Personalization and Youth Online:

Assessing Benefits, Risks, and Safeguards

JUNE 2026

FUTURE OF PRIVACY FORUM

AUTHORS

Daniel Hales, Policy Counsel, US Legislation
Holly Hawkins, Director for Youth Policy, US Policy

EDITORS

Tatiana Rice, Senior Director for US Legislation, **Jules Polonetsky**, CEO,
Matthew Reisman, VP for US Policy

ACKNOWLEDGEMENTS

The authors thank Camille Herring, Bailey Sanchez, Jim Siegl, Kelly Brandmeyer, and all other reviewers for their support on this report.

The Future of Privacy Forum (FPF) is a global non-profit organization that advances principled and pragmatic data protection, AI and digital governance practices. We convene leaders across industry, academia, and the public sector to provide expert analysis, benchmarking, and best practices that support responsible innovation and regulatory compliance. Learn more about FPF by visiting fpf.org.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Personalization—the use of personal data to tailor content and services to individual users—sits at the center of today's most contested debates about youth and digital technology. This document provides a comprehensive framework for understanding personalization in youth online experiences, examining its mechanisms, benefits, risks, mitigations, and the emerging legislative and regulatory landscape. The report identifies five key takeaways:

1. **Personalization is not a single practice.** It encompasses a range of mechanisms, each with distinct data practices, risk profiles, and youth-specific implications that require differentiated policy responses.
2. **Personalization delivers meaningful benefits for young people and is essential to making digital services safe and usable for young people.** Personalized systems help young people find relevant information and resources, connect with peers and communities, and access adaptive learning tools. It also enables functions such as age appropriate content filtering and accessibility accommodations.
3. **Personalization also carries significant risks for young people that vary by context, data practice, and implementation.** Key risk categories include privacy and profiling, content exposure and amplification, manipulation and autonomy, disruptions to well-being, blurred and persuasive commercial practices, and misrepresentation and bias—many of which are cross-cutting and can arise across content, advertising, and AI contexts simultaneously.
4. **No single mitigation addresses the full range of personalization risks, and each introduces tradeoffs.** Mitigation strategies around alternatives to personalized recommender systems, safety and well-being measures, data governance and transparency requirements, and activity-related governance for advertising and AI each target distinct risks while introducing distinct tradeoffs.
5. **The emerging legislative and regulatory landscape reflects genuine concern but significant divergence.** From existing privacy frameworks that establish baseline data constraints to emerging regulatory responses ranging from feature restrictions to outright platform bans for minors—the lack of consensus reflects the genuine difficulty of calibrating interventions to personalization's complex risk and benefit profile.

I. Introduction

Personalization—the use of personal data to tailor content and services to individual users—sits at the center of today's most contested debates about youth and digital technology. Though often treated as distinct issues in policy and legislation, many of the most contested debates about children's online experiences trace back to the same underlying practice: how entities collect and use minors' data to shape their experiences across digital contexts. It is a practice that now touches nearly every digital environment young people inhabit—from social media and streaming platforms to AI tools and tutoring assistants. That mechanism can serve young people well: personalizing learning, providing age appropriate digital experiences, surfacing relevant resources, and reducing friction in complex digital environments. But it also sits at the root of the field's most pressing policy concerns. For example, concerns about social media and adolescent mental health are, at their core, concerns about what algorithmically personalized feeds surface and amplify; systems aimed at providing tailored content feeds based on user interests may steer young users toward harmful content, filter bubbles, or escalating content sequences.

Young people encounter that same mechanism at a distinct moment in their development—one that shapes both the risks personalization poses and the benefits it can offer. Children and adolescents are still developing the cognitive and emotional capacities that adults use to critically evaluate content, recognize manipulation, and make informed choices about their digital lives, making them more susceptible to potential harms. But that same developmental context also makes personalization potentially more valuable: tailored educational tools can meet young learners where they are, personalized content can surface age-appropriate and enriching experiences, and adaptive systems can provide support that generic, one-size-fits-all services cannot. These realities make the stakes of personalization in youth contexts distinct—and make the need for nuanced, evidence-informed policy responses all the more pressing.

Because risk profiles vary widely depending on data practices and use cases, there is no single, one-size-fits-all approach to mitigating potential harms from personalized online experiences. Different interventions may be used to address context—and use-specific personalization risks, each with their own distinct benefits and tradeoffs.

To navigate that complexity, this document provides a comprehensive framework for understanding personalization in youth online experiences, across five areas:

(1) Definitions and common use cases, establishing a shared understanding of how minors experience personalized online services;

(2) Key benefits of personalization in youth digital experiences;

(3) Common risks of harm;

(4) Emerging mitigation proposals, including how those mechanisms are designed to address risk and the potential tradeoffs each presents for stakeholder consideration; and

(5) The emerging policy and regulatory landscape, examining how legislators and regulators globally are responding to personalization risks in minors' digital experiences and the range of approaches currently being pursued.

II. Personalization 101: Mechanisms and Use Cases

Personalization encompasses a wide range of data-driven practices—from social media feeds ranked by engagement history to adaptive learning tools and AI chatbots that remember prior conversations—each distinguished by what data is collected, how it is used, and to what end. Before evaluating its benefits, risks, and proposed mitigations, it is worth establishing what personalization actually looks like in practice—and specifically, how it operates across the digital services young people use every day.

This section defines personalization and its core mechanisms, then maps those mechanisms onto common youth-facing use cases across content, advertising, and AI-driven services. That foundation sets up a critical distinction, developed at the close of this section, between functional implementations that make services usable and age-appropriate, and higher-impact implementations that raise the more significant privacy, safety, and well-being concerns.

A. HOW PERSONALIZATION WORKS: CORE MECHANISMS

Broadly speaking, personalization is the process of using data about an individual to tailor or adapt content recommendations, interface functionality, and experiences within an online service to that specific individual. Typically, most personalized services involve some level of algorithmic automation—whether to carry out a user's expressed preferences or to adapt the experience based on observed behavior. How that process works in practice varies significantly depending upon the data sources, the categories of data collected and used, how data is retained and processed, and the purpose of processing within an implementation context. The table below outlines five core mechanisms that distinguish how personalization operates across the digital services young people encounter.

MECHANISM	DESCRIPTION	EXAMPLE IN YOUTH CONTEXT
Explicit Personalization	Uses information directly provided by a user—such as account settings, stated preferences, product feedback, or instructions—to tailor an experience.	*A student selects preferred subjects on an educational platform; the service surfaces relevant content and resources accordingly.*
Implicit-Behavioral Personalization	Adapts a user's online experience based upon implicit behavioral data, such as user history, service interactions, and behavior patterns, without requiring them to provide the information directly.[1]	*A streaming service tracks which videos a teen watches to completion and adjusts its recommendations accordingly.*
Contextual Personalization	Optimizes a service's utility by tailoring the experience to a visitor's immediate "context," such as general location, device type, language, or time.	*A search engine surfaces local after-school programs based on a student's general location.*
Demographic Personalization	Leverages a user's demographic data, such as age, gender, education, or location, to tailor content or experiences.	*A platform applies more protective default settings and content filters when it identifies a user as a minor.*
Cross-Platform Personalization	Draws upon data collected and shared across a user's online experiences to build a more comprehensive picture of a user's interests and behaviors, often through use of cookies.	*A young person's browsing activity across multiple apps and sites informs the ads and content they see on an unrelated platform.*

These mechanisms are not mutually exclusive—most personalized services draw on several simultaneously, and the combination of mechanisms at work in a given context significantly shapes its risk profile.

PURPOSES OF PROCESSING
Functional and Higher-Impact Implementations

Understanding how personalization works—the mechanisms outlined above—is only part of the picture. While the risk profile of personalization practices is strongly influenced by the data sources and categories of data collected and used, an equally important determination is the purpose those mechanisms serve. In other words, not all processing of user data for service personalization is inherently concerning or high-impact—the purposes of data processing for personalization play a crucial role in defining overall risk. Take implicit-behavioral personalization as an example: A children's education platform that tracks which lessons a student completes in order to surface the next appropriate exercise is using the same underlying mechanism as a social media platform that tracks every scroll, pause, and interaction to build an engagement profile designed to maximize time on platform. The data collection approach may be identical, but the purpose and risk profile are not.

Conversely, many personalization practices are necessary to improve, support, and maintain service functionality. These **"functional"** implementations of personalization leverage basic data sources and mechanisms, such as device type, IP address, location, or age data collected explicitly or contextually to implement local regulations, push region-specific service updates, adapt the user interface to the specific device used to access the service, or implement safety features and age-appropriate experiences for kids and teens. This kind of functional implementation differs from higher-impact implementations, which typically draw on more expansive uses of personal data to tailor content or experiences for engagement or monetization. Higher-impact implementations may present greater privacy, safety, and well-being risks, as they often involve broader and more varied data collection and processing practices—requiring more nuanced mitigation strategies to improve user protection. **As a result—even though the use cases and concepts discussed in this paper will intrinsically implicate certain low-risk functional implementations—this paper will focus primarily on risks and mitigations in higher-impact implementations due to the more nuanced risk considerations involved in those applications. Importantly, mitigations for high-risk functionality may not be appropriate for lower-risk functionality and may result in negative unwanted tradeoffs.**

B. PERSONALIZATION ACROSS YOUTH-FACING DIGITAL SERVICES

Personalization appears in nearly every digital service young people use—influencing the content they see, how information is ranked, which recommendations are surfaced, and what prompts or ads are presented.

