcome. But the current proposal focuses too much on how consent is collected and too little on when consent should be required. It would allow users to allow or refuse consent once in their browser, operating system, or digital wallet. While this change would reduce visible pop-ups, it risks shifting power to whoever controls the default.

The European Commission is also preparing a new Digital Fairness Act to regulate dark patterns, addictive design, unfair personalization, and opt-out mechanisms for personalized advertising. The concerns are real. But Europe risks adding procedures around user choice without addressing the fundamental problem that users are already asked to make too many choices that they do not fully understand.

A better answer is available. Europe can protect users from harmful tracking, manipulation, and unfair profiling while reducing pointless prompts and giving businesses clear rules. It should draw a sharp line between high-risk tracking and low-risk operational uses. Basic measurement, fraud prevention, security, ad-delivery checks, and service improvement should be allowed under strict safeguards: limited purpose, short retention, no reuse for profiling, and no onward use beyond the stated function.

This paper makes the case for that reset: reduced meaningless prompts, strengthened protection against real abuse, clear rules for business, and no new gatekeepers standing between European firms and their customers.

Timeline: How Europe Built the Pop-Up Economy

As the commercial web expanded in the late 1990s and early 2000s, websites, advertisers, analytics providers, and other third parties gained new ways to collect information about online behavior. Much of this happened invisibly, through cookies and similar technologies placed on devices. EU citizens could read a privacy policy, but in practice they had little real-time visibility over who was collecting their data, why it was being collected, or how it would be used.

For more than 20 years, the EU's answer to that problem has been consent. The idea was simple: if companies wanted to store information on a device, read information from it, or process personal data for certain purposes, the owner should be asked first.

But consent was asked to do too much. It now covers different activities: cookies that help a site work, tools that count visitors or check whether ads were shown, and tracking systems that follow people across websites to build detailed profiles. Europe ended up with a system that asks too often, explains too poorly, and treats low risk uses like high-risk tracking.

1994: The web finds its bargain

Online advertising begins to fund free content and services. Cookies emerge as a simple memory tool. They keep sessions alive, remember preferences, and help measure whether ads worked. The early web rests on a simple exchange: users gain access, publishers revenue, and advertisers evidence that their ads ran.

2003: The EU creates the trigger

The 2002 ePrivacy Directive required websites that wanted to save or read almost all cookies on a user's phone or computer to give the user clear information and a

right to refuse. This "device access" rule became the starting point for later fights over when websites must ask for consent.

2011: The cookie law takes effect

The EU later tightened the cookie rule. By May 25, 2011, governments had to apply new rules requiring websites to obtain consent before placing or reading most cookies, rather than merely informing users and offering them a way to refuse.

This "Cookie Law" marked the shift from a lighter information-and-refusal model toward a prior-consent model. The change made sense for invisible tracking, where information about what someone does online is collected in the background, often across websites, without the user clearly seeing who receives it or how it is used.

But over time, the same consent process began to cover ordinary website functions too: cookies that keep someone logged in, cookies that check whether an ad was shown, and cookies used to build a behavioral profile are very different.

2010s: Consent becomes infrastructure

A single click turned into a full compliance system. Websites no longer just had to ask users for consent. They also had to keep records showing what the user was told, what they accepted or refused, when the choice was made, and whether that choice was passed correctly to advertising and analytics partners.

This created an exhaustive layer behind the banner: consent logs, vendor lists, proof trails, and consent-management platforms. The click became less like a simple choice and more like evidence a business might need to defend later.

2018: GDPR locks in the model

The General Data Protection Regulation (GDPR) went into effect on May 25, 2018 changed the legal environment around consent.

GDPR made consent harder to obtain. Consent now had to be clear, specific, informed, freely given, and easy to withdraw. Businesses also had to keep evidence that consent was valid.

If the rules were broken, GDPR imposed large fines. Many firms moved toward the safest option: ask users for consent for even routine or low risk uses.

2019-2020: The courts raise the bar

The Court of Justice of the European Union reinforced the active-consent model. In the Planet49 case, the Court held that pre-ticked boxes did not represent valid consent for cookies. Users had to actively agree. Silence, inactivity, or a box already ticked for them was insufficient.

In the Orange România case, the Court again stressed that consent must be active, specific, informed, and freely given. Consent could not be buried in paperwork or inferred from a lack of objection.

While these rulings helped clarify users' rights, they added additional burdens on businesses.

Late 2010s to early 2020s: Europe improves the banner, but not enough of the system

Regulators focused on how consent is presented: whether "accept" and "reject" buttons are equally visible, whether users can refuse as easily as they can accept, whether choices are granular enough, whether withdrawal is easy, and whether interfaces manipulate users into agreeing.

These clarifications made banners longer and more complex. The tough question remained unresolved: which uses should require repeated prompts?

2017–2025: Reform gets stuck

The Commission proposed an ePrivacy Regulation in 2017 to replace the older ePrivacy Directive. It aimed to harmonize rules across the EU, including for cookies, electronic communications confidentiality, tracking technologies, and enforcement.

But the proposal stalled. Governments, privacy advocates, industry groups, publishers, telecom operators, and technology companies disagreed over how strict

the rules should be, which uses should be exempted, and how consent should work in practice.

In the process, the Commission tested soft tools, including the Cookie Pledge, a voluntary initiative launched in 2023, to reduce cookie fatigue. But voluntary commitments could not fix the underlying legal structure.

The ePrivacy Regulation proposal was eventually withdrawn in 2025, leaving the old framework in place.

2020s: Costs mount

Users confront banner fatigue. Smaller firms face fixed large compliance costs. Publishers struggle to measure audiences, prove ad delivery, prevent fraud, and maintain advertising value.

The same consent framework is asked to cover both intrusive tracking and ordinary website operations. That is the core problem: Europe built a system designed to protect users from hidden surveillance, but it now often treats too many routine digital functions as if they pose the same risk.

2026: A chance to fix the system

The Digital Omnibus recognizes the frustration with cookie banners. It proposes to fold parts of the cookie framework into the GDPR through new Articles 88a and 88b.

Article 88a would bring the rule on storing or accessing information on a device closer to the GDPR framework and create limited exemptions for certain low risk tasks. Article 88b would allow users to express consent or refusal through automated, machine-readable signals, including through browsers, operating systems, or digital wallets, rather than clicking through banners on every website.

