ITIF | INFORMATION TECHNOLOGY & INNOVATION FOUNDATION

How Personalization Drives Consumer Choice and Autonomy

DANIEL CASTRO | MAY 2026

As new technologies such as AI expand both user-directed and provider-driven personalization capabilities in digital systems, policymakers should ensure that personalization strengthens transparency, accountability, and user control rather than constrain its development.

KEY TAKEAWAYS

- Personalization in digital systems has evolved from simple user settings into advanced, context-aware AI systems. These technologies consistently help consumers manage complexity and reduce search costs.
- Consumers strongly favor digital systems that maintain context and remember preferences. Abstract privacy concerns rarely stop users from adopting tools that offer clear convenience.
- Strict regulatory approaches such as GDPR create heavy compliance burdens that stifle tech start-ups. This unintentionally concentrates market power and delays the rollout of new AI features.
- Policy frameworks should regulate actual harmful outcomes such as fraud and deception and avoid restricting core technology features such as memory, learning, and inference.

CONTENTS

Key Takeaways 1
Introduction 3
The Evolution of Personalization in Digital Technology 4
Stateless Computing and Early Persistence (Pre-1990s) 4
The Customization Era (1990s–2000s) 7
The Algorithmic Era (2000s–2010s) 8
The Predictive Era (2010s–Present) 10
Personalization as Consumer Empowerment 11
The Recurring "Personalization Panic" Cycle 13
Global Approaches To Governing Personalization 14
The Next Phase: AI Agents 16
Governing Harm Without Stopping Progress 18
Endnotes 21

INTRODUCTION

A new generation of digital systems—driven by growing adoption of artificial intelligence (AI)—is making personalization more visible, persistent, and integrated into daily life. Unlike earlier digital services that operated largely within isolated applications, emerging AI tools increasingly access users' email, calendars, phones, browsers, productivity tools, commerce platforms, social media, and financial services. These systems can retain context across interactions, remember user preferences, and provide recommendations or assistance tailored to individual needs and habits. Some emerging wearable AI systems can even continuously capture and process visual and audio information from users' daily environments, allowing digital assistants to remember conversations, view encrypted messages, recognize objects and locations, and provide context-aware assistance throughout the day.[1]

Major technology firms are rapidly expanding these capabilities. In 2025, OpenAI expanded ChatGPT's memory features to allow the system to reference past conversations and user preferences across sessions.[2] In 2026, Google introduced "Personal Intelligence" and memory tools in Gemini that connect across Gmail, Photos, Search, YouTube, and other services to deliver more context-aware assistance.[3] Microsoft has similarly announced memory and personalization features for Copilot that allow the assistant to retain information about users' preferences, workflows, and prior interactions.[4]

These developments have spurred debate about the future of personalization.[5] Critics often frame increasingly context-aware AI systems as a form of pervasive surveillance that could enable manipulation, reduce privacy, or concentrate excessive informational power in large technology firms.[6] Concerns surrounding AI memory, persistent context, and cross-platform integration reflect broader anxieties about how much digital systems should know about users and how that information should be governed.[7]

Digital services have long adapted to user preferences, from customized ringtones and homepage layouts to recommendation systems, search ranking, autocomplete, navigation apps, and streaming suggestions. What is changing is the depth, persistence, and breadth of personalization.

At the same time, consumer expectations are moving in the opposite direction. Many users increasingly expect digital systems to recognize them, remember prior interactions, and reduce repetitive tasks.[8] Consumers routinely express frustration when systems fail to retain preferences, lose conversational context, or require the same information to be repeatedly re-entered. The appeal of AI assistants is not simply automation; it is continuity. Systems that can remember dietary preferences, writing styles, travel habits, accessibility needs, or recurring work tasks can reduce friction and provide better user experiences.

This tension reflects two competing narratives about personalization. One views personalization primarily as a surveillance machine that extracts data to influence user behavior. The other views it as a digital concierge: a tool that helps users manage information overload, reduce search costs, and receive services that better match their individual needs. Both perspectives contain legitimate considerations. Personalized systems can create data governance challenges, particularly when data practices lack transparency or meaningful user control. But

personalization also delivers substantial consumer value by helping users find information, simplify decisions, and interact with increasingly complex digital environments more effectively.