This section outlines three categories of common personalization use cases: (1) content; (2) advertising; and (3) conversation and generative AI. Understanding where and how personalization appears in each context is essential for evaluating the benefits, risks, and mitigations discussed in the sections that follow.

1. Content

Content personalization shapes the digital environments young people inhabit more than any other category—influencing what information they find, what communities they discover, how their learning experiences adapt to their individual needs, and how interactive and commercial services tailor experiences.

Across these contexts, personalization is not a single uniform practice but a set of related functions that operate differently depending on the service, the data involved, and the purpose being served.

Information and Discovery

Search engines, news publishers, and streaming platforms all use personalization to organize what would otherwise be an unmanageable volume of content—surfacing what is most relevant to a specific user based on their location, prior activity, stated preferences, and behavioral signals. For young people, this organizing function is particularly valuable: personalization makes large and complex information environments navigable rather than overwhelming, and can surface age-appropriate content while filtering out material unsuitable for younger audiences. But the same systems that organize information also shape what young people are exposed to—and what they are not—making the design and purpose of these systems consequential for how young people understand the world around them.

PUBLISHER CONTENT	STREAMING MEDIA	SEARCH
Surfaces stories, links, and content based on a user's interests, history, or characteristics. For kids and teens, it can provide age-appropriate results and more robust privacy and safety protections.	Draws on a user's viewing history, engagement patterns, and contextual cues like time or device, to display relevant content and recommendations. Recommendation systems can also be used to recommend age-appropriate content for youth where such features are implemented.	Tailors search results and ranking to user information and characteristics, such as their account information, location, related search queries, and previous use. For kids and teens, it can provide age-appropriate results and more robust privacy and safety protections where such features are implemented.

Connection and Community

Beyond information discovery, content personalization also shapes how young people connect with others and find community online. Social platforms use personalization to determine not just what content young people see, but who they are encouraged to interact with and which communities they find.

SOCIAL PLATFORMS
Operates across feed ranking, connection recommendations, community discovery, and engagement prompts simultaneously. Can shape what young people see, with whom they are encouraged to connect, which communities they find, and how they are prompted to interact. May appear in other services such as gaming, education, and retail service.

Learning and Development

Content personalization also plays a distinct role in how young people learn. Educational personalization adapts content, pacing, and support to individual students in ways that generic instruction cannot.

EDUCATION AND LEARNING
Provides students with individualized learning plans tailored to a student's particular learning needs more specific and relevant tutoring support to supplement classroom instruction, and access to information. Can assist both under- and over-performing students, support teachers' ability to create more responsive learning plans, and help schools and education technology vendors meet compliance obligations regarding the collection and use of student data.[2]

Interactive and Commerce

Finally, personalization appears in gaming and retail environments young people use—contexts that sit at the boundary between content, community, and commerce.

GAMING	RETAIL AND E-COMMERCE, AND FOOD DELIVERY
Customizes in-game recommendations, matchmaking, and overall gameplay experiences. Analyzes a player's behavior, skill level, and interaction history to surface relevant games, modes, or challenges that align with their preferences, making it easier for players to discover content. Can also be used to create age-appropriate and safer experiences for younger players by ensuring that younger players are matched with players of a similar age.	Personalization examples include use cases such as using location data to facilitate proximity and navigation-dependent services and products. For instance, location personalization can help users submit in-app food and drink orders at the most proximal shop location to a user at a given time. Personalization may also tailor product recommendations and search results based on past purchases, expressed preferences, and search queries.

2. Advertising

Advertising is a primary economic engine of the digital ecosystem—funding many of the services young people use daily, from social media platforms and search engines to mobile applications and educational tools, often at no direct cost to users. That economic reality means that the same data-driven practices that raise concern in advertising are also what make many youth-facing services financially viable. At the same time, advertising presents distinct considerations in youth contexts. Young people are still developing the cognitive capacity to recognize and resist commercial persuasion—particularly when personalized ads are embedded in content feeds, influencer posts, or interactive environments in ways that blur the line between content and marketing.[3] And the data practices that power personalized advertising—drawing on behavioral signals, inferred interests, and cross-platform tracking—raise particular privacy concerns when applied to minors, who may have limited awareness of how their data is collected and used to reach them commercially.

Personalization shapes advertising at every stage of the lifecycle—from determining which users are reached and what "creative" (visual, auditory and text elements of advertisements) they see, to measuring how they respond and refining future campaigns accordingly.[4] Each of these stages involves distinct data practices and raises particular considerations for young users.

PERSONALIZED ADVERTISING	PERSONALIZED AD CONTENT	MEASUREMENT AND REPORTING
Uses data about users such as their inferred interests, location, and browsing behavior to make ads more relevant and help advertisers reach users more likely to be interested in their product or service. Differs from contextual advertising, which targets based on the content currently being viewed, coarse location, and basic technical and contextual data, rather than a user's profile or history. For young people, personalized ad targeting raises questions about the appropriateness of using minors' behavioral data, including data collected across platforms, for commercial reach decisions.	Personalizes ad content—dynamically rendering different images, copy, animation, video, product information, or calls to action based on available signals about the user, such as inferred interests. For young people, this raises a distinct concern: when ad creative is dynamically tailored to match the look, tone, and format of the surrounding content—a sponsored post styled like an organic creator video, for example—it becomes harder to recognize as advertising at all. Young people, who are still developing media literacy skills, may be particularly susceptible to this blurring of content and commerce.	Uses performance data, such as reach, clicks, and conversions, to assess how users respond to different ads, messages and creative variations. These insights can help advertisers determine effectiveness and adjust future campaigns across different user segments. For young people, measurement practices mean that their interactions with ads, including passive exposure, generate data that shapes how they are targeted in future campaigns.

Table Note: Adapted from Gao et al. (2023).

3. Conversational and Generative AI

Conversational and generative AI systems represent a distinct and rapidly evolving personalization context because unlike recommendation algorithms that select from existing content, AI systems generate responses, outputs, and actions tailored to a specific user in real time. For young people, this distinction matters because personalization embedded in AI systems is not just about what content is surfaced but how a system communicates, what it remembers, how it responds emotionally, and increasingly, what it does on a user's behalf. These capabilities offer meaningful benefits for young users, such as supporting learning, creativity, and information access, but they also introduce risks that are qualitatively different from those associated with content feeds or targeted ads.

Personalization in AI systems operates across three increasingly autonomous contexts, each with distinct implications for young users. Although these three contexts are presented separately for analytical clarity, in practice they often overlap—a chatbot is frequently powered by a generative AI system, and agentic AI typically builds on both. What distinguishes them is not the underlying technology but the nature of the personalization and the degree of autonomy involved.

CONVERSATIONAL AI (CHATBOTS)	GENERATIVE AI (CONTENT GENERATION)	AGENTIC AI
Draws on prior conversations, memory, saved preferences, and age signals to adapt tone, guidance, and responses in real time. For young people, it supports tutoring, information-seeking, and wellness access—but the emotionally responsive and individualized nature of these systems raises concerns around overreliance, dependency, and whether young users understand they are interacting with a tool rather than a person.[5]	Produces original content such as text, images, and audio that is tailored to a user's inputs, context, and interaction history—adapting complexity, tone, and subject matter accordingly. For young people, it supports learning and creative expression, but raises concerns around accuracy—systems may generate confident but incorrect outputs—and around consistent application of age-appropriate content standards.	Goes beyond generating responses to taking actions on a user's behalf, such as browsing, purchasing, managing communications. Draws on history, preferences, and explicit and contextual signals to determine what to do and how. For young people, this is the least settled context: questions around data retention, autonomous decision-making on behalf of minors, and the scope of consent remain largely unresolved in both policy and practice.

SPOTLIGHT
Personalized AI Chatbots

AI chatbots increasingly use personalization to maintain a user's context and intentions. This may include using prior interactions, saved preferences, memory, inferred interests, age signals, or other contextual information to shape the system's response, guidance, tone, recommendation, or overall experience. In this respect, the chatbot is the interface through which personalization is used to shape the interaction. This framing is consistent with the OECD's definition of an AI system as one that "infers, from the input it receives," how to generate outputs such as "predictions, content, recommendations, or decisions," and that may vary in its "autonomy and adaptiveness" after deployment.[6]

Personalization can shape how a chatbot explains information, what follow-up questions are asked, what resources are suggested, and how socially or emotionally responsive the system appears. For youth, these functions may support information-seeking, tutoring, creativity, wellness resources, or entertainment, as detailed in a Pew Research Center report. For example, a young person might ask a chatbot to explain a math concept in simpler terms or gain more information on a hobby.[7] At the same time, personalization used within AI chatbots may pose risks when they make interactions feel highly individualized or emotionally responsive. Risks may include overreliance, inaccurate or inappropriate advice, emotional dependency, exposure to harmful content, privacy concerns from sensitive disclosures, and confusion about the role of the chatbot as a tool rather than as a substitute for a human relationship or licensed mental health professional.[8]

Proposed mitigations for AI chatbots used by youth are further discussed in Section VI. These include a clear disclosure that the chatbot is not human, limits on designs that intentionally create emotional dependency; and age-appropriate experiences.