But this proposal risks reinforcing a few gatekeepers. Reducing friction will require more than clean presentation rules or technical signals. Europe needs to narrow the scope of consent and low risk tasks.

The opportunity is clear: reduced meaningless prompts, strong protection against harmful tracking, and clear rules for the businesses that fund and build the open web.

Chapter One: The Cookie Revolution

In October 1994, AT&T paid for a small banner on HotWired, the online offshoot of Wired magazine. It was a modest ad buy, but it helped establish the bargain that funded the commercial web. HotWired needed revenue to publish. Advertising, which had long funded newspapers and magazines, offered the answer. The content stayed free. There were no user profiles, no real-time auctions, no consent prompts.

The early web had another problem. Websites had almost no memory. Each click looked like a new visitor. A site could not reliably tell whether you were logged in, which language you had chosen, what was in your shopping cart, or whether you had already seen the same ad. Publishers and advertisers wanted something simple: basic memory and basic measurement.

Cookies solved that problem.

A cookie is a small piece of text that a website asks your browser to store and return on later visits. It usually contains a short identifier and simple instructions, such as how long it should last and which site may read it. It is not a program. It cannot, by itself, inspect your device or run code. A cookie is closer to a coat-check tag: useless on its own, but useful because the server can match it to a record on its side.

The term comes from an older computing idea called a “magic cookie”: a small token that a system gives you and that you later return as proof that you are the same user or session. Web cookies borrowed that logic.

Cookies quickly became part of the plumbing of the web. Some keep a site working. They log you in, remember what is in your cart, or keep your language setting. Others help with basic counting: how many people visited a page, whether an ad appeared, or whether the same user was shown the same ad too many times. Others help detect fraud, such as bots pretending to be people and generating fake clicks.

These uses are ordinary website operations. They are also easy to misunderstand because the same word, “tracking,” now covers too much.

Contextual advertising shows ads based on what a user is looking at now. A person reading about hiking boots may see an ad for hiking boots. Behavioral advertising goes further. It uses information about what someone has done over time, often across many websites, to predict what ad might work. That is the form of advertising that makes people feel followed around the internet.

A cookie that keeps a shopping cart alive, a cookie that checks whether an ad appeared, and a cookie used to follow someone across the web do not carry the same privacy risk. Treating them too similarly is where the policy problem begins.

The economic value came from measurement. A print ad could tell a business where the ad appeared. A digital ad could tell whether the ad had loaded, whether it reached a real user rather than a bot, whether the same person saw it too many times, and whether the campaign helped produce a sale. That information makes advertising less wasteful. It helps businesses that buy advertisements spend more carefully and helps publishers prove the value of the space they sell.

The evidence is clear. A study of 260 million ad impressions across more than 10,000 publishers found that ad prices drop by 18% to 23% in the EU when user identifiers or tracking signals are unavailable. A later field experiment found that removing third-party cookies reduced publisher revenue by 29.1%, while Google's Privacy Sandbox recovered only 4.2% of the lost revenue. Another study estimated that if advertisers lose the ability to optimize ad delivery with offsite data, the median cost per customer rises by 31%, with smaller advertisers hit especially hard.

This matters because lower ad prices do not only hurt ad-tech intermediaries. They directly reduce the revenue available to publishers, including news media, that rely on advertising to fund journalism and free content. They also make customer acquisition more expensive for smaller businesses, which depend on efficient advertising to reach buyers without the brand recognition, data reserves, or marketing budgets of larger firms.

That does not mean every form of tracking should be allowed. It means the distinction matters. Measurement, fraud prevention, frequency capping, and ad-delivery checks are part of how advertising-funded services work. Detailed

behavioral profiling raises different concerns and deserves stricter limits. “Actually, cookies are among the most privacy-friendly technologies we have,” argues Dr. Bernd Skiera of Goethe University Frankfurt.

Europe’s privacy law took a different route. The 2002 ePrivacy rules made the act of storing or accessing information on a user’s device the key legal trigger. But the trigger focused on the technical act of device access, rather than the risk created by the use of that information.

Cookies became legally sensitive even when the purpose was routine. Companies began requiring consent even for low risk cookies that keep a service running, count visitors, prevent fraud, or measure whether an ad appeared.

Exemptions remained narrow. A website could place cookies without asking only when they were strictly necessary to provide what the user requested: load the page, keep a cart, complete a login, remember a language setting. Everything else was left exposed. Audience measurement, analytics, ad delivery checks, and fraud prevention often sat in a grey zone. Many businesses treated that grey zone as a warning sign.

Consider fraud detection. Many websites use technical signals from a user’s device to detect fake accounts, bot traffic, payment fraud, credential-stuffing attacks, or other suspicious behavior. Those tools check whether a device has appeared before, whether activity looks automated, whether the same device is being used to abuse a service, or whether an ad impression or transaction is genuine.

Routine ad operations depend on signals: ad delivery confirmation for billing, performance measurement, brand safety, frequency capping, and ad fraud detection. If every such signal is treated as tracking that requires consent, the system becomes unmanageable.

Cookie rules on security and anti-fraud tools raise a perverse possibility: a website could appear to need the consent of the very person it is trying to investigate: the fraudster, warns Mikołaj Barczentewicz, a scholar of EU privacy and technology law and Senior Scholar at the International Center for Law & Economics.

Over time, the compliance burden moved from asking to proving. A simple “yes” or “no” was not enough. A website had to show what the user was told, what the user accepted, when the user accepted it, and whether the user later withdrew it. A fleeting click became an audit trail. Asking for consent became storing, managing, transmitting, and defending consent.

Online advertising made this harder because ads involve several parties. A publisher may sell space. An ad-tech company may help place the ad. Another company may measure delivery. Another may check fraud. Another may manage consent. The user is being asked to approve the whole chain.

Consent-management platforms emerged to run this system. They sit behind the banner, record the user’s choice, and pass a signal to advertising and analytics partners. They signal to those partners what the user has accepted or refused. This infrastructure helped companies comply. It did not make the user experience simple.

