Mozilla

Pathways to a Fairer Digital World: Mozilla's Views on the EU Digital Fairness Act

Mozilla

Executive summary

The **Digital Fairness Act (DFA)** is a defining opportunity to modernise Europe's consumer protection framework for the digital age. Mozilla welcomes the European Commission's ambition to ensure that digital environments are fair, open, and respecting of user autonomy.

As online environments are increasingly shaped by **manipulative design, pervasive personalization, and emerging AI systems**, traditional transparency and consent mechanisms are no longer sufficient. The DFA must therefore address how digital systems are designed and operated – from interface choices to system-level defaults and AI-mediated decision-making.

Mozilla believes the DFA, if designed in a smart way, will complement existing legislation (such as GDPR, DSA, DMA, AI Act) by closing long-recognized legal and enforcement gaps. When properly scoped, the DFA can **simplify** the regulatory landscape, **reduce fragmentation,** and **enhance legal certainty** for innovators, while also enabling consumers to exercise their choices online and bolster overall consumer protection. Ensuring effective consumer choice is at the heart of contestable markets, encouraging innovation and new entry.

Policy recommendations

To achieve these objectives, Mozilla recommends that the DFA:

Recognize and outlaw harmful design practices at the interface and system levels.

- Update existing rules to ensure that manipulative and deceptive patterns at both interface and system architecture levels are explicitly banned.
- Extend protection beyond "dark patterns" to include AI-driven and agentic systems that steer users toward outcomes they did not freely choose.
- Introduce anti-circumvention and burden-shifting provisions requiring platforms to demonstrate the fairness of their design and user-interaction systems.
- Harmonize key definitions and obligations across the different legislative instruments within consumer, competition, and data protection law.

Establish substantive fairness standards for personalization and online advertising.

- **Prohibit exploitative or manipulative personalization** based on sensitive data or vulnerabilities.
- Guarantee **simple, meaningful opt-outs** that do not degrade service quality.

- Require the use of **privacy-preserving technologies (PETs)** and data-minimisation by design in all personalization systems.
- Mandate **regular audits** to assess fairness and detect systemic bias or manipulation across the ad-tech chain.

Strengthen centralized enforcement and cooperation across regulators.

- Adopt the DFA as a Regulation and introduce centralized enforcement to ensure consistent application across Member States.
- Create formal mechanisms for cross-regulator coordination among consumer, data protection, and competition authorities.
- Update the “average consumer” standard to reflect real behavioral dynamics online, ensuring protection for all users, not just the hypothetical rational actor.

A strong, harmonized DFA would modernize Europe’s consumer protection architecture, strengthen trust, and promote a fairer, more competitive digital economy. By closing long-recognized legal gaps, it would reinforce genuine user choice, simplify compliance, enhance legal certainty, and support responsible innovation.

Mozilla envisions a digital economy where autonomy, transparency, and fairness are built into technology by design. The Digital Fairness Act can make this vision a reality – ensuring that users’ choices are respected, that companies compete on quality rather than manipulation, and that Europe’s digital transformation remains open, human-centered, and fair by design.

Mozilla

Introduction

The Digital Fairness Act (DFA) presents a pivotal opportunity to modernize Europe's consumer protection framework for a digital environment increasingly shaped by complex interfaces, data-driven personalization, and emerging AI systems. As digital services become ever more integrated into people's daily lives, the line between persuasion and manipulation has blurred. Protecting user autonomy now requires going beyond transparency to address how systems are built — from sticky defaults and integrations that steer behavior to the algorithms and AI agents that increasingly mediate choice.

The DFA can serve as a major step toward a fair, open, and trustworthy digital space: one where consumers are empowered to make genuine choices, where commercial practices are transparent and accountable, and fairness is by design. This means **providing legal certainty to businesses and builders by closing long-recognized gaps in EU law, all while addressing**: harmful design practices that fall outside current frameworks; personalization that exploits data and attention rather than respecting user expectations; and enforcement that remains fragmented across legal domains.

Mozilla's perspective is shaped by **our dual role as advocate and builder.** As a global non-profit technology company, Mozilla works to ensure the internet remains open and accessible to all. Through public policy advocacy, we promote privacy, security, and competition as the foundations of a healthy digital ecosystem. Through products such as Firefox and the Gecko browser engine, we put those principles into practice by **demonstrating that it is possible to innovate while protecting user autonomy and choice.** This dual approach informs our policy recommendations: rules that make fairness, transparency, and accountability the norm across digital services.

Our approach in this paper builds on Europe's existing legal foundations while addressing the gaps that remain. **We argue for a Digital Fairness Act grounded in consumer protection, one that complements rather than duplicates existing instruments, such as** the GDPR, DSA, DMA, and AI Act. We acknowledge the European Commission's dedication to simplifying existing rules to boost innovation and competitiveness. This strategic goal remains necessary and we are concerned that it has been exploited by some stakeholders to promote a *'deregulation'* agenda that aims to weaken and remove safeguards that promote and support consumer protection, autonomy and choice. In other words, ***true simplification in the consumer protection space means greater legal certainty for businesses and builders, better enforcement, clearer institutional coordination, and a lighter burden on consumers to protect themselves***. ***Only if consumers have effective choice, transparency, and the ability to switch can contestable markets flourish and encourage innovation and new entrants.***

The **first section** of this paper explores how harmful design practices operate not only at the interface level but deep within system architecture, where tricky defaults that are hard to navigate or change, integrations, and now, AI-driven assistants can quietly restrict user freedom and entrench dominant ecosystems.