III. Benefits of Personalization for Youth

Personalization delivers meaningful value to young people, and understanding that value is essential to evaluating any proposed intervention. The use cases described in the previous section illustrate where and how personalization operates across the digital services young people use. This section examines the benefits those use cases deliver for young people—and why those benefits are a necessary consideration in evaluating any proposed intervention or mitigation.

The benefits fall into two distinct categories of value to young users. The first set of benefits shape the substantive experience a young person has on a service: what content they encounter, what communities they find, and how their learning, creative, and informational needs are met. The second set are functional and protective, operating largely in the background to make services usable, accessible, and age-appropriate. Both categories matter: interventions that address risks of personalization without accounting for its benefits risk foreclosing value that young people—particularly those with fewer offline resources and support—may rely on.

A. YOUTH-FACING BENEFITS

Young people interact with personalized systems across nearly every digital context they inhabit—and in many of those contexts, personalization is what makes the experience substantively valuable rather than merely functional. The benefits described below are not incidental features of digital services; they reflect what personalization, when well-designed, can uniquely offer young users. Those benefits are examined here through three of the key use case contexts detailed in the previous section: information and discovery, connection and community, and learning and development.

Information and Discovery

The information and discovery use cases described in the previous section—search, streaming, and publisher content—share a common function: organizing vast information environments into something navigable and relevant for a specific user. For young people, that organizing function carries stakes that go beyond convenience. Access to information is not uniform across young people's lives—offline resources, trusted adults, and institutional support vary widely depending on family context, geography, and community. For a teen exploring questions about their health, identity, or mental well-being where the social cost of asking a parent, teacher, or peer is high, personalized systems may be among the most accessible paths to accurate, age-appropriate resources.[9] That is not a marginal use case—it is a consequential one that broad interventions targeting personalization data practices should account for.

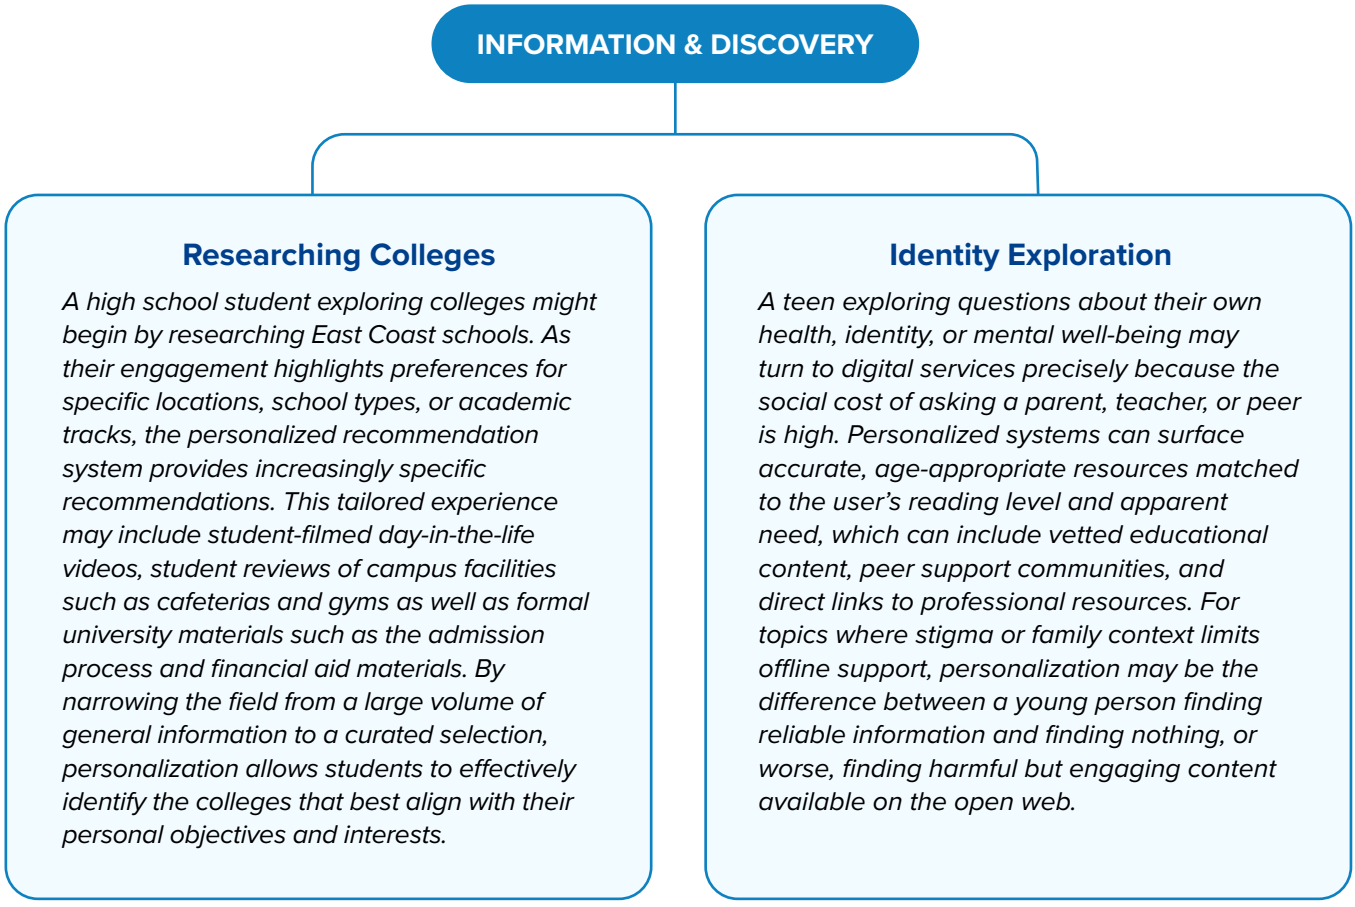

Connection and Community

The connection and community use cases described in the previous section, like social platforms and gaming communities, illustrate how personalization shapes not just what young people see but who they find and connect with. For many young people, that distinction is significant due to obstacles in geography, language, and identity. Personalization can open avenues for connection and belonging that would otherwise be out of reach: a peer community, an identity-affirming space, or an audience for creative work with no local equivalent.

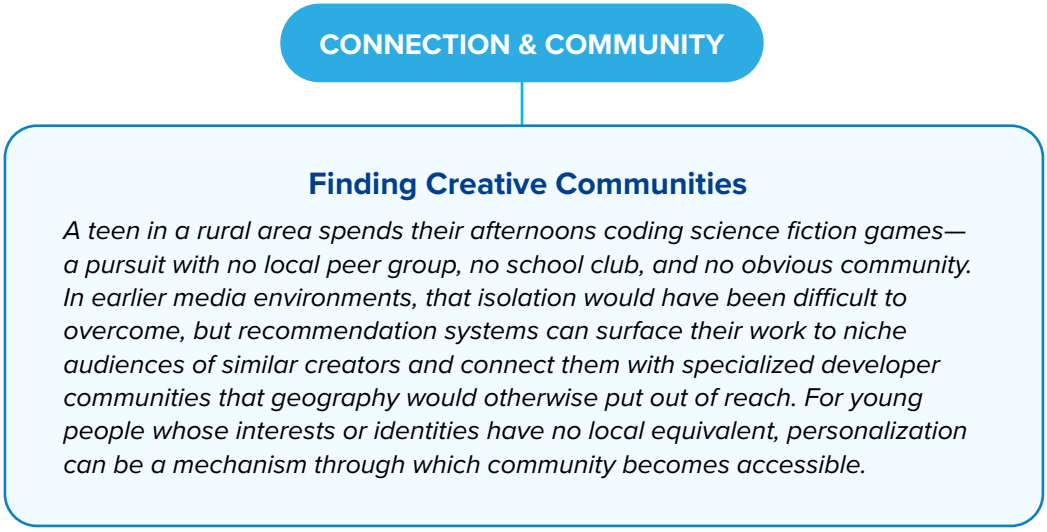

Learning and Development

The learning and development use cases described in the previous section illustrate how personalization adapts content, pacing, and support to individual students in ways generic instruction cannot. Research on adaptive learning systems generally finds that personalization to a student's pace and level improves outcomes as compared to one-size-fits-all instruction.[10] For students who need additional support, personalized tools can identify where understanding breaks down and adjust accordingly; for students ready to move ahead, they can extend the pace and depth of instruction to meet individual needs that a single teacher managing a full classroom may not be able to address.[11]