Importantly, personalization is not a new phenomenon created by AI. Digital services have long adapted to user preferences, from customized ringtones and homepage layouts to recommendation systems, search ranking, autocomplete, navigation apps, and streaming suggestions. What is changing is the depth, persistence, and breadth of personalization. AI systems are evolving from largely stateless tools into systems capable of maintaining ongoing context across time and applications.

At its core, personalization is fundamentally about recognition and continuity. History shows that consumers generally prefer systems that remember prior interactions, anticipate needs, and reduce unnecessary repetition. The central policy challenge is therefore not whether personalization should exist, but how to ensure that it develops in ways that preserve transparency, trust, competition, and user control. Policymakers should focus on governance frameworks that empower consumers and address legitimate harms, without constraining the development of technologies that many users increasingly value and expect.

To achieve this balance and guide the future of context-aware technology, this report proposes a proportional governance framework built around six core policy principles:

- **Transparency.** Ensure that consumers are able to understand when personalization occurs and view significant data practices.
- **User Control.** Empower users to modify preferences, manage or delete memory features, and opt out when appropriate.
- **Interoperability and Portability.** Encourage data portability so profiles can transfer across services, thereby preventing ecosystem lock-in.
- **Outcome-Focused Accountability.** Target legal and regulatory scrutiny toward demonstrable harms rather than technical capabilities.
- **Innovation Flexibility.** Avoid rigid, preemptive restrictions that stifle competition and slow beneficial advancements.
- **Proportionality.** Distinguish low-risk features that improve everyday usability from higher-risk practices involving manipulation.

THE EVOLUTION OF PERSONALIZATION IN DIGITAL TECHNOLOGY

Stateless Computing and Early Persistence (Pre-1990s)

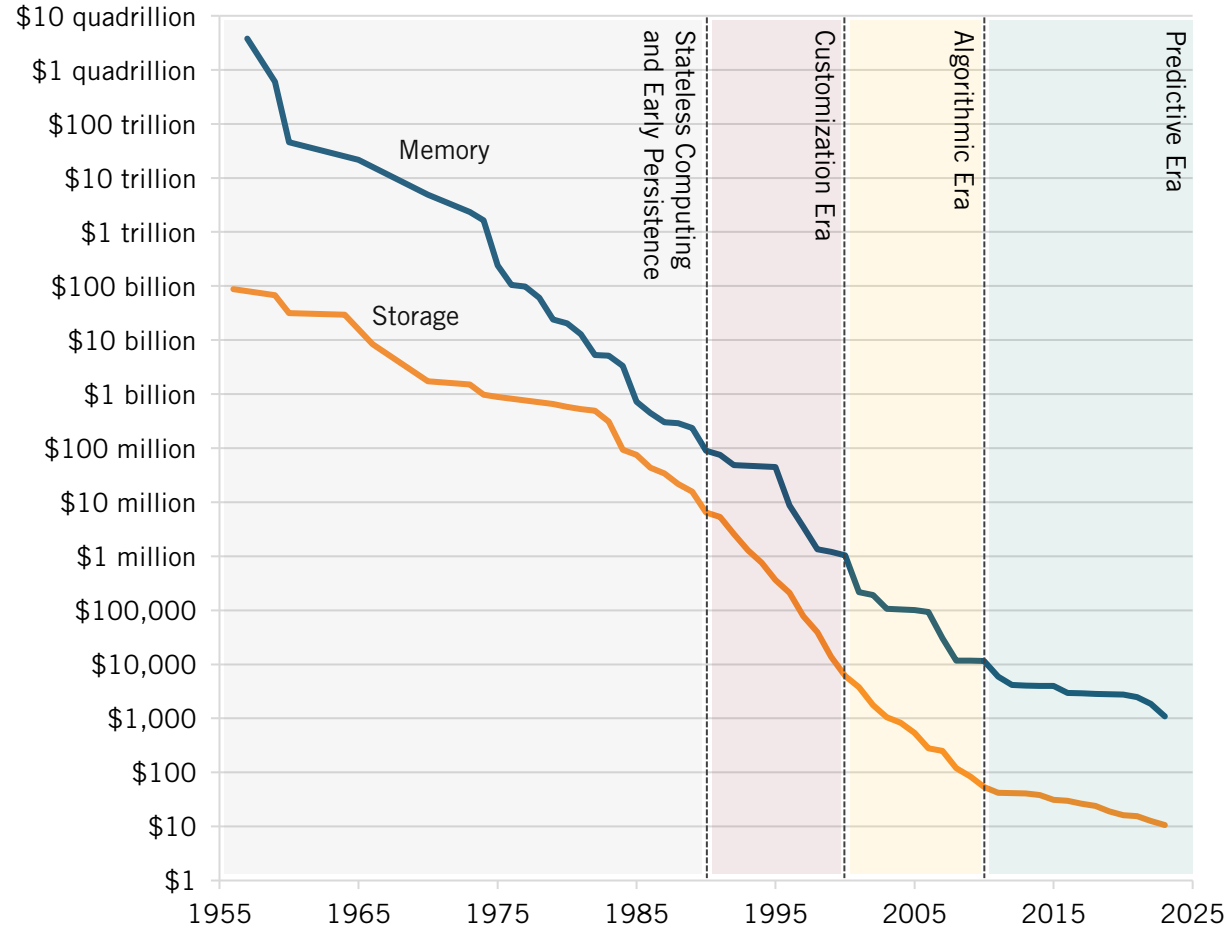
Personalization in digital technology has gradually evolved as computing systems have become more capable of retaining information about users across interactions. The history of personalization is therefore closely tied to the history of storage, memory, and processing.