The **second section** examines personalization and online advertising through a fairness lens, highlighting how opaque profiling and targeting distort consumer choice and proposing measures such as meaningful opt-outs, accountability through audits, and privacy-preserving design by default.

The **third section** turns to legal and institutional design, calling for a directly applicable regulation with centralized, coordinated enforcement, effective cross-regulator cooperation, and an updated "average consumer" benchmark that reflects how people actually behave online.

Taken together, these recommendations outline a vision for a Digital Fairness Act that is both evidence-based and forward-looking — one that secures consumer autonomy, strengthens trust, and ensures that fairness, openness, and transparency become defining features of Europe's digital economy.

Mozilla

Addressing Harmful Design Practices to Increase Consumer Protection & Choice

In today's interconnected world, the design of digital interfaces significantly influences our daily interactions, decisions, and overall well-being. An observed troubling trend in the digital realm is the proliferation and pervasive use of harmful design practices. These practices, often termed *'dark patterns'* or *'deceptive interfaces'*, subtly (or sometimes aggressively) coerce people into decisions they might not have otherwise made, compromising the fundamental principles of user autonomy and transparency.

These designs are more than mere annoyances; they represent a stealthy influence on people's behavior, exploiting a range of practices to undermine their autonomy to the benefit of online services. This phenomenon ranges from frustrating mazes and sneaky designs in user experiences to tricks like false scarcity claims in e-commerce. More insidiously, these designs often target the most vulnerable, exploiting personal characteristics such as disability, age, health, income, or digital literacy, as well as temporary states of vulnerability.

While taxonomies are always evolving to reflect changing practices, researchers have categorized harmful design patterns into three broad types: **coercive design** (which restricts or forces user choices against their interest), **manipulative design** (which subverts or pressures user decision-making), and **deceptive design** (which gives users a false impression or misleads their understanding). Any design pattern falling into these categories can be considered **"harmful design."**

At the same time, these practices are considered '**harmful'** because they can **directly or indirectly impact consumers in their online experiences.** These harms can take the form of financial loss, violations of privacy, subverting consumer choice and autonomy by steering users toward outcomes that do not reflect their preferences, as well as more broadly distorting market dynamics and giving structural advantages to dominant platforms or services, reinforcing lock-in and weakening fair competition.

The prevalence of these practices is well-documented. A European Commission study in 2022 found that 97% of popular websites and apps used at least one deceptive design pattern. Beyond the interface level, **harmful design practices also lie deeper in the system's architecture**. Research by Mozilla has found that operating systems utilise online choice architecture to discourage users from selecting their preferred browser, which not only hinders choice but also has cascading effects related to privacy, cybersecurity, and competition.

Recognizing the harms caused by 'dark patterns', regulators have begun to act. The EU has adopted laws that aim to curb deceptive *interface* designs. This is an encouraging development. However, until now these efforts remain a patchwork, while leaving critical gaps, predominantly when it comes to harmful practices that lie deeper at the level of system architecture.

Legislative landscape and existing regulatory gaps

The **Unfair Commercial Practices Directive (UCPD)** is a crucial piece of legislation that covers harmful design, especially in the context of online advertising and commercial practices. The Directive prohibits unfair practices that could mislead consumers and affect their economic decisions. The **Digital Services Act (DSA) and the Digital Markets Act (DMA)** complement the regulatory framework, specifically targeting deceptive techniques that distort user choice. Last but not least, **the AI Act** includes a number of provisions and prohibitions against practices that "exploit vulnerabilities of specific groups of persons" or "use subliminal techniques beyond a person's consciousness" to manipulate individuals into making decisions that they may not have otherwise made. It further includes a series of transparency obligations for certain AI systems to ensure that natural persons are informed that they are interacting with an AI system.

The **European Data Protection Board (EDPB) and the Data Act** have provided definitions and categorizations of dark patterns, emphasizing their manipulative nature and the resultant harmful outcomes for consumers. Although the **General Data Protection Regulation (GDPR)** and the ePrivacy Directive do not explicitly mention dark patterns, they form part of the legal framework regulating these practices. For instance, the collection of consent under the GDPR or the ePrivacy Directive could involve harmful design techniques. The **EDPB's guidelines** on dark patterns for social media platforms offer practical recommendations for assessing these practices, highlighting their potential to hinder users' ability to provide informed consent.

Amidst all this legislative activity, there is an ongoing debate around the extent and immediacy of further legislative action in this area. Some argue for a 'wait and see' approach, emphasizing the need to allow current regulations to be enforced fully and to identify gaps before introducing additional legislation. Additionally, there are arguments that such regulation might hinder innovation or undermine a seamless online consumer experience.

In our view, action is needed to ensure a user-centric approach that facilitates innovation. While existing EU laws represent meaningful progress, they tend to focus on isolated interface-level abuses or specific sectors. What they *miss* are the **more structural, system-level design choices** that can be just as harmful. For example, an operating system (OS) or app ecosystem can be configured in a way that steers user behavior at every step – without any single pop-up or dialog crossing a legal line, the

overall architecture can still severely constrain user freedom. **The net effect is a regulatory blind spot**: *manipulative design can occur at the system architecture level, not just in one-off UI elements, yet our regulations haven't squarely tackled these architectural dark patterns.*

The European Commission's recent ***"Digital Fairness" fitness check*** report describes deceptive patterns as *"commercial practices deployed through the structure, design or functionalities of digital interfaces or* ***system architecture*** *that can influence consumers to take decisions they would not have taken otherwise."* This broad definition recognizes that the structure and default workings of a system – not just individual interface tricks – can manipulate users. Building on the fitness check report and findings, Mozilla strongly supports this direction: only a harmonized, updated legal framework can address the full spectrum of manipulative design, especially the system-level tactics that currently slip through the cracks.