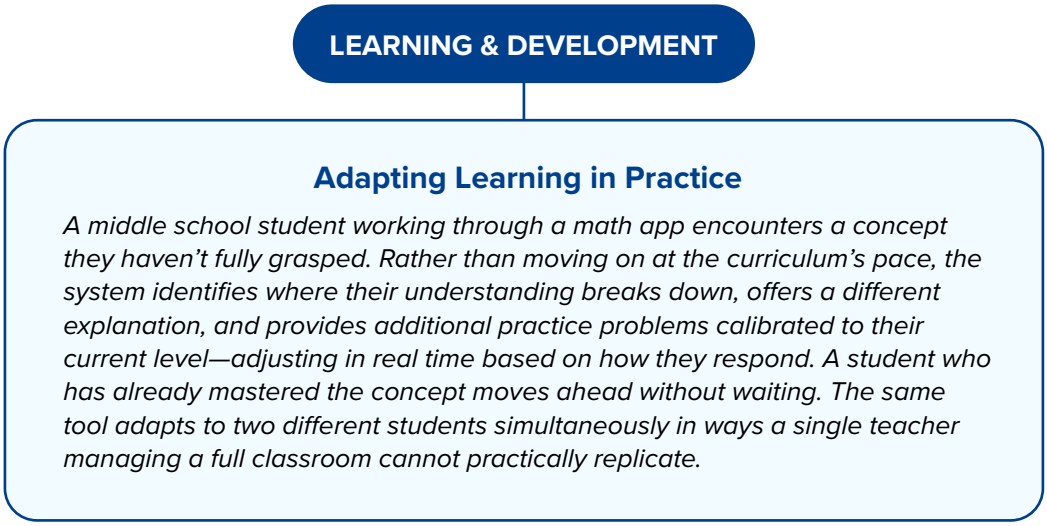

B. SUPPORTIVE FUNCTIONALITY AND PROTECTIONS

Alongside the visible, experiential benefits personalization delivers to young users, a second category of benefits operates in the background—not shaping what young people choose to engage with, but determining whether digital services are usable, safe, and appropriate for them at all. These functional and protective implementations are often the least visible in policy debates, yet they are among the most consequential: they enable platforms to distinguish a 13-year-old from an adult user, apply more protective default settings for younger users, surface age-appropriate content while filtering out harmful material, and make services accessible to young people with different abilities and needs.

Broad interventions that restrict data collection and use without distinguishing these implementations from higher-impact ones risk an important unintended consequence: disabling the very protections they are designed to strengthen. The two functional and protective benefits examined below—age-appropriate experiences and protections, and accessibility and inclusion—illustrate how personalization operates in this capacity and why that distinction matters for policy.

Age-Appropriate Experiences and Protections

As noted across the use cases described in the previous section—from content and social platforms to gaming, search, and AI systems—the ability to deliver age-appropriate experiences is a recurring function of personalization across nearly every digital context young people inhabit. Enabling platforms to customize safeguards, content filters, default settings, and safety features to match a user's age draws on demographic signals, explicit data, and contextual indicators to implement protections appropriate for younger users. Without that data, platforms cannot meaningfully distinguish a 13-year-old from an adult user, for instance, and the protections that follow from that distinction, such as more conservative default privacy settings, filtered content, limited exposure to higher-risk social interactions, cannot be applied.

AGE-APPROPRIATE EXPERIENCES & PROTECTIONS

Age-Appropriate Defaults in Practice

A 13-year-old joins a platform that identifies them as a minor and applies more protective default privacy settings, limits their exposure to mature content, and restricts higher-risk social interactions by default. When they attempt to adjust those settings, they receive just-in-time messaging explaining how each change affects their privacy and visibility. Those protections are not static rules applied uniformly—they are personalization functioning as a safeguard, using what the platform knows about the user to shape an experience appropriate to their age and developmental stage.

Accessibility and Inclusion

Unlike the youth-facing benefits examined above, accessibility and inclusion benefits are not concentrated in specific use case contexts—they cut across nearly every digital environment young people inhabit. Personalization adapts how information is presented and how users interact with a service to support different language needs and physical, cognitive, and sensory abilities—making digital services usable for young people whose needs are not met by generic, one-size-fits-all design.[12]

For young people with disabilities or language differences, accessible design is not an enhancement to an otherwise adequate experience—it is what determines whether participation is possible at all.[13] Personalization that identifies their needs and adapts accordingly is what makes inclusion a supportive functionality rather than an incidental benefit.

ACCESSIBILITY & INCLUSION

Personalization as Accessible Design

A teen who is hard of hearing creates an account on a streaming platform and, through a combination of stated preferences and observed behavior, signals their accessibility needs to the system. Rather than requiring them to manually filter for captioned content each time they search, the personalized system learns from those signals and automatically surfaces captioned content by default across their experience. A visually impaired teen using the same platform receives a parallel adaptation—image descriptions and screen-reader-compatible content prioritized automatically based on their device settings and interaction patterns. For both users, personalization is what converts an accessibility feature that exists on the platform into one that works for them specifically and automatically—reducing the burden of manually navigating a system not designed with their needs as the default.

IV. Personalization Risks in Youth Contexts

The same mechanisms that deliver the benefits described in the previous section can, depending on how they are designed and the purposes they serve, also produce significant risks for young people. Those risks are not uniform—they vary depending on how systems and features are designed, what data they use, and the outcomes they are intended to achieve. They are also not isolated to specific service contexts: as the use cases described earlier illustrate, personalization operates across content, advertising, and AI-driven services simultaneously, and many of the risks identified below can arise across multiple contexts rather than within a single one.

The table below outlines six common risk themes, describing the practices that give rise to each and how those risks may manifest specifically for young users.[14] The section that follows examines proposed mitigations for each category—and the tradeoffs those mitigations introduce for the benefits personalization can also provide.

RISK CATEGORY/ THEME	PRACTICES	THE RISKS FOR YOUNG PEOPLE
Privacy and Profiling	Collection and use of excessive or unnecessary personal data, including data from unexpected sources, and the combining or use of that data to personalize experiences which may lack transparency regarding how that data is collected, combined, or used.	May be profiled based on inferred characteristics—interests, behaviors, emotional states, vulnerabilities—in ways they cannot perceive or contest. Those profiles can persist across platforms and over time, with data collected during childhood and adolescence potentially informing consequential decisions—college admissions, employment, financial services—long after the behaviors or characteristics that generated it are no longer representative of who that person is.
Content Exposure and Amplification	Use of algorithmic ranking and recommender systems that utilize a user's personal data to suggest, promote, or rank content. This content includes material from other users, hashtags, and other media.	May contribute to the formation of "filter bubbles," which can limit exposure to diverse ideas and reinforce existing preferences, potentially leading to a distorted view of reality. May also progressively surface more extreme or harmful material through linked content sequences, or "rabbit holes"—a risk heightened for young people, whose developing capacity for critical evaluation makes them more susceptible to the influence of repeated, escalating content exposure.
Manipulation and Autonomy	Use of design elements to influence user behavior, including personalized nudges, time-pressure signals, and sycophantic response patterns that provide agreeable outputs to retain engagement."[15]	May steer young people toward interactions, disclosures, or decisions they do not fully understand or might not otherwise make. Particularly in AI chatbot contexts, personalized systems designed to be emotionally responsive and agreeable may reinforce harmful thought patterns, encourage sensitive disclosures, or escalate conversations in directions that contribute to self-harm or harm of others.

RISK CATEGORY/ THEME	PRACTICES	THE RISKS FOR YOUNG PEOPLE
Disruptions to Well-Being	Use of personalized design elements that drive engagement and continuous use such as message notifications, infinite scroll, and platform feedback signals.	May contribute to prolonged use, repeated feed-checking, and difficulty disengaging—disrupting sleep and crowding out offline activities in patterns that some researchers and policymakers have characterized as addictive. Young people, whose neurological development makes self-regulation and impulse control more difficult, face heightened vulnerability to design patterns engineered to sustain engagement.
Blurred and Persuasive Commercial Practices	Use of automated systems to collect and analyze user data; infer preferences, interests, or behaviors; and use those inferences to tailor advertisements or other commercial content.	May blur the line between content and marketing, making it harder for young users to identify commercial messages, including when embedded in influencer content. It may also take advantage of developmental vulnerabilities at a stage when they may be less able to recognize or resist persuasive marketing influences. In addition, when sensitive personal data is used to target ads, it may result in advertising that reveals information they did not intend to share (e.g., health, sexual orientation).
Misrepresentation, Discrimination, and Bias	Use of biased training data or algorithmic processes that generate inferences about user characteristics in ways that reinforce bias or produce unfair outcomes.	May reinforce stereotypes and produce inequitable outcomes—surfacing lower-quality content, fewer opportunities, or less accurate information to young users from marginalized groups. For young people, biased personalization during a formative period of development can shape educational trajectory, self-perception, and access to opportunity in ways that extend well beyond the digital environment.

Note: *Information compiled in this draws upon a series of closed stakeholder convenings discussing potential risks of personalization practices.*

The risks identified above are not uniformly distributed across personalization contexts—some arise primarily in content and social environments, others in advertising or AI-driven services, and many cut across multiple contexts simultaneously. That variation is important for determining mitigation strategies, where interventions designed to address one risk category may be poorly suited to another, and measures that effectively constrain higher-impact implementations may inadvertently foreclose the functional and protective uses of personalization described in the benefits section. The section that follows examines proposed mitigations with those tradeoffs in view.

V. Strategies for Risk Mitigation

No single mitigation addresses the full range of personalization risks identified above—and any effective response must also account for the benefits that well-designed personalization delivers to young people. The same mechanisms that give rise to the risks identified in the previous section also underpin the functional, protective, and youth-facing benefits described earlier—a tension that runs through every mitigation discussed below, where measures designed to address one risk category may amplify another, place implementation burdens on parties poorly positioned to bear them, or foreclose beneficial uses of personalization alongside riskier ones.