Some of these changes responded to real abuses. Users should not be tricked into consent. Refusing should not be harder than accepting. But each new layer also made the system more complex. The more choices a user sees, the less likely the user is to understand them. The more separate chances a user gets to refuse, the more likely routine business functions are switched off.

But that “no” has consequences. For a publisher, measurement is how advertising is priced and proved. It shows how many people saw an ad, whether they were real users, whether the campaign ran, and whether the same person was shown the same ad too often. When measurement becomes harder, ad inventory loses value. When ad inventory loses value, the free-access model becomes harder to sustain.

Europe tried to stop invisible access to users’ devices and give people more control. But by making device access the trigger, and by reading exemptions narrowly, the system pushed many ordinary website functions into the same consent process now used for higher-risk tracking. The result was predictable: more banners, more records, more legal caution, and less clarity for users.

Chapter Two: GDPR Locks in Cookie Monster

By the late 2000s, cookies were being used to recognize users across websites, measure advertising, target ads, and build behavioral profiles. European wanted to stop invisible tracking and offer a meaningful chance to object.

In May 2011, they moved cookies from a light notice-and-refusal model toward a prior-consent model. Websites could no longer simply tell users that cookies were being used and offer a way to object. For most non-essential cookies, they had to ask first.

But the strict consent rules failed to offer clear guidelines. Governments and regulators failed to read the rules in the same way. In Germany, regulators argued that legitimate interests could support some forms of advertising-related processing. In Spain, the data protection authority at one stage accepted broad user conduct, such as continued browsing or mouse movement, as a sign of consent. Across Europe, businesses faced the same basic question with different practical answers: when exactly must a website ask?

The 2018 General Data Privacy Regulation changed the risk calculation.

The regulation does not regulate cookies. It regulates personal data. ePrivacy governs storing or accessing information on a user's device. GDPR governs the personal data that may be collected or processed as a result. The two collide. Where consent is used, GDPR sets the standard. Consent must be freely given, specific, informed, unambiguous, and easy to withdraw. It also has to be proved.

Under GDPR, a banner could no longer be treated as a simple notice. It became evidence. A company needed to show what the user saw, what options were available, what purposes were listed, when the choice was made, and whether refusal was as easy as acceptance. GDPR also raised the stakes by allowing fines of up to €20 million or 4% of global annual turnover, whichever is higher.

Consent became the safest route for many firms. Although legitimate interest remained available in principle, it required a balancing test and carried litigation risk. Most companies ended up choosing consent because it is easier to gather evidence of compliance than to defend a contested balancing test. Dr. Bernd Skiera describes the logic: once a user clicks yes, the firm is usually "on the safe side." That is why consent became so attractive. It did not necessarily reflect the best user experience. It reflected legal risk management.

Case law reinforced the shift. In 2019, the Court of Justice of the European Union ruled in Planet49 that pre-ticked boxes do not constitute valid consent. Regulators also became more prescriptive about banner design, layers, withdrawal, and user interaction. The goal was to make consent more genuine. The effect was also to make consent procedural.

News Media Europe describes the result as a framework "anchored in burdensome procedures." That phrase captures the daily reality for many publishers. GDPR was meant to be risk-based. In practice, much of the work became documentary: interfaces, logs, consent strings, vendor lists, proof trails, and internal records.

The tension is obvious. Regulators want consent to be clear, specific, and freely given. Users want to get to the page. Businesses want predictable rules. The more regulators demand detail, the longer and more technical the banner becomes. The more choices users receive, the less likely they are to read them. A privacy tool becomes a compliance exercise.

The evidence on consent design confirms the problem. One large field experiment found that hiding rejection choices behind extra clicks reduced selection of those choices by 70%. It also found that 22% of users closed banners without making a choice, while holding different beliefs about what closing the banner meant. The same study estimated that users spend 6.6 minutes each week on consent interactions. Time costs, not only privacy preferences, shape the system.

GDPR also changed firm performance. One global analysis found that companies most exposed to GDPR saw profits fall by about 8% and sales by about 2%. The impact was not evenly shared. Large technology companies remained relatively

unaffected, while small technology companies suffered more sharply. That finding fits what businesses described in our interviews: compliance costs behave like fixed costs. Large firms spread them across huge user bases. Small firms feel them directly.

Other studies suggest that reduced advertising effectiveness can lower investment, increase concentration, and raise prices. Evidence from Apple's App Tracking Transparency policy shows how this can play out in practice: when advertising effectiveness falls, small e-commerce firms are hit hardest because they rely on targeted advertising to find new customers.

The economic fallout is widespread. Privacy rules do not only affect platforms. They affect the businesses that use platforms to advertise, the publishers that sell advertising, and the small firms that depend on measurable customer acquisition.

FEDMA, representing data and marketing firms, describes the Single Market problem in one phrase: "up to 27 interpretations." The law is European, but day-to-day enforcement remains national. A startup may be compliant in Belgium and still worry about the Netherlands or Luxembourg. FEDMA says firms hesitate to expand "by fear that they're not compliant with their neighbors' legislation."

That uncertainty changes behavior. A company does not need to be fined to become cautious. It only needs to believe that a regulator may later disagree with its interpretation. FEDMA gives the lesson companies draw when they invest around one reading of the law and an authority later rejects it: "don't build."

Once firms build the systems, hire the lawyers, buy the tools, and redesign the user interface, the machinery becomes hard to unwind. Consent-management platforms, legal templates, compliance dashboards, audit trails, and vendor frameworks become part of the web's infrastructure. They help firms survive the rules, but make the system feel permanent.

For publishers, the cost is direct. Advertising pays the bills, and advertising depends on being able to show advertisers that ads reached real people and worked. That requires audience measurement, ad-delivery checks, fraud prevention, frequency

capping, and ad verification. These tools are ordinary for a media business. Yet they are often pulled into the same consent debate as profiling.

Over time, the system turned risk-based privacy into interface law. The legal question became less about the actual harm of the data use and more about the journey through the banner. Did the user see the right wording? Were the buttons symmetrical? Were the toggles separate enough? Was the second layer clear? Could the user withdraw as easily as they accepted?

A user can receive a perfectly symmetrical banner and still have no meaningful understanding of what is happening. A company can maintain a clean consent log and still face uncertainty across borders. A publisher can comply with granular rules and still lose the measurement needed to fund journalism.