Examples of harmful design in system architecture

Harmful design practices today go far beyond a misleading button here or a confusing dialog there. Increasingly, entire platforms and ecosystems are architected to constrain user agency, shape user journeys, and cement incumbents' advantages. Below, we highlight several examples that illustrate how manipulation is woven into system architecture rather than just the interface level:

Operating systems and impact on users' default choices

One clear example is how operating system design can undermine consumer choice of web browsers, a critical gateway to the internet. Mozilla's 2022 report ***"Five Walled Gardens: Why Browsers are Essential to the Internet and How Operating Systems Are Holding Them Back"*** sets out the importance of browsers and the various ways in which Apple, Google, Microsoft and others leverage their position operating system providers to restrict browser choice and competition, including through harmful design tactics.

Diving deeper, Mozilla commissioned an independent report in 2024 titled *"**Over the Edge: How Microsoft's Design Tactics Undermine Browser Choice,**"* which documents how Windows 10 and 11 steer users toward Microsoft's own Edge browser at multiple levels. When a user attempts to download and install an alternative browser, Windows deploys a host of manipulative design patterns – for instance, using preselected options, visually intrusive prompts, misleading wording, and even advertisements disguised as system messages – all to skew the user toward staying with Edge. Should the user persist and try to switch their default browser to a rival, the system throws up further obstacles and friction: Microsoft imposes an obstruction pattern by making

the default-change workflow complex and discouraging (e.g., requiring numerous confirmation steps), and in some cases refuses to transfer certain link associations to the new default. Even after a user manages to switch their default browser, Windows continues to push them back toward Edge through repeated prompts and “nag” screens – employing visual interference and other nudges whenever the user engages with the non-Microsoft browser.

Even when users successfully set a new default browser at the system level, operating systems often fail to respect the user's default choice consistently. There are instances where users are required to repeat the process within specific apps or services that override the system default. This fragmented design ***effectively makes the “default” meaningless***, turning what should be a one-time user choice into a persistent struggle against pre-configured OS behaviour. Such patterns illustrate how defaults, once set by users, are routinely ignored or reset in practice, reinforcing the dominance of platform-controlled applications.

Taken together, these design tactics exemplify coercive, manipulative, and deceptive design: **the user’s intent to choose a different service is silently undermined by the very architecture of the OS**. This is harmful equally to consumer autonomy (the user ends up using a product they actively tried to avoid) and to competition, as these design practices distort the level playing field (making it vastly harder for a competing browser to gain or retain a user).

These real-world examples also highlight **how current EU legislation remains fragmented and insufficient** in addressing such manipulation. When users encounter ‘dark patterns’ in browsers—such as the example above—these manipulative designs reveal persistent regulatory gaps across existing EU law. The **DSA** and Article 25 specifically do not apply, as browsers are not intermediary services; the **Unfair Commercial Practices Directive (UCPD)** can only address such tactics as *misleading or aggressive commercial practices* linked to promotion, leaving structural manipulation in software interfaces largely untested. The **DMA** applies only insofar as a browser or operating system has been designated as a core platform service. For instance, as illustrated above, Mozilla’s research found that Windows and Edge use repeated prompts, confusing workflows, and preselected defaults to steer users back to Microsoft’s own browser. This behaviour undermines user choice, but *the DMA provisions can only apply if Edge is designated as a core platform service*, which is not the case (contrary to the Windows operating system, which is designated as such). Meanwhile, the **GDPR** governs only when personal data or consent mechanisms are directly implicated, and the **AI Act** covers manipulation solely through AI systems that cause identifiable harm.

In practice, **this patchwork leaves no unified legal avenue to challenge manipulative system-level or interface design**. This complexity favours only companies and

platforms whose practices fall through the cracks, while consumers have no clear path to take action.

In-App Browsers

Another modern phenomenon is the use of in-app browsers within popular apps like Instagram, Facebook, and TikTok. This is a design choice that subverts user choice by keeping them within a controlled environment. When a user clicks a link in some of these apps, instead of opening the phone's default web browser (which a user might have configured for privacy or preference), the app opens the link in its own embedded browser window. This seemingly minor design decision has significant implications. Users are often not given a clear option to open the link in their preferred browser. The result is that users are kept on the platform's terms – and they may not realise that this is the case. In fact, a user might wrongly assume they are using their default browser and have the protections that this brings.

This concern has also been highlighted in the European Commission's Study to support the fitness check of EU consumer law on digital fairness (p. 235). In practice, in-app browsers can record keyboard inputs, cookies, and browsing data, and even inject tracking scripts, enabling platforms to monitor user interactions—including text entries, passwords, or payment details—without explicit user understanding or consent. Terms of service for these apps rarely explain the implications of in-app browsing.

Moreover, by forcing users to stay within the platform's own browsing layer, these apps effectively block privacy-enhancing tools such as content blockers or tracking protection (e.g. Safari with content blockers, or Firefox with Enhanced Tracking Protection). **This practice not only erodes user autonomy and transparency but also cements the platforms' power over user experiences**.

Emerging AI Assistants and Agents

The latest development in online choice architecture is **the rise of AI-powered assistants or "AI agents" integrated into platforms and operating systems**. Examples include Microsoft's Windows Copilot, which is deeply baked into Windows, Edge, and Microsoft 365 apps, as well as standalone AI-based browsers or assistants like Perplexity AI and others.