To help stakeholders navigate that complexity, the discussion below examines mitigations across four areas:

(A) Alternatives to Personalized Recommender Systems;

(B) Safety and Well-Being;

(C) Data Governance and Transparency; and

(D) Activity-Specific Considerations in Advertising and AI

Each entry identifies the tradeoffs a given mitigation introduces for the benefits personalization can also provide.

A. ALTERNATIVES TO PERSONALIZED RECOMMENDER SYSTEMS

Algorithmic recommendation systems are among the most studied and contested aspects of personalization in youth online experiences. The three approaches detailed below each target the risks of content exposure and amplification and disruptions to well-being identified in the previous section. The first two approaches—chronological and non-personalized feeds, and user-centered personalized algorithms—propose alternative designs for how content is selected and ranked. The third approach, concerning restrictions on design features like infinite scroll and autoplay, targets the mechanisms that amplify algorithmic curation rather than the curation itself. Each reduces reliance on personal data to varying degrees, but each also introduces tradeoffs for the benefits that well-designed recommendation systems deliver.

1. ***Chronological and Non-Personalized Feeds*** display content by publication date and time, or by popularity, recency, or location rather than individual behavioral data—reducing reliance on the engagement-driven algorithmic curation that gives rise to risks of content exposure and amplification and disruptions to well-being.[16] Both were common design elements of early online services and have been proposed by policymakers as alternatives to engagement-driven algorithmic curation.[17] Research suggests, however, that these approaches do not eliminate the harms associated with algorithmic curation, and they may increase exposure to spam, abuse material, and misinformation that algorithmic systems would otherwise suppress.[18] They also make it difficult to find relevant content in the large and complex information environments that have grown exponentially since the rise of those early online services. Critically for youth contexts, both approaches limit a platform's ability to apply age-appropriate safeguards, since those protections depend on knowing something about the user. Even non-personalized feeds require some personal data—such as location, device type, or IP address—to deliver a minimum quality experience, meaning the distinction between personalized and non-personalized is one of degree rather than kind.

2. ***User-Centered Personalized Algorithms*** reorients algorithmic curation around explicit user feedback and deliberate preference signals rather than implicit behavioral data, addressing risks of content exposure and amplification and manipulation and autonomy.[19] Explicit preference selection may not, however, capture the full range of a user's interests, and adoption rates for ongoing feedback mechanisms tend to be low.[20] For young people, the cognitive demand of actively managing personalization preferences may be a barrier, and preferences set at account creation may not reflect how interests evolve over time.[21]

User-Centered Personalized Algorithm in Practice

When creating a new account on Substack, users are prompted to select content categories that shape personalized recommendations provided within their account. Prompting users to select topic preferences at account setup—which then directly informs content recommendations—puts users in control of their digital experience.

Screenshot taken May 20, 2026.

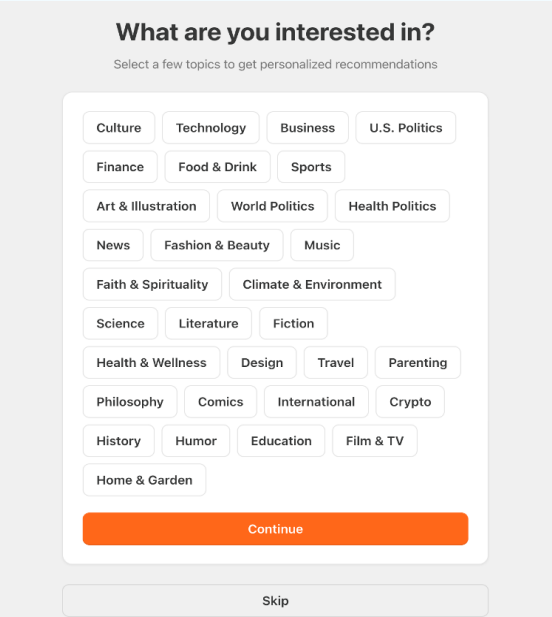

3. ***Restrictions on Design Features*** such as restrictions on infinite scroll, autoplay, and algorithmically triggered notifications—may address risks of disruptions to well-being and manipulation and autonomy by targeting the mechanisms that extend and intensify exposure to personalized content. Depending on how broadly these terms are defined, some captured features may also serve functional and accessibility purposes: reduced-friction content delivery or web access supports aid users with certain motor or cognitive needs, and autoplay can support passive consumption appropriate for educational or long-form content.[22] Stakeholders considering feature restrictions may therefore wish to distinguish between engagement-driving applications in social and content contexts and functional applications in educational or accessibility contexts.

Beyond Design Alternatives

Taken together, the design alternatives examined in this section represent meaningful but partial responses to the risks of algorithmic recommendation systems. None eliminates the underlying tension between personalization's capacity to surface relevant, age-appropriate content and its potential to amplify harmful or engagement-driven material. The safety and well-being mitigations examined in Part B address complementary dimensions of that tension through user-level controls and platform-level protections.

B. SAFETY AND WELL-BEING

While the design alternatives examined in Part A target the architecture of recommendation systems, the mitigations below operate at the level of user experience and platform obligation—providing individuals, families, and platforms with tools to manage how personalization affects young people's safety, well-being, and autonomy.

The five approaches below address risks of privacy and profiling, manipulation and autonomy, disruptions to well-being, and blurred commercial practices. They range from parental oversight mechanisms and user data controls to preference selection and digital literacy initiatives—each targeting a different dimension of how personalization can affect young users, and each carrying distinct tradeoffs for the benefits personalization can also provide.

1. ***Parental Oversight and Control*** enables parents to understand and manage how personalization operates in their child's digital experience. This may include parental consent processes that limit an entity's collection and use of their child's data for personalization purposes, or controls around content filters, data use restrictions, and time limits. Though this approach looks to address risks of privacy and profiling, manipulation and autonomy, and disruptions to well-being, effectiveness may vary based on parental digital literacy and family context. In cases where parental and child interests diverge, the same controls can foreclose the protective and identity-supportive functions of personalization for youth in unsupportive or abusive family environments. Stakeholders considering default parental control mandates should account for contexts where such defaults could reduce rather than enhance a young person's safety and well-being.[23]

2. ***Control Over Personal Data Use*** enables individuals to delete, limit, or specify how their data is collected and used for personalization purposes. This may include the ability to restrict data processing to specified activities, limit data sources, or adjust explicit preferences about how personalization operates. Its effectiveness depends on user awareness of available rights and the data practices of the platforms they use. For young people specifically, controls that exist on paper may go unused without meaningful, accessible notice about what data is collected and how it shapes their personalized experience.

3. ***Content Filtering and Preference Selection*** gives users more explicit control over what content is displayed or prioritized by a personalized service. By giving users direct input into what personalization surfaces, this approach reduces the degree to which algorithmic systems can expose young people to harmful, age-inappropriate, or commercially persuasive content without their awareness. people to harmful, age-inappropriate, or commercially persuasive content without their awareness. However, optional controls have generally seen lower adoption; and user-directed filters can reinforce narrower information environments which could restrict a young person's access to content that supports identity development, health information, and community connection. Platform-level filtering addresses some of these limitations but introduces questions about who determines what content is age-inappropriate and how those judgments account for variation across families, cultures, and communities.

4. ***Increasing Digital Literacy*** aims to equip young people with a better understanding of how personalization operates and how to manage their digital experiences more deliberately. It looks to address risks of manipulation and autonomy and blurred commercial practices by reducing susceptibility to design features, which could include education about how algorithmic systems work, how personal data is collected and used, and what tools are available to manage personalization settings. Nonetheless, digital literacy is not a standalone solution—it addresses awareness without constraining the underlying design features that generate risk—and is most effective when paired with structural mitigations rather than treated as a substitute for them.

5. ***Measures to Improve Digital Well-Being*** provide users with tools to manage their engagement with personalized services, such as with time-use limits, sleep and nighttime protections, notification controls, and the disabling of engagement signals such as likes and follower counts for youth accounts. These tools address the engagement-driving dimension of personalization by giving users mechanisms to disengage more easily. The primary limitation is reliance on users' willingness to adopt and consistently use available tools. Making well-being features defaults rather than opt-ins can improve adoption, though default settings alone do not address the underlying design features that generate engagement-driven risk.

SAFETY AND WELL-BEING MITIGATIONS
Scope and Limitations

The mitigations examined above address personalization risks at the level of user experience and platform obligation—providing families, users, and platforms with tools to manage how personalization affects young people's safety, well-being, and autonomy. Most, however, place meaningful implementation burdens on users and parents without constraining the underlying data practices and system design that generate those risks in the first place. The data governance and transparency mitigations examined in Part C address those upstream dimensions directly.

C. DATA GOVERNANCE AND TRANSPARENCY

While the mitigations examined in Parts A and B address how personalization affects user experience and feed design, the measures below operate at a more foundational level by targeting the data collection and processing practices that make personalization possible in the first place. As the functional vs. higher-impact framework established earlier in this document makes clear, the risk profile of any personalization practice depends heavily on what data is collected and to what end—making data governance a particularly consequential lever for addressing personalization risks in youth contexts.