GDPR locked in Cookie Monster because it raised the standard for valid consent and the proof burden where consent is used. Even when policymakers admit the banner experience is broken, the instinct is often to redesign the same machinery: clearer buttons, better layers, stricter symmetry, longer cooling-off periods, more granular choices.

True reform requires a narrower consent surface. Consent should remain strict where tracking is intrusive, sensitive, or used to build detailed behavioral profiles. Routine, low-risk functions should not be forced through the same process when safeguards can protect users more effectively: limited purpose, short retention, no reuse for profiling, and no onward use beyond the stated function.

Publishers Lament

DER STANDARD

Source: Wikimedia Commons

The leading Austrian newspaper Der Standard thought it was offering a fair deal.

Readers could pay for journalism the old-fashioned way—buying a subscription for a monthly fee of €9.90—or they could read for free, allowing advertising to pay the bill. The paper was upfront. Choose a subscription and get a clean site with no ads and no data-driven targeting. Choose the free version and accept ads and the data processing that makes those ads valuable.

Then GDPR arrived, and with it, a legal morass.

Austrian privacy activist Max Schrems filed suit in 2018. The founder of the privacy NGO None of Your Business already had led scorched earth campaigns against Facebook for its privacy violations and the alleged transfer of European personal data to US security services. Against Der Standard, Schrems argued that the pay or consent model did not give readers a "free choice," because almost all readers chose the free, ad-filled version, resulting in "a North Korean consent rate of 99.9%."

The Austrian privacy watchdog agreed. It ruled in 2023 that the pay or consent model was illegal, because it only allowed binary consent or rejection. Under GDPR, readers should have the option to consent to specific types of data processing. Instead of asking readers to choose between paying for subscription or receiving ads with a single button, the regulators demanded three separate choice buttons.

- *First: embedded social media content. If an article includes a tweet or a social post inside the page, readers need to receive a separate choice - strip it out and use a screenshot.*
- *Second: website analytics. Publishers measure traffic to understand what content performs, and to show advertisers reach and delivery. Regulators said bundling analytics under the general consent screen was unacceptable. Readers needed to be offered an opportunity to opt out.*
- *Third: ad-space analysis—proof that ads were delivered. Der Standard describes this as the most absurd part of the ruling. You can't sell advertising if you can't measure whether readers view it. "If you want to sell advertisements you need to "how many visitors you have. You need to show a client that the campaign ran."*

Der Standard rebuilt its website and allowed subscribers to choose whether or not to see social media plugins in articles. If readers want them, they get a separate consent request. Readers are allowed to opt out of website analytics and opt out of ad-space measurement—choices Der Standard says are bad for their business.

The saga goes on. The Austrian Data Protection authority is now reviewing whether the changes were implemented correctly.

Chapter Three: Consent Becomes Cost

The cookie banner looks like a nuisance. In practice, it acts like a tax: a small toll collected at every visit.

Some firms can afford it. Others cannot.

Start with the merchant. A small business does not have many chances to win a customer. A person looking for a niche product may open one website, then another, then a third. Each site asks for consent. By the third banner, one e-commerce platform interviewed said, users are often “done with it… tired… fed up.” They close the page before seeing the product.

For small merchants, that lost attention becomes lost sales. They do not have the brand power of a large marketplace. They do not have thousands of products to keep people browsing. They also do not have legal and engineering teams constantly adjusting banners, records, and settings.

The platform calls cookie banners a “do-not-entry sign for small businesses.” The phrase captures the problem. The banner appears before the merchant has had a chance to show the product, earn trust, or explain why basic data use may improve the service.

The numbers from merchants are stark. “We lose about 10% of potential customers the moment our cookie banner appears,” says Artur Wagner, Chief Digital Officer of the German leather goods company Braun Büffel. As banners proliferated, Wagner says refusal rates surged. “Five years ago, 10% of users rejected cookies. Today, it’s 50%.”

The consequences are measurable. Wagner says conversion rates are 50% higher among users who accept cookies. Users who refuse may lose helpful features such as wish lists, remembered preferences, or smoother navigation. The business then pays twice: first through lost sales, then through customer support.

The platform sees this pattern across smaller sellers. When users refuse cookies, some features work less well. Customers then ask why a function is missing, why a friend sees something they do not, or why the site feels different. Someone has to answer. Someone has to maintain the “no-cookie” version of the site. For a small merchant, that is time, support, and technical maintenance that could have gone into the product.

The fixed costs make the problem worse. Even a small firm needs a consent tool to block tags until a user clicks, store the user’s choice, and keep scanning the site as cookies and vendors change. At the time of writing, Cookiebot’s published pricing lists €30 per month, per domain for accounts with fewer than four domains and up to 3,500 subpages, before implementation and maintenance. Consentmanager’s “Essential” package starts at €59 per month and includes one million pageviews, with extra charges above that.

And that is just the tool. Small firms still need developer time, legal review, banner updates, consent records, and troubleshooting as regulators and enforcement expectations change. The cost does not shrink neatly with the size of the business. It lands as a fixed bill plus ongoing work.

Advertising costs also rise when measurement weakens. A small company buys ads to find customers. If it cannot tell which ads worked, it spends more and learns less. Recent research on offsite tracking data found that when advertisers lose the ability to optimize ad delivery with offsite data, the median cost per customer rises by 31%. The burden falls especially hard on smaller advertisers.

That is why many entrepreneurs see targeted advertising as a growth tool, not a luxury. Startups need “a rifle, not a shotgun,” says Peter J. Kofler of Danske Iværksættere Startup Association. A large brand can afford broad, wasteful marketing. A startup cannot. It needs to find the right customers quickly, with limited money.

The point is not that every form of targeting should be allowed. The point is that advertising restrictions do not fall only on large platforms. They also hit the businesses that use those platforms to reach customers.

Small companies need advertising to acquire new customers, reach local audiences, and increase sales. Across six markets, most SMEs using targeted or personalized advertising said it helped them compete with larger or more established companies.

Publishers sell advertising to fund content. That advertising is worth less if they cannot measure audiences, prove that ads were shown, stop fake traffic, or limit repeated ads. NewsMediaEurope says publishers are often “required to obtain the consent of the user” for audience-measurement tools, even where the output is aggregated statistics used to understand audiences and improve content.