These AI agents promise to help users by providing answers, recommendations, or automating tasks via natural language. However, they also introduce new risks of manipulation and opacity. When an AI assistant becomes the intermediary between the user and the web, it can obscure the transparency and agency that users have traditionally had. Instead of the user actively searching or browsing (where they might

see a list of choices or URLs), the AI may present a single synthesized answer or perform an action directly. The concern is that users might not understand how or why an AI-driven recommendation was chosen, or what options were omitted. This creates fertile ground for *systemic bias and manipulation*: if the AI's underlying system or training data has a bias toward certain content or services, the user's journey can be subtly skewed without any obvious "dark pattern" or skewed choice architecture to blame.

The risk is amplified when these AI agents are tied to closed ecosystems or defaults. For example, Windows Copilot is integrated in a way that any web search or action it performs is routed through Microsoft's own Bing search engine and rendered in Edge by design. A user asking Copilot to, say, "find me a good restaurant" or "open my banking site" might be unwittingly kept within Microsoft's ecosystem (e.g., Bing results, etc.), even if that user normally prefers a different default browser or search engine. Because Copilot is a system-level feature, users have little ability to redirect those queries to a competitor's service – the system architecture has made that choice for them. **This is effectively the power of defaults writ large**: the AI assistant reinforces the dominance of the operating system/AI agent provider service in a way that is even less visible to the user than a traditional default setting.

In all these examples above, the common thread is that ***harmful design is woven into the very architecture of digital products and platforms***. It's not just a deceptive button here or a misleading pop-up there – it's the overall system design orchestrated to influence and limit user choice. Such practices can materially impair consumer autonomy, steer users into making choices against their preferences, and distort competitive dynamics by favoring the platform's own services.

All of these point to a clear conclusion: **we need to address harmful design at the system level, not just the interface level, in order to fully protect consumers in the digital age.** The forthcoming Digital Fairness Act should therefore fill these gaps by creating a coherent, horizontal framework to address manipulative design practices, ensuring that users' autonomy and freedom of choice are protected consistently across all digital environments.

Policy recommendations

To effectively counter harmful design practices – both at interface and system architecture levels – **policy and regulatory action must evolve**. Mozilla strongly supports the adoption of a *Digital Fairness Act (DFA)* that modernizes consumer protection law to explicitly outlaw manipulative design practices in the following ways:

Update and harmonize existing rules

The upcoming DFA proposal should aim to **define and prohibit harmful design patterns** comprehensively. Crucially, it should ***harmonize definitions*** **across the EU**, so that concepts like "dark patterns," "deceptive design," and "addictive design" are clearly defined and consistently addressed as unfair or harmful practices. Today's fragmented approach – where the DSA, DMA, GDPR, etc., each use different terms and cover different pieces – results in unnecessary complexity and uncertainty. The DFA should consolidate these efforts, close known gaps, and ensure that even subtle harmful designs (which might not trigger current bright-line rules) can be tackled.

In essence, this law would bring coherence and clarity by making harmful design unlawful as a consumer protection violation in its own right, rather than relying on indirect provisions. It should also **include an anti-circumvention clause** to prevent companies from getting around new rules by merely tweaking designs while maintaining harmful ecosystems. By updating the legal toolkit in this way, regulators can more easily pursue perpetrators of harmful design. This is a future-forward step to keep the digital market fair and user-centric.

Recognize in law system-level manipulation

The DFA should explicitly acknowledge that manipulative design can occur at the system or architectural level – not just in user interface elements. For example, default technical settings, product integration decisions, or the design of an entire user flow can be just as coercive or deceptive as a mislabeled button. Recognizing this in consumer protection law means framing *"deceptive online practices"* to include those *"deployed through the structure or functionalities of a system's architecture"* (to borrow language from the European Commission's own analysis). This could be done by clarifying that unfair design practices cover more than just visual tricks. They include **any design of technology that materially distorts user choice or undermines user control**.

Such recognition is essential for enforcement, as it empowers authorities to scrutinize aspects such as an operating system's default configurations or an app's overall design strategy, rather than just isolated UI components. It also sends a message to the industry that *product and service design will be viewed holistically*. If the overall system is built to nudge or trap users in unfair ways, it may be deemed unlawful from a consumer protection perspective. Ultimately, legal recognition of system-level harmful design will help ensure that regulators can pursue the kinds of integrated tactics we described (from OS-level browser bundling to in-app browser traps) directly as unfair practices.

Shift the burden of proof onto platforms

Given the difficulty of detecting and proving manipulative design (especially when it's baked into complex systems), we recommend a **reversal of the burden of proof** for platforms and companies when it comes to harmful design. In other words, the onus should be on these companies to **demonstrate that their systems and interfaces are fair and not manipulative**, rather than on consumers or regulators to first prove harm. This could be operationalized in a proportionate way by requiring companies above a certain size or market share to conduct *independent audits or assessments of their user interface designs and system flows* to certify they are compliant with fairness principles. If a design feature is called into question (for example, a complicated process to change a default setting), the platform should provide evidence that this design is necessary, proportionate, and not intended to frustrate user choice.

Reversing the burden would significantly empower enforcement: currently, regulators must investigate dark patterns on a case-by-case basis, which is resource-intensive. If instead companies knew they **must demonstrate the absence of manipulation,** they would build safer designs from the start, or face legal risk. It would also help address information asymmetry – platforms have the data and UX research that regulators often lack. A burden-shifting framework makes them share that information or face consequences. Such an approach would streamline enforcement and allow consumers to benefit, ensuring that the DFA achieves its goals.