1. ***Data Minimization*** limits data collection, use, and disclosure to what is necessary for a specified purpose—restricting practices (except to the extent consent is provided) such as cross-platform behavioral tracking, collection of data from unexpected sources, and the combining of data categories to build detailed user profiles that enable higher-impact personalization. As a feature of most major privacy frameworks globally, data minimization is often paired with purpose

limitations that bar covered entities from repurposing personal data in ways that contradict the context in which it was collected. Its central tradeoff in the personalization context is that the same data inputs that enable engagement-driven personalization also support functional and protective uses, such as age estimation, accessibility accommodations, and the suppression of age-inappropriate content. Minimization rules applied without reference to the purpose of processing can foreclose beneficial personalization alongside higher-risk applications; purpose-based limitations may better preserve protective uses while constraining riskier ones.

2. ***Transparency*** provides users with meaningful insight into what data is collected, how it is used for personalization, and what controls are available, addressing risks of privacy and profiling and blurred commercial practices by giving users visibility into the data practices that drive their personalized experience. Lengthy, technical, or poorly timed disclosures are unlikely to be read or understood—particularly by young users—and transparency alone does not constrain the underlying data practices it describes. For young people specifically, meaningful transparency requires disclosures designed to be comprehensible and actionable for a younger audience—a design challenge that existing transparency frameworks do not always adequately address.

DATA GOVERNANCE
A Necessary Foundation for Other Mitigations

Data governance and transparency measures address personalization risks at their source—constraining the data practices that make higher-impact personalization possible. The commerce and marketing mitigations examined in Part D address a specific and consequential application of those data practices: the use of personal data to target advertising to young people.

D. ACTIVITY-SPECIFIC CONSIDERATIONS FOR ADVERTISING AND AI

The mitigations below address the specific risks that arise from two specific and consequential applications of personalization—commercial advertising and conversational AI—rather than from personalization practices generally. Each targets a distinct context in which the youth-specific stakes of personalization are particularly acute and where the risk profile differs meaningfully from the broader content and discovery contexts examined in Parts A through C.

1. ***Advertising Restrictions and Disclosures:*** Limiting personalized advertising to contextual advertising—based on search queries, content currently being viewed, coarse location, and basic technical and contextual data, rather than personal data collected over time or across services—addresses risks of blurred and persuasive commercial practices and privacy and profiling. Contextual advertising restrictions are commonly proposed alongside complementary measures: clear labeling of ads and marketing content so young people can distinguish between content and advertising, filtering of age-inappropriate advertising, and avoidance of dark patterns that obscure commercial intent. However, contextual advertising does not eliminate commercial persuasion: a separate line of policy argument holds that commercial persuasion directed at minors raises concerns, regardless of whether it relies on personal data, and that

contextual advertising still uses the signal of a minor's presence on a particular page or service. Policymakers considering this argument must also weigh its potential impact on the role of advertising in funding free or low-cost services that minors rely on, including educational tools and content from public-interest publishers.

2. ***Responsible Data Practices and Safety-by-Design in AI:*** Safety-by-design addresses risks of manipulation; threats to autonomy, and privacy; potentially harmful profiling; and misrepresentation and bias by building protective measures into AI systems from the outset. Core measures include informing users about personalization features and controls, limiting data collection to what is necessary, and designing systems to avoid inferring sensitive characteristics in ways that reinforce bias. For youth specifically, these measures should account for the distinct risks emotionally responsive AI systems pose during adolescence: disclosing that the system is not human to prevent parasocial attachment, avoiding age-inappropriate content, incorporating interaction time limits to reduce dependency, and providing support resources when conversations touch on self-harm or harm to others. Systems should be tested through red-teaming to identify features that may encourage unhealthy emotional attachment.[24] Each of these measures nonetheless carries tradeoffs: restrictions on sensitive characteristic inference can limit age-appropriate experiences; strict content filtering can foreclose access to health and crisis information that youth in restrictive offline environments rely on; and on-device data processing protects privacy but may reduce model quality—constraints that risk shifting rather than eliminating harm.

3. ***Mental Health and Crisis Response in Chatbots:*** Conversational and generative AI systems—given their personalized and emotionally responsive design—may become a point of contact for young people experiencing distress or seeking mental health support.[25] Designing these systems to identify signals of distress and direct young users to timely support resources, including crisis helplines and trusted support networks may help reduce risks to well-being. These measures rely, however, on the accuracy of algorithmic distress detection—which may produce false positives or miss genuine signals of harm—and raise questions about data handling in crisis reporting contexts. Responses that involve third parties, including parents and law enforcement, may not always serve a young person's interests, particularly where family dynamics are complicated or abusive.[26]

AI MITIGATION: Managing Risk Judiciously

Effective safeguards for AI personalization require careful attention to which constraints reduce overall risk and which simply shift it. Restrictions that prevent one harm—biased inference, inappropriate content, emotional dependency—can foreclose protections or resources that young people, particularly those in restrictive offline environments, may rely on AI systems to provide. Mitigation design in this context is not a checklist but a calibration.

The mitigation strategies examined above reflect the range of approaches currently being considered—and in some cases enacted—by legislators and regulators globally, whose responses to personalization risks in youth contexts are examined in the section that follows.

VI. The Emerging Policy and Regulatory Response

Despite renewed energy around youth online safety, regulation of personalization practices is not entirely new. Existing privacy and data frameworks have long imposed constraints on how personal data can be collected and used in ways that directly affect personalization. What is new are recent proposals that expand personalization requirements in product and service design—and their associated data practices—further than before, targeting specific features and design elements through which personalization operates, and aiming to address online safety risks including harms to minors' mental health and well-being.[27]

The key policy and regulatory responses examined in this section reflect, in large part, the risks of privacy and profiling, content exposure and amplification, manipulation and autonomy, disruptions to well-being, and blurred commercial practices identified earlier in this document. This section examines that landscape across four areas: existing privacy and data minimization frameworks; emerging legislation targeting personalization design and features; federal and state regulatory activity in the United States; and enforcement activity in the European Union under the Digital Services Act (DSA).

A. EXISTING FRAMEWORKS: PRIVACY AND DATA MINIMIZATION

Existing privacy frameworks establish the foundation on which emerging personalization regulation builds. Youth and comprehensive consumer privacy laws in the U.S. commonly restrict the use of a minor's personal data for profiling or targeted advertising without proper consent or authorization from parents.[28] Others, like Maryland's Online Data Protection Act (MODPA), go a step further by banning targeted advertising to and profiling of known minors.[29]

Many frameworks also include data minimization requirements that restrict a covered entity's collection, use, and retention of personal data to what is necessary for a lawful purpose, and pair these restrictions with "purpose limitations" that bar covered entities from repurposing personal data in incongruous or unexpected ways. States, like Maryland and California, tie these standards to whether data practices are "reasonably necessary and proportionate" to the requested service or align with consumers' "reasonable expectations."[30] Internationally, global privacy regulations similarly include limitations on data collection and use. For example, the GDPR limits data collection to only what is directly relevant and necessary to accomplish a specified purpose and requires that data is only processed in accordance with a specific lawful basis.[31] Additionally, comparable to Maryland's ban on targeted advertising, the DSA prohibits display of ads through profiling based on the data of minors.[32]

B. EMERGING REGULATION TARGETING PERSONALIZATION DESIGN AND FEATURES

Emerging personalization regulations range from outright bans on access to certain services to restrictions on specific personalized features and design elements, often paired with transparency mandates, parental control obligations, and user setting requirements. At the most restrictive end, Australia and Indonesia have implemented regulations prohibiting social media platforms from allowing minors under 16 to create or maintain accounts, citing in part harms associated with personalization practices and design features.[33] At a Global Age Assurance Summit in 2025, Australia's eSafety Commissioner specifically referenced algorithmic manipulation and design features like infinite scroll as factors behind the prohibitive approach—language that maps directly onto the manipulation and autonomy and disruptions to well-being risk categories identified earlier in this document.[34]

Other laws condition minors' access to certain design elements on user or parental consent rather than restricting access unconditionally. California and New York have enacted laws restricting covered platforms from offering minors algorithmically curated feeds without first obtaining parental consent—reflecting the parental oversight and user-centered personalization mitigations examined in the previous section.[35] Other jurisdictions require platforms to offer opt-outs—or for minors, opt-ins—for recommendation systems not based on explicit user preferences, mandate transparency around personalization practices, or provide at least one alternative recommender system option not based on profiling.[36]

C. UNITED STATES—FEDERAL AND STATE REGULATORY ACTIVITY

Beyond enacted legislation, regulatory and enforcement activity has shaped the legal landscape around personalization in youth contexts—reflecting growing institutional attention to the risks of engagement-driven algorithmic design and data-intensive personalization practices.