Early computing systems were largely stateless. Mainframe and terminal-based systems treated each session as independent, with little or no capacity to retain individualized preferences over time. Computing resources were expensive, storage was scarce, and systems were primarily designed around shared institutional use rather than individualized consumer experiences. Users

often had to repeatedly enter commands, configure environments, and navigate identical setup processes every time they logged in.

Technical limitations heavily constrained early personalization. In the 1950s and 1960s, digital storage costs were extraordinarily high, making persistent user-specific information economically impractical at scale. Over time, however, storage costs fell dramatically as semiconductor technology improved and computing hardware became more affordable. As shown in figure 1, the cost of computer memory and storage declined substantially over several decades, fundamentally changing the nature of digital systems.[9]

Figure 1: Cost per terabyte of memory and storage, 1956–2023 (log scale)



The spread of personal computers in the 1980s and 1990s marked a major transition from shared computing toward individualized computing environments. Local hard drives allowed systems to save files, user settings, interface layouts, and application preferences directly on personal devices. Consumers no longer needed to fully reconfigure systems every time they used them as they had to do with such early personal computers as the Apple Macintosh (see figure 2). Features such as saved desktop settings, browser bookmarks, stored preferences, and saved login credentials became commonplace.

Figure 2: Apple Macintosh 128K (Credit: Apple, Inc.)



These developments introduced one of the central goals of personalization: reducing friction through continuity. Even relatively simple forms of persistence improved convenience by eliminating repetitive tasks. Saving a preferred printer, retaining font settings, or remembering frequently visited websites reduced the cognitive and time costs associated with using digital systems.

The rise of the commercial Internet accelerated this shift from stateless to persistent experiences. Early websites initially functioned much like earlier computing systems, treating each page request as independent. The web itself lacked built-in memory. This limitation became increasingly problematic as online commerce expanded because websites could not easily remember shopping carts, login sessions, or user preferences across visits.

The introduction of HTTP "cookies" in 1994 became a foundational breakthrough in web personalization.[10] Netscape engineer Lou Montulli developed cookies to allow websites to retain state information between page requests, originally to support persistent shopping carts for e-commerce applications. Cookies enabled websites to recognize returning users, maintain login sessions, remember language preferences, and preserve basic account settings across visits.

This seemingly modest technical innovation fundamentally changed the nature of the web. Websites could now provide continuity rather than treating every interaction as entirely new. "Remember me" functionality became widespread across email services, retailers, financial platforms, and online communities. Persistent sessions reduced repetitive authentication burdens and made online services significantly easier to use.