This is how a privacy procedure becomes a media funding problem. When users refuse consent, publishers often lose access to data used for ordinary measurement. When measurement weakens, ad inventory loses value. When ad revenue falls, publishers search for alternatives: subscriptions, paywalls, consent-or-pay, more ads, or lower investment in journalism.

Roughly 20% to 25% of users refuse consent, says Thomas Lue Lytzen, Director of Ad Sales and Technology at JP/Politikens Media Group. Without consent publishers lose or limit tools for tracking, frequency capping, measurement, and anti-fraud. Readers are pushed back to “the internet in the early 90s,” Lytzen laments.

Publishers did not invent the choice of either paying for content or consent for free articles because they wanted to annoy readers. They adopted it because the old bargain became harder to sustain. If a reader refuses advertising-related data use, and if that refusal reduces the value of the ad-funded model, the publisher has to ask for payment somewhere else.

The same cost pattern appears in compliance. Large platforms can internalize the burden. They have privacy teams, engineers, lawyers, and product managers. They can test flows, absorb complaints, and adjust across markets. Small firms often buy a tool and hope it is enough.

Users do not benefit much from this complexity. The theory is that more detail gives people more control. In practice, many users do not read the notices, do not understand the choices, or do not know what happens when they close the banner.

A large field experiment found that 22% of users closed banners without making a choice, while holding different beliefs about what closing meant. The same study estimated that users spend 6.6 minutes each week on consent interactions.

That is time lost across millions of people. It is also a poor form of privacy protection. A user can be asked many questions and still not understand the data flows. A business can collect many clicks and still not build trust.

The consent-first cookie system creates three burdensome costs.

It raises the cost of advertising. Businesses pay more to reach customers and learn less from campaigns.

It raises the cost of operating. Firms build compliance layers, support flows, and technical workarounds that do not improve the product.

It raises the cost of entry. Smaller players hesitate to expand, test, or launch because the legal and operational burden is hard to price.

Europe says it wants more competition, more startups, stronger media, and a more competitive digital economy. The current cookie system pulls in the opposite direction. It rewards scale, punishes uncertainty, and makes the open web harder to fund.

The fix is not to abandon privacy. Users need protection from invisible tracking, dark patterns, sensitive-data abuse, and intrusive profiling. But routine, low-risk functions should not be treated like the enemy. Basic measurement, fraud prevention, security, ad-delivery checks, and service improvement can be allowed under strict limits: clear purpose, short retention, no reuse for profiling, and no onward use beyond the stated function.

Reform should reduce the toll. Few meaningless prompts. Clear rules. Less fragmentation. Improved protection where the real risks are.

Calculator Cookies

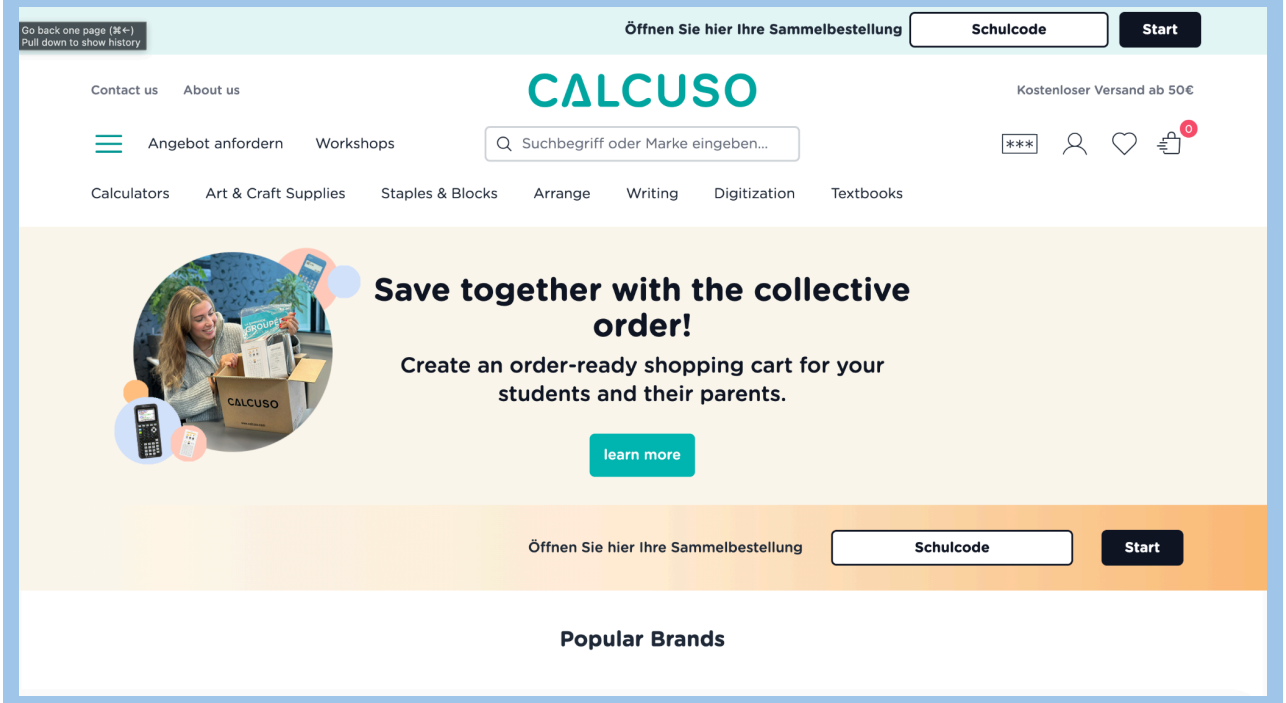

Source: https://www.calcuso.com/

It started in 2009, with a calculator. Alexander Giersz found a device that helped him with his exams at the University of Freiburg in Germany. He needed to earn money and "didn't want to work as a waiter in a restaurant." The 18-year-old student began importing calculators and selling them online.

The business thrived. Giersz dropped out of his law studies and launched CALCUSO on Amazon, C-Discount, Bol.com, and other marketplaces, offering a broad range of educational and school supplies for teachers, students, and their parents. Few regulations on e-commerce marketplaces existed at the time, and CALCUSO grew fast, expanding to three million customers, 50 employees, and €25 million in sales.