On the federal level in the US, the FTC held a workshop in June 2025 examining the "attention economy," which included, in part, consideration of "addictive" algorithms or design features like infinite scroll and auto-play in digital services.[37] The workshop featured calls from lawmakers to enact legislation such as the Kids Online Safety Act (KOSA) to restrict certain personalization practices, with specific concern raised about algorithms designed to drive minors' engagement in ways that harm their mental health.[38] Likewise, while Chairman Ferguson reaffirmed the FTC's commitment to enforcing COPPA violations related to unlawful data practices in this domain, other speakers urged lawmakers to grant the agency broader authority to combat "predatory behaviors" affecting children's online safety—including use of "algorithms that [do not] have [childrens'] best interests in mind."[39]

On the state level, Attorneys General (AGs) have been especially active in recent years, employing novel efforts to enforce against alleged online safety deficiencies through consumer protection laws. Through these lawsuits, state AGs broadly argue that companies engaged in unfair and deceptive trade practices by misrepresenting the safety of digital services even though they were aware of high risks of harm—including from "addictive" algorithms purportedly designed to excessively drive minors engagement, among other harms—that were not sufficiently mitigated.[40] Dozens of lawsuits following this litigation formula have been filed against a variety of technology platforms, including social media platforms, AI companies, gaming companies, and streaming services.[41]

D. EUROPEAN UNION—DSA ENFORCEMENT

DSA enforcement has also contributed to rising regulatory pressures over platform personalization practices. For example, in a 2025 ruling, a Dutch court found that a social media platform violated the DSA's requirement to provide users with the ability to opt-in to a non-personalized content feed.[42] Although the court acknowledged the company's design changes, which attempted to comply with the DSA across its platforms and services, the court found that the company ultimately failed to ensure the mechanisms were "direct and easily accessible" to users across all platforms and applications.[43] More recently, in February 2026, the European Commission preliminarily found another social media platform in violation of the DSA, in part due to the design of its "highly personalised recommender system," which the Commission described in a press release as one of several "key addictive features" on the platform.[44] The Commission's investigation in this case remains ongoing and no final outcome has yet been reached.

With a range of global legislative and regulatory pressures and a variety of approaches to addressing harms from personalization, there is no clear consensus legislative or regulatory approach for addressing personalization risks to young people, creating a complicated and sometimes diverging patchwork of obligations, protections, and regulator expectations.

VII. Conclusion

Personalization is neither inherently beneficial nor inherently harmful to young people—its effects depend on what data is collected, how it is used, and the purposes it serves. The same mechanisms that enable harmful content amplification also enable adaptive learning, community discovery, and the age-appropriate protections that make digital services safer for younger users. That duality is why it is so important to carefully calibrate regulatory responses so as to avoid foreclosing benefits alongside harms.

The risks are real and documented: content exposure and amplification, disruptions to well-being, manipulation, blurred commercial practices, and privacy concerns each represent meaningful harms that policy and product interventions can address. But the young people most likely to bear the cost of poorly calibrated interventions are often those who rely most heavily on personalization's benefits: youth in isolated communities, young people navigating identity without offline support, students in under-resourced schools, and young people with disabilities for whom adaptive design is a condition of participation, not a convenience.

No single mitigation addresses the full range of personalization risks, and each introduces tradeoffs for the benefits personalization can also provide. The most effective approaches share a common characteristic: they constrain higher-impact implementations while preserving the functional and protective implementations that serve young people's direct interests. The emerging legislative landscape reflects the genuine goal of many policymakers to address real risks—but also significant divergences that will persist absent a shared analytical framework for recognizing benefits alongside harms, and distinguishing among use cases that warrant stronger safeguards and those that do not.

Effective policy in this space requires moving beyond whether personalization should exist in youth-facing digital services toward more precise questions: how it is implemented, what data it draws on, and what purposes it serves. A risk-proportionate approach is better positioned to reduce harm without foreclosing the value that well-designed personalization delivers to young people.

Appendix A: Sample of Relevant Legislation and Regulations

The following Appendix supplements this report with a representative sample of laws and regulations that directly or implicitly restrict certain personalization practices with respect to minors. This chart does not encompass the entire universe of existing legal requirements, but rather provides a snapshot of the types of requirements and regulatory approaches affecting companies' personalization practices within digital experiences provided to minors. Where applicable, this table also notes important emerging legislative or regulatory trends related to the enacted legislation included in the table.

LEGISLATIVE APPROACH	BILLS
Laws or regulations **restricting personalized features and/or requiring alternative feed options and transparency disclosures** for minors.	Digital Services Act (DSA, 2023; requiring non-personalized feed options and restricting profiling of minors); NY S7694A ("SAFE for Kids Act," 2024; restricting minor access to personalized feeds without parental consent); CA SB 976 ("Protecting Our Kids from Social Media Addiction Act," 2023; restricting minor access to personalized feeds without parental consent); NE LB 504 (AADC, 2025); SC SB 3431 (AADC, 2026; requiring transparency regarding how information is provided in recommendation feeds and opt-out processes); MN HF 4138 (2026, requiring feed transparency); ID H.542 (2026, prohibiting profile-based advertising within feeds and use of "addictive interface features" such as personalized feeds); SC HB 4591 (2026, prohibiting profile-based advertising within feeds and use of "addictive interface features" such as personalized feeds); CT SB 5 (2026, prohibiting personalized social media feeds unless parental consent is provided).
Frameworks with data protections that impact or restrict various personalization practices, such as targeted advertising, profiling, and/or limitations on data collection and use.	US State Comprehensive Privacy Laws (captured in linked resource, current up to December 2025; covering data minimization standards and profiling/targeted advertising restrictions); GDPR (EU, 2016); Brazil's General Law on the Protection of Personal Data (LGPD, 2018); COPPA; NY S7695B (New York Child Data Protection Act, 2024); AR HB 1717 ("Arkansas's COPPA 2.0," 2025); SC SB 3431 (AADC, 2026; restricts data collection and use to only the specific elements of the covered online service with which a minor has "knowingly engaged"—this is a developing standard in minor privacy and online safety legislation). Some proposed US state legislation in 2026 would have taken transparency requirements and restrictions on minor data use a step further—for example, frameworks proposed in Kansas (SB 499) and Georgia (SB 495) would have required covered businesses to prominently disclose how minors' data is used across service features and algorithmic recommendation systems, restricted the use of minors' personal data for selecting or prioritizing media in online services, and permitted collection and use of minors' personal data only to the extent necessary to provide a service with which a minor was "actively and knowingly engaged."
Laws and/or regulations **restricting minors' access to services**, in part due to concerns regarding harms of algorithmic or personalized tailoring.	Digital ECA (Brazil, 2025); Australia Online Safety Amendment Act 127, 2024 (Social Media Minimum Age; prohibiting minors under 16 from holding social media accounts); Indonesia Government Regulation 17 of 2025 (Social Media Ban for Children, March, 2026). Several jurisdictions have recently considered proposals that included some level of prohibition on minor access to social media services, including the UK (ban on under 16s), France (ban on under 15s), Hawaii (ban on under 16s), and the US Congress (ban on under 13s).

ENDNOTES

1 Although user experiences may be adapted to ***relate to*** the user's history or behavioral patterns, the overall relation should be conceptualized more generally rather than as a direct connection between data or use patterns and personalized tailoring. In other words, content and experiences provided through implicit-behavioral personalization may share some common thread that relates to historical contexts and behavioral data sources, but can still differ significantly from those patterns.

2 U.S. Dep't of Educ., Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations 9–10, 18–23, 44 (2023), https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf.

3 Packer, J., H. Croker, A. L. Goddings, E. J. Boyland, C. Stansfield, S. J. Russell, and R. M. Viner. "Advertising and Young People's Critical Reasoning Abilities: Systematic Review and Meta-analysis." *Pediatrics* 150, no. 6 (December 2022): e2022057780. https://doi.org/10.1542/peds.2022-057780.

4 Gao, B., Wang, Y., Xie, H., Hu, Y., & Hu, Y. (2023). Artificial intelligence in advertising: Advancements, challenges, and ethical considerations in targeting, personalization, content creation, and ad optimization. *SAGE Open*, *13*(4), 1–20. https://doi.org/10.1177/21582440231210759

5 Internet Matters, *Me, Myself & AI: Understanding and Safeguarding Children's Use of AI Chatbots* 13–15, 25-42 (July 2025), https://www.internetmatters.org/wp-content/uploads/2025/07/Me-Myself-AI-Report.pdf.

6 OECD, *Explanatory Memorandum on the Updated OECD Definition of an AI System* 4 (2024), https://www.oecd.org/en/publications/explanatory-memorandum-on-the-updated-oecd-definition-of-an-ai-system_623da898-en.html.

7 Pew Research Center, *How Teens Use and View AI* 3-7 (Feb. 24 2026) https://www.pewresearch.org/internet/2026/02/24/how-teens-use-and-view-ai/.

8 Internet Matters, 13–15, 25-42 (July 2025).

9 Common Sense Media & Hopelab, *Getting Help Online: How Young People Find, Evaluate, and Use Mental Health Apps, Online Therapy, and Behavioral Health Information* (2024), https://www.commonsensemedia.org/sites/default/files/research/report/2024-getting-help-online-hopelab-report_final-release-for-web.pdf.

10 Angélique Létourneau et al., *A Systematic Review of AI-Driven Intelligent Tutoring Systems (ITS) in K-12 Education*, 10 *npj Sci. Learning* art. 29 (2025), https://www.nature.com/articles/s41539-025-00320-7.

11 John F. Pane et al., *Continued Progress: Promising Evidence on Personalized Learning*, RAND Corp. (Nov. 2015), https://www.rand.org/content/dam/rand/pubs/research_reports/RR1300/RR1365/RAND_RR1365.pdf.