Future-proof the proposed rules for the agentic AI era

The regulatory framework should anticipate how AI-driven features and assistants will reshape user interaction and online choice. **If the DFA limits its scope to current user interfaces, it risks overlooking one of the most transformative shifts in digital markets**: the emergence of AI agents capable of autonomously performing tasks, transactions, and decisions on behalf of consumers. Ignoring this development could render the law outdated before it even enters into force. As AI agents and generative assistants become embedded in operating systems, browsers, and other services, users may increasingly rely on them to navigate the web or make decisions.

As these systems evolve, many of today's safeguards—such as consent prompts or transparency notices—will lose relevance, since interactions will increasingly occur between agents rather than between users and visible interfaces. To remain effective, the DFA must embed fairness and transparency "by design," ensuring that consumer protection principles apply not only at the interface level but also within agent-to-agent interactions and system APIs. This requires **extending the notion of manipulative design to include algorithmic and architectural choices that shape or constrain the behaviour of AI intermediaries**.

To preserve consumer autonomy and genuine choice in this new context, users should be able to select and use the AI agents or assistants of their choice—not only those pre-installed or tied to a dominant ecosystem. In line with its objective to modernise consumer protection law, this freedom of choice should be **reinforced in the DFA through clear openness and interoperability requirements,** so that consumers can replace or combine agents without being locked into a single ecosystem or technical standard. Regulators should also be empowered to monitor the evolution of AI-driven design patterns and update the framework over time, so that the DFA continues to safeguard fairness, transparency, and choice in the era of agentic AI.

Personalization practices & online ads: ensuring consumer protection and innovation incentives

Digital services today increasingly personalize what users see – from targeted advertisements and product recommendations to curated social media feeds. In theory, personalization can improve user experiences by surfacing relevant content. In practice, however, **many personalization practices rely on pervasive tracking and profiling models** that collect extensive personal data and leverage opaque algorithms. A key harm of this status quo is discriminatory impact: **surveillance-driven targeting and delivery can segment, exclude, or steer people into inequitable options**—e.g., different prices, withheld opportunities, or predatory offers aimed at vulnerable users. In other words, manipulation and coercion often translate directly into unfair outcomes.

Legislative landscape and existing regulatory gaps

Significant steps to counter these issues have been taken in recent years with the adoption of the **Digital Services Act (DSA)**, which introduces new transparency requirements for recommender systems and advertising. The DSA obliges very large online platforms (VLOPs) and search engines (VLOSEs) to assess and mitigate systemic risks linked to the design and functioning of their recommender systems (Articles 34 and 35), and to offer users at least one option that does not rely on profiling (Article 38). It also requires platforms to disclose key information about ads shown to users, including the advertiser, the main targeting parameters, and whether the ad was based on profiling (Article 26). Together, these measures mark an important step towards greater accountability and user agency, helping users understand why they see certain ads or recommendations.

These transparency obligations largely stop at the surface. The DSA does not address the underlying data flows, targeting logic, or decision-making systems that shape how ads are selected and delivered. **Most of the ad-tech ecosystem** — including intermediaries such as demand- and supply-side platforms, data brokers, and measurement providers — **remains outside its scope,** leaving the mechanics of online advertising opaque to both users and regulators. Similarly, the **General Data Protection Regulation (GDPR)** provides a strong foundation for protecting personal data, anchored in principles such as lawfulness, fairness, transparency, and purpose limitation. However, the GDPR's fairness principle primarily concerns the fairness of data processing itself i.e. how data is collected, used, and disclosed, rather than the fairness of the broader effects that processing may have on individuals' choices or

behaviour. It does not explicitly determine when the use of data to influence, persuade, or steer users becomes unfair or manipulative.

Therefore, **neither the GDPR nor the DSA necessarily regulates personalization and advertising practices through a consumer fairness lens**. The GDPR addresses how data is collected and processed, *not* how that data is used to influence user decisions. The DSA focuses on platform transparency and systemic risks but stops short of establishing substantive fairness obligations on how recommender systems or targeted ads can be designed or deployed. In practice, this means that while platforms must explain how their algorithms work, they remain free to optimise for engagement or profit in ways that exploit user attention or bias — as seen in social media feeds that amplify addictive or polarising content, or in ranking systems that promote a platform's own products under the guise of relevance.

Similarly, online advertising is primarily governed by disclosure *("Why am I seeing this ad?")* but lacks fairness rules on how targeting or delivery systems profile and prioritise users, **leaving much of the ad-tech ecosystem, including intermediaries and data brokers, outside regulatory reach**. For example, neither framework directly imposes a non-discrimination duty on ad targeting or delivery, and both leave room for proxy profiling (e.g., by location, device, or behavioral signals) that can yield disparate impacts even without explicit use of sensitive data. Transparency alone does not prevent discriminatory targeting or delivery and substantive fairness constraints are needed. The widespread trade in sensitive and inferred data by ad-tech and data brokers magnifies these risks.

This creates a **regulatory blind spot since** large parts of the online advertising and personalization supply chain remain outside the reach of any clear consumer protection framework. The Digital Fairness Act should therefore close these gaps by establishing horizontal consumer protection rules for personalization and advertising practices. Specifically, it should:

- **Extend fairness and transparency obligations** to personalization systems and advertising actors outside the DSA's scope, including intermediaries and data brokers.
- **Prohibit manipulative or exploitative personalization practices**, particularly those that use profiling to distort or limit user choice.
- **Ensure cross-regulator cooperation** (between consumer, data protection, and competition authorities) to assess personalization practices holistically across legal domains.