Then Giersz hit the cookie wall. He needed to advertise to build his business. But because of European privacy rules, introduced after 2017, he couldn't tell if his ads were working and bringing in customers. The reports he received from Google and Amazon, he believed, contain inaccurate, incomplete information.

“I feel blind,” he recalls. “You spend thousands of euros, and you don’t have data. I fear money is being thrown out of the window.”

The ambitious entrepreneur doesn’t need personal data. He didn’t need to know the identity or address of the potential buyer. He didn’t need to know if the buyer was a woman or a man, a teacher or a student, even if that would have been interesting. He just needed a cookie to know whether someone who clicked on his ad ended up buying something on his site.

A privacy-sensitive solution should be possible if the anonymous data of buyers were considered non-personal data. Unfortunately, the current European Commission’s The Digital Omnibus doubles down on the “consent”-only paradigm and mandates the use of centralized privacy controls through browsers, OS, or digital wallets that would manage user choice, adding to his woes because even more potential buyers would reject cookies.

The rules “end up hurting consumers,” he insists. “Since I have to spend more on advertising, my costs are higher than necessary, and I have to raise prices.”

Future European entrepreneurs will suffer. Giersz built CALCUSO before the brunt of European regulations hit e-commerce. His wife recently considered launching her own website to sell cosmetics, only to hear her husband warn her about the regulatory obstacles.

“She ended up listening to me and picking another profession, " he says.

Chapter Four: A Fix That Still Needs Fixing

Europe is trying to simplify digital rules in the name of competitiveness. That is the right ambition. The Commission has recognized that cookie fatigue is real, that users are tired of banners, and that businesses spend too much time managing consent rather than building better services.

The Digital Omnibus represents a welcome opening. It puts the problem on the table. But the current proposal risks preserving the system that created the problem. It focuses heavily on how consent is collected, while doing too little to narrow when consent is required.

That matters because the cookie problem is not only a design problem. It is a scope problem. If too many routine functions still require consent, the banner survives. It may look cleaner. It may offer a clearer refusal button. It may be linked to machine-readable signals. But the underlying machinery remains.

Article 88a of the Commission's proposal would create more detailed rules for consent screens. Websites would have to let users refuse with a single click, through a clear "reject all" option at the banner level. They would also have to respect that refusal for a set period, commonly described as a six-month "do-not-re-ask" period.

The consumer logic is easy to understand. Refusing should not be more difficult than accepting. Users should not be pushed into agreement by confusing design. Civil society groups are right to insist that consent should be real, not extracted through tricks.

But more rules on the shape of the banner do not solve the deep problem. They can even make compliance more procedural. Businesses then fight over the exact placement of buttons, the timing of repeated requests, the wording of options, and the legal meaning of silence, refusal, or withdrawal.

One large e-commerce platform warns that this kind of hard-coded interface law can create “UX lawyering.” Small firms become vulnerable to claims over tiny design choices, even when the underlying data use is low risk. For a large platform, that is manageable. For a small merchant, it represents another reason to hesitate.

The most difficult question is which functions should sit behind the banner. The Omnibus does not draw a clear line between intrusive tracking and basic operations. Without bold reform, businesses will keep treating consent as the safest cover. Routine operations will keep appearing in the same pop-up as profiling. Users will keep clicking through lists they do not read.

The second major concern is centralization.

Article 88b would allow users to express consent, refusal, and objections through automated, machine-readable signals. In practice, that points toward choices set in browsers, operating systems, apps, or digital wallets. The appeal is clear: set a preference once, stop seeing banners everywhere.

But centralization creates a new risk. Whoever controls the default can shape the data flows of the open web. If the consent layer moves into the browser or operating system, the browser or operating system becomes a powerful intermediary between publishers, merchants, advertisers, and users.

NewsMediaEurope warns that centralized settings would “disintermediate” publishers from their audiences. Publishers would lose the chance to explain, in context, why they use data and what readers receive in return. A national newspaper, a niche magazine, a local merchant, and a global platform would all be filtered through the same technical layer.

The risk is practical. Google, Apple, and Microsoft operate the most important browsers and operating systems. These companies already shape the advertising market. A browser-level consent system could give them even greater influence over who receives data, who can measure advertising, and who can fund content.

One European marketing association says centralized consent would “give disproportionate power to a few players.” It also argues that such a move sits uneasily

beside the EU's own competition agenda. Brussels would be trying to limit gatekeepers in one file while giving them a new gatekeeper function in another.

The European Commission has tried to answer some of these concerns through a media exemption. The intention is understandable. News publishers play a special role in democratic life, and policymakers do not want to damage journalism.

But the exemption may create its own problems. NewsMediaEurope warns that if users discover that "everyone has to follow their consent settings except media websites," they may feel "cheated." That would hurt the very trust publishers need.

The exemption also raises liability questions. Someone in the advertising chain may have to decide who qualifies as a media service provider. Many intermediaries will not want to carry that risk. A carve-out meant to help publishers could become another compliance puzzle.

Commercial broadcasters make a related point. They argue that the list of purposes allowed without consent is too narrow and does not reflect the low-risk processing that audiovisual media services need to operate. In other words, the problem is not only who gets an exemption. It is whether the exemption covers the functions that actually keep media businesses working.

The Digital Fairness Act debate shows the same unresolved tension. The Commission is looking at dark patterns, addictive design, unfair personalization, influencer marketing, and consumer vulnerability. Those are real concerns. Some online practices manipulate users and should be restricted.

But the policy debate is becoming confused. On one track, Europe tightens consent and opt-in expectations. On another, it considers opt-out mechanisms for personalized advertising. Both are presented as ways to give users control. Together, they show that Brussels knows the current consent model is not working, but has not yet settled on an improved architecture.

More procedures around user choice will not solve a system already overloaded with user choice. The answer should be more targeted: ban harmful practices, require

strict consent for intrusive profiling, and allow low-risk operational uses under clear safeguards.

The Omnibus can still become a useful reform. To do that, it should move beyond banner design and focus on practical changes.

Lunch Cookies

Source: Blockhütte

Malte Meurer feared getting fat. He was eating too many fast food meals and wanted to prepare healthy lunches for himself. But all the lunch boxes he saw were made of plastic. They leaked. They were difficult to clean. Maurer came up with the idea of a sleek steel substitute ("the Porsche of lunchboxes," he quips, dubbed it Blockhütte, and, in 2019, set up his shop in his mother's basement.