12 World Wide Web Consortium (W3C), Web Accessibility Initiative, *Support Adaptation and Personalization*, Supplemental Guidance to WCAG 2 (Apr. 29, 2021), https://www.w3.org/WAI/WCAG2/supplemental/objectives/o8-personalization/.

13 Royal Society. *Disability Technology: How Data and Digital Assistive Technologies Can Support Independent, Fulfilled Lives,* 6–7, (2025). https://royalsociety.org/-/media/policy/projects/disability-technology/disability-technology-report.pdf

14 It is important to note, however, that the supporting research is not uniform and some risks have been more studied and documented than others. This table compiles common personalization practices and associated risks identified across enacted and proposed legislation, and in consultation with stakeholders. See table below for full list of sources.

15 Georgetown Law, Tech Institute, *Tech Brief: AI Sycophancy & OpenAI* (n.d.), https://www.law.georgetown.edu/tech-institute/research-insights/insights/tech-brief-ai-sycophancy-openai-2/.

16 *E.g.*, Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act), 2022 O.J. (L 277) 1, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R2065.

17 KGI Expert Working Group on Recommender Systems, *Better Feeds: Algorithms That Put People First* 18 (Mar. 2025), https://kgi.georgetown.edu/wp-content/uploads/2025/02/Better-Feeds_-Algorithms-That-Put-People-First.pdf; Appendix A.

18 *Id*.

19 *Id*. at 2. **Note:** it's common practice for platforms to offer some level of user-feedback methods within personalized products and services today (e.g., "like" buttons serve as a form of preference selection; content category selections at account setup allow users to control prioritization of content categories).

20 *Id*.

21 *Id*.

22 Lisa Seeman & Michael Cooper, *Cognitive Accessibility Roadmap and Gap Analysis, W3C Working Draft* (Dec. 11, 2018), https://www.w3.org/TR/coga-gap-analysis/#multi-modal-content-delivery.

23 Danielle Keats Citron & Ari Ezra Waldman, *Rethinking Youth Privacy*, 111 Va. L. Rev. 1429 (2025), https://virginialawreview.org/articles/rethinking-youth-privacy/.

24 Daniel Berrick & Stacey Gray, *Concepts in AI Governance: Personality vs. Personalization 17–20* (Sept. 2025), https://fpf.org/wp-content/uploads/2025/09/Concepts-in-AI-Governance_-Personality-vs.-Personalization.pdf.

25 Am. Psych. Ass'n, *Use of Generative AI Chatbots and Wellness Applications for Mental Health Purposes* (2025), https://www.apa.org/topics/artificial-intelligence-machine-learning/health-advisory-chatbots-wellness-apps.

26 See *supra* note 23, Citron & Waldman, Rethinking Youth Privacy, (2025).

27 *See* New York Attorney General, *Attorney General James Releases Proposed Rules for the Safe Kids Act to Restrict Addictive Social Media Feeds for Minors* (Jan. 15, 2025), https://ag.ny.gov/press-release/2025/attorney-general-james-releases-proposed-rules-safe-kids-act-restrict-addictive and Senator Andrew Gounardes, *Governor Signs Internet Safety Legislation* (Dec. 22, 2023), https://www.senatorgounardes.nyc/gov-signs-internet-safety-legislation. *See also* Julie Inman Grant, *A Fairer Fight: Protecting Childhood and Adolescence in a Digital World* (eSafety Commissioner Feb. 12, 2024), https://www.esafety.gov.au/newsroom/blogs/a-fairer-fight-protecting-childhood-and-adolescence-in-a-digital-world.

28 Children's Online Privacy Protection Rule, 89 Fed. Reg. 33814 (Apr. 22, 2025), https://www.federalregister.gov/documents/2025/04/22/2025-05904/childrens-online-privacy-protection-rule; N.Y. S. 7695-A, 2023–2024 Leg. Sess., https://www.nysenate.gov/legislation/bills/2023/S7695/amendment/A ; Jordan Francis, *Anatomy of a State Comprehensive Privacy Law* (Future of Privacy Forum 2025), https://fpf.org/wp-content/uploads/2025/12/FPF-Anatomy-of-a-State-Comprehensive-Privacy-Law-Report.pdf; Regulation (EU) 2016/679 of the European Parliament and of the Council, 2016 O.J. (L 119) 1, https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng.

29 Md. S.B. 541, 2024 Gen. Assemb., Reg. Sess., https://mgaleg.maryland.gov/2024RS/bills/sb/sb0541E.pdf.

30 Jordan Francis, *Data Minimization: A Comprehensive Overview* (Future of Privacy Forum 2025), https://fpf.org/wp-content/uploads/2025/06/FPF_Data-Minimization.pdf; Jordan Francis, *The Old Line State Does Something New on Privacy* (Future of Privacy Forum, June 2025), https://fpf.org/blog/the-old-line-state-does-something-new-on-privacy/; California Office of the Attorney General, *California Consumer Privacy Act (CCPA)*, https://oag.ca.gov/privacy/ccpa and Cal. Code Regs. tit. 11, § 7001 (Westlaw), https://govt.westlaw.com/calregs/Document/I7C5D69409E0F11F09C6BF97E55B516E3

31 Regulation (EU) 2016/679 of the European Parliament and of the Council, 2016 O.J. (L 119) 1, https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng.

32 Gabriela Zanfir-Fortuna & Vasileios Rovilos, EU's *Digital Services Act Just Became Applicable: Outlining Ten Key Areas of Interplay with the GDPR*, Future of Privacy Forum (Aug. 31, 2023), https://fpf.org/blog/eus-digital-services-act-just-became-applicable-outlining-ten-key-areas-of-interplay-with-the-gdpr/.

33 *See* eSafety Commissioner, Social Media Age Restrictions (2024), https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions and Indonesia Ministry of Communication & Informatics, *Permenkominfo No. 9 Tahun 2026* (2026), https://peraturan.bpk.go.id/Details/346040/permenkomdigi-no-9-tahun-2026 (regulation prohibiting minors under 16 from creating or maintaining social media accounts).

34 Julie Inman Grant, *A Fairer Fight: Protecting Childhood and Adolescence in a Digital World* (eSafety Commissioner Feb. 12, 2024), https://www.esafety.gov.au/newsroom/blogs/a-fairer-fight-protecting-childhood-and-adolescence-in-a-digital-world.

35 Cal. S.B. 976, 2023–2024 Leg., Reg. Sess., https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB976; N.Y. S. 7694-A, 2023–2024 Leg. Sess., https://www.nysenate.gov/legislation/bills/2023/S7694/amendment/A. .

36 KGI Expert Working Group, *supra* note 12; Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act), 2022 O.J. (L 277) 1, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R2065; Daniel Hales, *Paradigm Shift in the Palmetto State: A New Approach to Online Protection by Design* (May 23, 2024), https://fpf.org/blog/paradigm-shift-in-the-palmetto-state-a-new-approach-to-online-protection-by-design/

37 Federal Trade Commission, *Attention Economy: How Tech Firms Exploit Children* (June 5, 2025), https://www.ftc.gov/news-events/events/2025/06/attention-economy-tech-firms-exploit-children.

38 Federal Trade Commission, *Transcript: The Attention Economy—How Big Tech Firms Exploit Children and Hurt Families* 8 (June 4, 2025), https://www.ftc.gov/system/files/ftc_gov/pdf/Transcript-The-Attention-Economy-How-Big-Tech-Firms-Exploit-Children-and-Hurt-Families-06-04-25.pdf.

39 *Id*. at 4, 14-15.

40 Francis, et al., *U.S. Privacy Enforcement in 2025*, at 8–9 (Future of Privacy Forum 2026), https://fpf.org/wp-content/uploads/2026/02/FPF-U.S.-Privacy-Enforcement-in-2025.pdf.

41 *See id*., Kevin Grout, *AG Coleman Sues AI Chatbot Company for Preying on Children* (Ky. Att'y Gen. Jan. 8, 2026), https://www.kentucky.gov/Pages/Activity-stream.aspx?n=AttorneyGeneral&prId=1857 , and Texas Attorney General's Office, *Attorney General Ken Paxton Sues Netflix for Spying on Texas Kids and Consumers by Illegally Collecting Users' Data Without Their Knowledge or Consent* (May 11, 2026), https://www.texasattorneygeneral.gov/news/releases/attorney-general-ken-paxton-sues-netflix-spying-texas-kids-and-consumers-illegally-collecting-users.

42 *Stichting Bits of Freedom v. Facebook Netherlands B.V., Meta Platforms Ireland Ltd. & Meta Platforms Inc.*, District Court of Amsterdam, Private Law Div., C/13/774725 / KG ZA 25-687 (Oct. 2, 2025), https://www.bitsoffreedom.nl/wp-content/uploads/2025/10/20251002-District-Court-of-Amsterdam_en.pdf

43 *Id*. at 10-11.

44 European Commission, Press Release IP/26/312, *Commission preliminarily finds TikTok's addictive design in breach of the Digital Services Act* (Feb. 5, 2026), https://ec.europa.eu/commission/presscorner/detail/en/ip_26_312.

1350 EYE STREET NW, SUITE 350 | WASHINGTON, DC 20005 FPF.ORG | info@fpf.org