Challenges with existing personalization practices

A core issue with today's personalization ecosystem is how it **subverts user expectations about data use**. Data provided in one context is routinely reused in another, without meaningful understanding or control. This violates the principle of **contextual integrity**, which holds that information should flow only within its intended context[1]. For instance, a user sharing an email address for an order confirmation may later find it used to target them across social media, or notice ads "following" them across the web based on unrelated browsing activity. These experiences highlight a deeper breakdown of trust: personalization that ignores contextual boundaries erodes users' ability to predict or influence how their information shapes their online environment. Mozilla advocates for **contextual relevance** as a design principle (e.g. limiting data use to what makes sense in context) to prevent such opaque, cross-context profiling at its source.

Personalization also frequently intertwines with **harmful interface design**. Many platforms or websites use ***'dark patterns'*** **to extract consent or steer users toward more profitable choices**. Opt-out mechanisms are often hidden behind multiple clicks or framed in discouraging language, while *"Agree"* buttons for personalized tracking are brightly displayed and frictionless. Such designs give the illusion of choice while effectively coercing consent. Take for example, the design of cookie banners and the use of high and low-contrasting colors. While this nudging technique is not explicitly illegal under existing EU rules, Data Protection Authorities, like the Danish DPA, have **argued** that the use of colors when choosing options can influence visitors to make certain choices.

The above practices fall between legal regimes. The **GDPR** governs the lawfulness, transparency, and fairness of personal data processing, but it remains unclear whether this extends to assessing the fairness or manipulative effects of personalization practices themselves. A platform or website can legally obtain "consent" under GDPR through a confusing banner, yet still engage in targeting that exploits user behaviour. The **DSA** introduces transparency duties and certain prohibitions (e.g. bans on profiling minors and sensitive traits, and requiring large platforms to offer a non-personalized feed), but it does not apply to the entire personalization and advertising chain, nor does it set substantive fairness obligations on how recommender systems or ad targeting are designed. Meanwhile, **enforcement** where rules exist, such as under GDPR or ePrivacy, **remains slow, uneven, and inconsistent**, especially against large platforms. As BEUC **observed,** the GDPR's *"enforcement [has been] its Achilles heel,"* allowing data-driven targeting practices to continue despite legal prohibitions. This weak enforcement has left users navigating a system stacked against them, where

[1] In online advertising, this means that data collected in one context (like a health app) should not be used to inform ads in another (like a retail site) without violating the established norms of those spheres, which can occur with technologies like third-party cookies.,

manipulative consent flows and exploitative personalization persist as standard practice.

The forthcoming **Digital Fairness Act** should fill these regulatory and enforcement gaps by clearly defining and prohibiting **unfair personalization tactics**, such as cross-context profiling, manipulative consent design, and personalization that exploits user vulnerabilities. Empowering consumer protection authorities to address such practices as **unfair commercial conduct** would lead to faster and more consistent remedies.

Policy recommendations

To address these challenges, Mozilla recommends a series of policy measures to embed **fairness, transparency, and user choice** into personalization practices, without duplicating existing obligations. The Digital Fairness Act should focus on empowering consumers and reining in manipulative tactics, as follows:

Provide meaningful opt-out and control

Empower users with an easy, genuinely effective right to opt out of personalized content and targeting. Users should have granular controls to influence the degree and type of personalization they receive. In practice, this means **mandating simple and accessible settings** (no buried menus or confusing toggles) to turn off behavioral targeting or to switch to a non-personalized feed. Crucially, **opting out should not degrade the core service.** Companies must not coerce consent by making "no personalization" modes unnecessarily limited or inferior. This approach aligns with emerging norms under the DSA (which requires very large platforms to offer a non-profiling recommender option) and strengthens them while levelling the playing field with products and services that are not directly covered under exciting rules. At the same time, the DFA should ensure that **all actors involved** (i.e., websites, online platforms, etc.) **will respect and effectively apply the consumer preferences and privacy choices made through any means of technical expression of users' choices.** For example, universal opt-out signals such as Global Privacy Control (GPC) can only be effective and respected by websites if there is the necessary regulatory backing and actors are not allowed to ignore such signals.

Prohibit exploitative personalization based on sensitive data and vulnerabilities

The DFA should declare that certain bases for personalization are presumptively unfair. This would include any targeting or personalization that relies on a user's sensitive characteristics or exploits their vulnerabilities. **Using sensitive personal data** (e.g. ethnicity, religious belief, health status, sexual orientation, political affiliation) **to personalize content or ads should be presumed *unfair and prohibited*.** Profiling

people's most private traits to influence their behavior is inherently manipulative and can lead to discrimination or predatory practices. Mozilla has advocated for prohibiting ad targeting based on sensitive categories. Under a fairness lens, personalized offers that differentiate or prey on users due to their personal vulnerabilities (for example, higher prices for a user believed to be affluent, or gambling ads shown to someone profiled as struggling with addiction) would be outright banned as unfair.

Ensure fairness and transparency in ad delivery systems

Beyond targeting inputs, ad delivery mechanisms themselves can create unfair and discriminatory effects. Even when data are lawfully processed under the GDPR, algorithmic ad delivery can determine who ultimately sees an ad, how frequently, and under what conditions, reinforcing social or economic bias. The DFA should therefore **require that ad delivery systems be designed and tested to avoid discriminatory or exclusionary outcomes, and that meaningful transparency be provided on how audiences are segmented and reached.** Users should be given clear, understandable explanations of why a particular ad was shown to them, and regulators should be able to conduct fairness testing to detect indirect discrimination or manipulation in delivery systems. Such obligations would complement existing transparency provisions under the DSA by extending scrutiny to ad delivery logic and its behavioural impact.