Importantly, early persistence features were generally user facing and convenience oriented. Consumers actively valued systems that reduced repetitive actions and retained useful information. Browser bookmarks helped users navigate a rapidly expanding web. Saved passwords reduced login friction. Persistent shopping carts allowed consumers to continue purchases across sessions.

Cookies quickly evolved beyond session management into tools for analytics, advertising, and behavioral tracking. By the late 1990s, policymakers and privacy advocates began raising concerns about how persistent identifiers could be used to monitor users across websites. Some privacy-sensitive users adapted by using browser plug-ins, virtual private networks (VPNs), and other services that increased their online anonymity, although most users generally preferred having sites remember useful information and reduce unnecessary repetition.

The evolution of personalization was not solely driven by firms seeking more data. It was also driven by consumer expectations for more seamless, efficient, and individualized digital experiences.

The Customization Era (1990s–2000s)

As computing became more persistent in the 1990s and early 2000s, personalization evolved beyond simple continuity features into explicit user customization. Consumers increasingly configured digital systems to reflect their preferences, identities, and habits. Unlike later forms of algorithmic personalization, this phase was largely user directed. Individuals actively chose how devices, interfaces, and online profiles would look and behave.

The growth of personal computing, mobile phones, and consumer Internet services created new opportunities for individualized digital experiences. Consumers customized desktop themes, wallpapers, screen savers, browser homepages, avatars, ringtones, and application layouts. Personalization became part of mainstream consumer technology because it offered clear and immediate benefits: greater convenience, faster access to preferred content, stronger self-expression, and a greater sense of ownership over digital environments.

This period marked an important shift in how consumers related to technology. Earlier computing systems were largely functional and standardized. By contrast, the customization era allowed digital devices and online spaces to become extensions of personal identity.[11] Consumers increasingly expected technology to adapt to them rather than having to adapt to technology.

One of the clearest examples was the rapid growth of mobile phone personalization. Downloadable ringtones emerged in the late 1990s and quickly became a global phenomenon. Global ringtone revenues reached several billion dollars annually by the mid-2000s.[12] Consumers willingly paid for personalization because it allowed phones to reflect personal tastes, music preferences, and social identity.

Personalization also became central to the early web experience. Yahoo! launched My Yahoo in 1996, allowing users to build customized homepages with selected news feeds, weather, sports scores, finance information, and email modules.[13] AOL later introduced similar customizable portal experiences through MyAOL.[14] These popular services gave users direct control over what information appeared and how it was organized. Instead of navigating generic interfaces, consumers could create personalized dashboards tailored to their routines and interests.

Personalization became part of mainstream consumer technology because it offered clear and immediate benefits: greater convenience, faster access to preferred content, stronger self-expression, and a greater sense of ownership over digital environments.

Desktop operating systems similarly expanded personalization features throughout this period. Microsoft Windows popularized customizable wallpapers, color schemes, icon packs, and desktop themes. Screen savers evolved from simple technical utilities into aesthetic and expressive features. Mobile devices added custom backgrounds, themes, and menu layouts. Online forums, chat services, and early social media platforms introduced customizable avatars, profile pages, and user-generated visual identities.

Notably, these forms of personalization were transparent and voluntary. Users decided what information to provide and actively controlled the customization process. Consumers were often willing to trade time and effort for more personalized experiences. Building custom profiles, selecting themes, organizing homepage widgets, and purchasing personalized mobile content all required active participation. Yet consumers embraced these features because the benefits outweighed the friction involved in configuring them.

The Algorithmic Era (2000s–2010s)

As digital services expanded in scale and complexity during the 2000s and 2010s, the limits of explicit personalization became increasingly apparent. Manual customization required users to continuously manage settings, preferences, and interfaces themselves—a model that became difficult to sustain as online content and digital ecosystems grew substantially larger. In response, personalization increasingly shifted from explicit configuration toward automated inference. Rather than relying primarily on users to specify preferences directly, digital platforms began learning from observed behavior such as searches, purchases, clicks, viewing histories, and engagement patterns. Personalization evolved from a model in which "the user tells the system" to one in which "the system learns the user."