Confident in their product's quality, the key was alerting hungry potential customers. They were selling only online. This meant they needed digital advertising. And here's where they crashed into Cookie Monster.

They spent thousands of euros but couldn't tell which of their ads were effective. They spent thousands more hiring lawyers to come up with some workarounds. "Big companies could afford to waste money overspending on advertising," Meurer says. "Small companies like us could not."

Despite the difficulties, the customers who did buy the Blockhütte liked the product, and the company now sells in 17 European countries. But its fastest growing market is the US, where few restrictions existed on cookies. Their American ads outperformed their European cousins by an estimated three times. Since customers were brought to the product through effective advertising, American customers returned their purchases at half the level of European customers.

Like other small merchants, Meurer insists he doesn't need to know personal information about potential customers. He isn't a sports apparel company targeting a specific demographic, such as athletically minded women between the ages of 35 and 40. He just wants to avoid customers looking for plastic lunch boxes from clicking on his site. "Every click is expensive," he says. "I want to reach the right customers who are interested in high-quality, upmarket lunch boxes."

The upshot? Meurer is considering relocating to the US, even though two-thirds of Blockhütte's business remains in Europe. The US business is growing faster than the European business. Taxes are lower than in Europe. He already makes several trips a year to Los Angeles, where the company's US subsidiary is located.

"It's sunny there all the time," he says, contrasting with his current Hamburg home.

Chapter Five: Fixing Cookie Monster

Recommendation One: Shrink the consent surface for low risk uses.

Stop treating every device-based operation as if it were behavioral profiling.

Define a small set of low-risk functions that do not require repeated consent prompts when safeguards are met. At minimum:

- **Basic measurement** (counting visits and verifying delivery, not building profiles);
- **Fraud prevention and security** (bot detection, payment fraud signals);
- **Strictly functional site operations** (necessary to provide what the user asked for).

It means drawing a line between routine operations and profiling. It also means reducing the number of times users are asked the same question. Keep the core "strictly necessary" principle but explicitly recognize common low-risk operational uses (measurement, delivery verification, fraud/security) when they are first-party, proportionate, and not used to build profiles. More broadly, Privacy-Enhancing Technologies (PETs) for advertising is an area of active development and innovation that may in time enable new privacy-safe approaches to core ad-related processing.

Recommendation 2: Don't centralize consent to create gatekeepers.

If consent signals are centralized, the design must avoid turning browsers and operating systems into permission gatekeepers for the entire web.

At minimum:

- Centralized controls should not override the ability of websites to explain data use in context.
- Default settings should not be controlled by actors with commercial incentives in advertising markets.
- The system should not impose liability on intermediaries that forces them to "police" who qualify for exemptions.

Recommendation 3: Clarify consent-or-pay rules for publishers.

Publishers need legal certainty. The current situation produces litigation and redesign cycles.

If policymakers accept that journalism needs funding, they should state the conditions under which consent-or-pay is lawful for publishers: what counts as a genuine choice, what minimum alternatives exist, and what level of granularity is required.

Recommendation 4: Reduce fragmentation with consistent enforcement.

One EU rulebook that behaves like 27 creates predictable damage.

Europe should narrow divergence through:

- clearer EU-level positions on contested issues (analytics, consent-or-pay, device-access scope);
- faster consistency mechanisms;
- and less reliance on soft-law guidance that differs by country.

Europe can protect privacy without turning the web into a pop-up obstacle course. But it will require one uncomfortable move: admitting that consent cannot carry the entire burden.

The proposed move to bring the ePrivacy cookie rules into the GDPR framework is a sensible step toward more streamlined and harmonized enforcement. Today, cookie enforcement remains fragmented. National authorities can take different views on issues such as analytics, cookie scope, legitimate interest, consent-or-pay, and the boundary between ePrivacy and data protection law. Businesses need predictable rules across the Single Market, not 27 versions of the same obligation.

Centralizing cookie enforcement within the GDPR system would also bring it closer to the GDPR's one-stop-shop principle, under which cross-border cases are generally handled by a lead data protection authority. That would reduce the current patchwork, where some national authorities can use ePrivacy rules to reach conduct that would otherwise fall outside their ordinary GDPR jurisdiction. The result should

be clearer institutional responsibility, fewer parallel interpretations, and a more coherent enforcement framework for businesses operating across Europe.

Recommendation 5: Consider radical reform.

Imagine a scenario where ex-post ad transparency, not GDPR, becomes the locus of "user choice"? Companies would have the benefit of doubt that they can process data for ad-related—including personalized ad-related – purposes without consent.

The Digital Service Act shows a potential solution. Under it, publishers who are "online platforms" that show ads need to provide detailed information about those ads. One clicks on "Ads" or "Sponsored" in the corner of an ad which leads to another window with more information, and typically the opportunity to further "Manage Ads Preferences / Settings". An advertisement, not the website, would become the focus of privacy concerns.

With one bold stroke, such a reform would slay Cookie Monster.

Bibliography

1. Aridor, Guy, Yeon-Koo Che, Brett Hollenbeck, Maximilian Kaiser, and Daniel McCarthy, “Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple’s App Tracking Transparency,” Management Science, 2025.
2. Article 29 Working Party, Guidelines on Consent under Regulation 2016/679, WP259 rev.01, adopted 10 April 2018.
3. Article 29 Working Party, Opinion 04/2012 on Cookie Consent Exemption, WP 194, adopted 7 June 2012.
4. Autoriteit Consument & Markt, Market Study into Mobile App Stores, 2019.
5. Autoriteit Persoonsgegevens, Normuitleg grondslag gerechtvaardigd belang, 2019.
6. Center for Data Innovation, The Value of Personalized Advertising in Europe, November 2021.
7. Chen, Chinchih, Carl Benedikt Frey, and Giorgio Presidente, “Privacy Regulation and Firm Performance: Estimating the GDPR Effect Globally,” Oxford Martin School Working Paper No. 2022-1, 2022.
8. CNIL, Guidelines and Recommendation on Cookies and Other Trackers, revised guidance adopted 17 September 2020.
9. CNIL, Restricted Committee Sanction SAN-2021-023 of 31 December 2021 concerning Google LLC and Google Ireland Limited.
10. CNIL, Restricted Committee Sanction SAN-2021-024 of 31 December 2021 concerning Facebook Ireland Limited.
11. Court of Justice of the European Union, Bundesverband der Verbraucherzentralen und Verbraucherverbände — Verbraucherzentrale

Bundesverband eV v Planet49 GmbH, Case C-673/17, Judgment of 1 October 2019.