Strengthen accountability across the ad supply chain

Responsibility for unfair or manipulative personalization should not end with the platform displaying the ad. **The DFA should ensure accountability across the ad supply chain**, including advertisers, intermediaries, and data brokers. Each actor should be required to exercise due diligence to prevent unfair targeting and delivery practices, with contractual and technical safeguards in place to ensure compliance. Regulators should have access to documentation showing **how targeting parameters, bidding systems, and optimization algorithms operate in practice,** and how they respect user choices and avoid manipulative amplification. This shared accountability model would close existing enforcement gaps and make it possible to trace how consumer data and profiling decisions move across the advertising ecosystem.

Incentivize privacy-enhancing technologies and approaches by default

Where personalization is offered (especially in "free" services that effectively trade user attention or data for personalization), **the default technical approach should be privacy-preserving.** In practice, this means mandating the use of Privacy Enhancing Technologies (PETs) and data-minimising methods in personalization systems. Services that personalize content or ads must do so in a way that minimizes personal data collection and exposure. For example, instead of building rich individual profiles on a

central server, local, on-device solutions could be used so that raw data never leaves the user's device. The DFA can require that any platform or website engaging in personalization demonstrates compliance with privacy-by-design: *data minimization, purpose limitation, and security* must be baked in. PETs should be evaluated not only for privacy, but also for their propensity to reduce discriminatory outcomes.

Accountability through audits and assessments

Companies deploying personalization should be accountable for how those systems impact users. The DFA should require **regular audits of personalization algorithms and their surrounding design** (user interface, consent flows, etc.) to verify that they align with consumer protection principles. Independent auditors or regulators should be empowered to inspect and test these systems – for example, examining if a recommendation algorithm tends to amplify harmful content, or if an e-commerce site's personalization leads to systematically higher prices for certain demographics.

Such audits would assess whether personalization practices respect users' reasonable expectations and autonomy, and whether any manipulative or deceptive effects are present. If a platform claims its personalization is in users' interest, it should be able to demonstrate that – and an audit is a tool to hold them to that claim.

Behavioral experiments (such as A/B tests of different interface designs) can be used by regulators to detect dark patterns or coercive effects, and the results should be made available for oversight. Overall, continuous auditing and monitoring will create a feedback loop of accountability, ensuring that the design and deployment of personalization tools remain aligned with consumer rights and do not cross the line into manipulation.

Making fairness work in practice: effective enforcement and cooperation

As the European Commission develops the Digital Fairness Act (DFA), it is crucial to incorporate key structural features that ensure effective and uniform protection for consumers across the EU. In particular, the DFA should establish a **centralized enforcement framework**, take the form of an EU **regulation** and create clear mechanisms for **cross-regulator cooperation**. Furthermore, the DFA presents an opportunity to **modernize core consumer protection concepts**, such as redefining **the notion of the "average consumer,"** to better reflect real-world consumer behavior in the digital age. Incorporating these elements will not only close existing enforcement gaps but also promote a more coherent and fair digital marketplace.

Centralized enforcement for consistency and impact

A centralized EU enforcement framework for consumer protection rules is crucial to overcoming the current patchwork of national enforcement and to addressing cross-border digital practices effectively. Under the current status quo, enforcement of consumer protection largely falls to national authorities, resulting in uneven outcomes and difficulties in scaling action against EU-wide online practices. Member State agencies often face resource constraints that hinder proper enforcement of digital consumer laws across borders. To tackle this, the European Commission **should propose a more harmonized approach with greater EU-level enforcement powers,** moving beyond the limitations of minimum harmonization and fragmented national action. Shifting certain enforcement responsibilities from Member States to a coordinated EU framework would help ensure that **rules against manipulative design and other unfair practices are applied uniformly and no violator can escape oversight due to jurisdictional gaps**. Centralized enforcement would strengthen the Digital Single Market by providing a level playing field and consistent consumer protection across all Member States. We recommend that the DFA explicitly include provisions for such an EU-level or jointly coordinated enforcement system – for instance, empowering the European Commission or a network of national authorities led by the Commission – to investigate and sanction infringements of the new rules. This would address the currently insufficient enforcement undermining consumer protection and ensure that companies face equal scrutiny regardless of where they operate.

The case for a regulation rather than a directive

To maximize effectiveness, the legal instrument for the Digital Fairness Act should be an EU **Regulation, not a Directive**. Choosing a regulation means the rules will be directly applicable and **uniform across all Member States**, avoiding the delays and

divergences that come with national transposition. A directive, by contrast, could **exacerbate fragmentation**, as each country might implement the requirements differently. Indeed, the very rationale for the DFA is to simplify and harmonize consumer protection rules in the digital environment, given that the current **patchwork of national laws creates fragmentation and compliance burdens** for cross-border services. A regulation would tackle this head-on by establishing one set of binding rules for all, immediately enforceable and interpreted consistently throughout the EU. This approach aligns with the Commission's single market justification for the initiative: only a **fully harmonized rulebook can simplify existing rules, remove legal uncertainty, and prevent a race to the bottom** in consumer protection.

Cross-regulator cooperation and coherent oversight

Effective enforcement of digital fairness rules requires not only stronger powers and resources for authorities but also coordinated oversight across regulatory domains. Issues such as dark patterns, manipulative personalization, and exploitative algorithms often straddle consumer protection, data protection, and competition law. When enforcement remains siloed, important aspects of these practices risk falling through the cracks or being inconsistently addressed.

The DFA should therefore **establish clear and formal structures for cross-regulator cooperation at both national and EU levels**. Consumer authorities, data protection regulators, and competition enforcers should be empowered, and where appropriate mandated, to share information, coordinate investigations, and take joint action when a digital practice simultaneously affects privacy, consumer rights, and market dynamics. Existing networks, such as the Consumer Protection Cooperation (CPC) and the European Competition Network (ECN), could provide useful foundations for this, but a dedicated digital fairness task force or liaison mechanism should be established to bridge these communities. The goal should be a **consistent interpretation of overlapping rules, joint prioritisation of cases, and a unified enforcement front that prevents companies from exploiting jurisdictional or legal gaps**.

Models like the **UK's Digital Regulation Cooperation Forum (DRCF)** — which brings together competition, consumer, and data regulators — and the **ICO–CMA joint position** on harmful design show the value of aligning enforcement approaches and shared messaging. In the EU context, regular cooperation among the European Data Protection Board (EDPB), the CPC network, and competition authorities would help ensure that manipulative design, profiling, and targeting practices are addressed comprehensively. The DFA should also **mandate the publication of joint guidance** clarifying how instruments such as the UCPD, the DSA's Article 25 ban on dark patterns, and data protection rules can complement one another in enforcement. The

recent joint guidelines from the EDPB and the European Commission on the interplay between the GDPR and the DMA serve as a useful model.

Finally, coordination must be backed by meaningful deterrence. **Regulators should consider stronger penalties or remedies when design practices harm both consumers and competition**, reflecting the broader systemic impact. By fostering a culture of cooperation and robust enforcement, the DFA can ensure that harmful design and manipulative personalization are not merely identified but actively deterred — and that users experience consistent protection regardless of which law or regulator is engaged.

Redefining the 'average consumer' standard for the digital age

DFA should **modernise how EU law interprets the "average consumer."** The traditional benchmark — someone "reasonably well-informed, observant and circumspect" — no longer reflects real behaviour in today's digital environment, where complex interfaces and behavioural design shape decisions. Users often act on impulse, skip reading terms, or are nudged by design cues; even informed individuals can be momentarily vulnerable depending on context. This high threshold leaves many unprotected, as it assumes a level of rationality and attention few can consistently maintain.

The problem is amplified by **personalization**, where firms tailor offers or messages to individuals' traits or vulnerabilities. In such cases, the idea of a single "average" consumer becomes meaningless. The **Dutch Authority for Consumers and Markets has already proposed** that in highly personalized practices, traders should consider the actual characteristics of the targeted group, not an abstract average. In addition to this, mobile devices make it even easier for consumers to split attention.

To align with digital realities, the DFA should update the average consumer standard as follows:

- **Adopt a realistic benchmark:** Recognize that consumers routinely rely on heuristics and that harmful design can influence even diligent users.
- **Exclude personalization from the generic test:** Where practices are individually tailored, enforcement should assess them against the targeted user or group, not a hypothetical average.
- **Address design practices that mislead at scale:** Regulators should be able to act when dark patterns or coercive designs mislead a significant share of users, without needing to prove that every "average" person would be deceived.
- **Broaden the protection threshold:** Acknowledge that even cautious consumers can be misled when confronted with manipulative timing or interface design.

> The law should guard against such **situational vulnerabilities** rather than assume perfect attentiveness.

By aligning the legal standard with actual consumer behaviour, the DFA can raise protection levels for all users and ensure that digital fairness rules reflect how people really interact — not how an idealised “average consumer” is presumed to.

Mozilla

Conclusion

Mozilla stands unwavering in its commitment to an internet that empowers users, respects privacy, and promotes fairness. We believe that strong, forward-looking regulation is essential to achieving that vision. The forthcoming **Digital Fairness Act (DFA)** represents a unique opportunity to strengthen consumer protection and choice in a digital ecosystem increasingly shaped by personalization, opaque design, and emerging AI systems.

At the heart of Mozilla's message is the principle of **user autonomy**: individuals should be able to make genuine choices, understand when and how they are being influenced, and trust that digital products are designed in their interest. Consumers deserve services that are open, transparent, and respect contextual integrity, ensuring that data provided for one purpose is not repurposed for another without clear consent or understanding. Our mission has always been to put people first, and the DFA can translate that principle into law by hardwiring fairness and accountability into digital markets. To achieve this, we urge policymakers to craft a **Digital Fairness Act** that:

- **Bans harmful and manipulative design practices**, ensuring that dark patterns and system-level coercion are explicitly outlawed;
- **Addresses manipulative personalization and profiling** as matters of consumer fairness, filling the gaps left by data protection and platform laws;
- **Rebalances accountability**, placing the burden on companies to prove that their interfaces and architecture are not designed to mislead or deceive;
- **Establishes a centralized and coherent enforcement framework**, with strong cooperation between consumer, data protection, and competition authorities to close existing enforcement gaps;
- **Modernises the "average consumer" benchmark**, aligning it with how people actually behave online, acknowledging behavioral biases and situational vulnerabilities;
- **Future-proofs the framework for the AI era**, ensuring that AI assistants and agents respect openness, interoperability, and user choice rather than reinforcing lock-in.

When users know that products are designed fairly and their choices are respected, confidence and innovation flourish. When companies compete on usability, quality, and trust — not manipulation — the result is a healthier market and a more sustainable digital economy. Looking ahead, as AI agents, immersive environments, and new interfaces reshape the way people interact with technology, the **DFA must ensure that consumer protections remain consistent across all digital contexts**. Users should always have ***transparency*** (e.g. knowing when they are dealing with automated systems or ads), ***control*** (e.g. clear and simple ways to opt out), and ***recourse*** (e.g. effective

remedies when things go wrong). Embedding these rights into law will help future-proof the EU's consumer protection framework. By enacting a **Digital Fairness Act** that closes legal gaps, harmonizes enforcement, and anchors fairness in the design of digital products and systems, the EU can set a global standard for consumer protection in the digital era.