12. Court of Justice of the European Union, IAB Europe v Gegevensbeschermingsautoriteit, Case C-604/22, Judgment of 7 March 2024.

13. Court of Justice of the European Union, Meta Platforms Inc. and Others v Bundeskartellamt, Case C-252/21, Judgment of 4 July 2023.

14. Deisenroth, Daniel, Utsav Manjeer, Zarak Sohail, Steven Tadelis, and Nils Wernerfelt, "Digital Advertising and Market Structure: Implications for Privacy Regulation," NBER Working Paper No. 32726, July 2024.

15. Deloitte, The Value of Advertising, January 2017.

16. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector, OJ L 201, 31 July 2002.

17. Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009, amending Directive 2002/58/EC, OJ L 337, 18 December 2009.

18. Dubé, Jean-Pierre, John G. Lynch Jr., and co-authors, "The Intended and Unintended Consequences of Privacy Regulation for Consumer Marketing," Marketing Science, 2025.

19. EDPB–EDPS, Joint Opinion 2/2026 on the Proposal for a Regulation as regards the simplification of the digital legislative framework (Digital Omnibus), adopted 10 February 2026.

20. European Commission, Commission Staff Working Document: Fitness Check of EU Consumer Law on Digital Fairness, SWD(2024) 245 final, 3 October 2024.

21. European Commission, Proposal for a Regulation as regards the simplification of the digital legislative framework (Digital Omnibus), 19 November 2025.

22. European Data Protection Board, Guidelines 05/2020 on consent under Regulation 2016/679, adopted 4 May 2020.

23. European Data Protection Board, Guidelines 2/2023 on Technical Scope of Art. 5(3) of ePrivacy Directive, Version 2.0, adopted 7 October 2024.

24. European Data Protection Board, Opinion 08/2024 on Valid Consent in the Context of Consent or Pay Models Implemented by Large Online Platforms, adopted 17 April 2024.

25. European Parliament, Digital Fairness Act, Legislative Train Schedule, 2026.

26. European Parliament, Policy Department for Economic, Scientific and Quality of Life Policies, Regulating Targeted and Behavioural Advertising in Digital Services, Study requested by the IMCO Committee, 2021.

27. European Public Policy Partnership, Economic Impact of Digital Ads for CEE Companies, 2025.

28. Farronato, Chiara, Andrey Fradkin, and Tesary Lin, "Designing Consent: Choice Architecture and Consumer Welfare in Data Sharing," NBER Working Paper No. 34025, 2025.

29. GfK, Europe Online: An Experience Driven by Advertising, September 2017.

30. Goldfarb, Avi, and Catherine Tucker, "Privacy Regulation and Online Advertising," Management Science, Vol. 57, No. 1, 2011, pp. 57–71.

31. Goldfarb, Avi, and Verina F. Que, "The Economics of Digital Privacy," NBER Working Paper No. 30943, February 2023.

32. Gu, Zhengrong, Garrett Johnson, and Shunto Kobayashi, "Can Privacy Technologies Replace Cookies? Ad Revenue in a Field Experiment," SSRN Working Paper, 2025.

33. IAB Europe, AdEx Benchmark 2024 Report, 2025.

34. IAB Europe, What Would an Internet Without Targeted Ads Look Like?, April 2021.

35. IHS Markit, The Economic Contribution of Digital Advertising in Europe, September 2017.

36. Johnson, Garrett A., Scott K. Shriver, and Samuel Du, “Consumer Privacy Choice in Online Advertising: Who Opts Out and at What Cost to Industry?” Marketing Science, 2020.

37. Laub, René, Klaus M. Miller, and Bernd Skiera, “The Economic Value of User Tracking for Publishers,” SSRN Working Paper, revised 2024.

38. Marotta, Veronica, Vibhanshu Abhishek, and Alessandro Acquisti, “Online Tracking and Publishers’ Revenues: An Empirical Analysis,” Working Paper, 2019.

39. Miller, Klaus M., and Bernd Skiera, “Economic Consequences of Online Tracking Restrictions: Evidence from Cookies,” International Journal of Research in Marketing, Vol. 41, No. 2, 2024, pp. 241–264.

40. Mustri, Eduardo Schnadower, Idris Adjerid, and Alessandro Acquisti, “Behavioral Advertising and Consumer Welfare: An Empirical Investigation,” SSRN Working Paper, 2023.

41. OECD, Competition Assessment Toolkit: Volume 2 — Guidance, 2019.

42. OECD, Competition in Digital Advertising Markets, 2020.

43. Public First / IAB UK, The Digital Dividend: The Value of Digital Advertising to the UK Economy, Its Businesses and Its People, September 2023.

44. Regulation (EU) 2016/679, General Data Protection Regulation, OJ L 119, 4 May 2016.

45. Regulation (EU) 2022/1925, Digital Markets Act, OJ L 265, 12 October 2022.

46. Regulation (EU) 2022/2065, Digital Services Act, OJ L 277, 27 October 2022.

47. Sun, Tianshu, Zhe Yuan, Chunxiao Li, Kaifu Zhang, and Jun Xu, “The Value of Personal Data in Internet Commerce: A High-Stake Field Experiment on Data Regulation Policy,” Management Science, Vol. 70, No. 4, 2024, pp. 2645–2660.

48. UK Competition and Markets Authority, Competition Impact Assessment: Guidelines for Policymakers, Part 2, 2015.

49. UK Competition and Markets Authority, Online Platforms and Digital Advertising: Market Study Final Report, 1 July 2020.

50. Wernerfelt, Nils, Anna Tuchman, Bradley T. Shapiro, and Robert Moakler, “Estimating the Value of Offsite Tracking Data to Advertisers: Evidence from Meta,” Marketing Science, 2024/2025.