Several technological developments enabled this transition. Cloud computing dramatically expanded storage and processing capacity, allowing firms to analyze enormous volumes of behavioral data at relatively low cost. Mobile devices and broadband Internet generated continuous streams of user activity data, including searches, purchases, viewing histories, clicks, locations, and engagement patterns. Advances in machine learning and recommender systems made it increasingly feasible to identify patterns across millions of users and predict what individuals might want to see, purchase, or consume next.

One of the most influential developments was collaborative filtering.[15] Collaborative filtering uses similarities between users to generate recommendations. Since users with similar viewing or purchasing histories tend to prefer the same products or content, systems recommend items

based on those patterns. Rather than requiring consumers to manually specify every preference, platforms could infer likely interests from observed behavior.

Amazon became one of the first major commercial adopters, introducing item-to-item collaborative filtering in 1998 to recommend products based on purchasing and browsing histories.[16] The system allowed Amazon to generate highly scalable recommendations across millions of products and users. These recommendation systems quickly became central to digital platform business models. Amazon later reported that recommendation systems generated roughly 30 percent of page views on the platform, while other estimates suggested recommendations drove as much as 35 percent of sales.[17] Recommendation engines helped consumers navigate increasingly large online marketplaces by surfacing products likely to match individual interests.

Streaming services adopted similar approaches. Netflix invested heavily in recommendation systems throughout the 2000s, culminating in the launch of the Netflix Prize competition in 2006 to improve predictive recommendation accuracy. By 2016, Netflix had reported that more than 80 percent of content viewed on the platform had been driven by recommendations rather than direct search.[18] Personalized recommendations reduced the time users spent browsing and increased the likelihood that consumers would find content relevant to their interests.

Search and social media platforms also became increasingly personalized during this period. Google introduced personalized search in 2004, tailoring results using search histories and prior user interactions.[19] Search results increasingly reflected individual behavior patterns rather than presenting identical rankings to every user. Facebook similarly transformed social content distribution in 2011 through algorithmically ranked News Feeds that prioritized posts based on predicted relevance and engagement.[20]

The introduction of recommendation algorithms fundamentally changed how consumers interacted with digital information. Earlier personalization models required active configuration. Algorithmic systems instead adapted continuously and automatically.

Music streaming services further accelerated this trend. Spotify's Discover Weekly playlist, launched in 2015, used listening histories and machine learning models to generate personalized music recommendations for individual users.[21] By 2020, Spotify had reported that users had streamed more than 2.3 billion hours of Discover Weekly content.[22] Personalized discovery tools allowed consumers to navigate enormous content libraries more efficiently while exposing them to artists and genres they might not have otherwise encountered.

The introduction of recommendation algorithms fundamentally changed how consumers interacted with digital information. Earlier personalization models required active configuration. Algorithmic systems instead adapted continuously and automatically. Consumers no longer needed to manually organize every preference or explicitly identify every interest. Systems increasingly learned from user signals in the background.

These systems delivered substantial consumer benefits. Personalized recommendations reduced search costs by helping users find relevant products, services, and content more quickly. As digital marketplaces and content libraries expanded, algorithmic personalization helped

consumers manage information overload. Recommendation systems also improved discovery by surfacing niche products, smaller creators, and specialized content that users might otherwise never encounter.

Importantly, algorithmic personalization emerged in response to growing digital abundance. By the late 2000s, consumers had far more information, products, and media choices than earlier generations of digital systems were designed to manage. Recommendation systems helped organize this complexity into more usable and individualized experiences.

The Predictive Era (2010s–Present)

By the 2010s, personalization had become increasingly predictive, contextual, and integrated into real-world environments. Earlier personalization systems primarily responded to explicit preferences or historical behavior. Newer systems began adapting continuously in real time using location data, behavioral history, environmental signals, device activity, and predictive modeling. Personalization increasingly shifted from recommending content to actively assisting users and automating routine decisions.

Several technological developments enabled this transition. Smartphones created continuous streams of contextual data, including location, movement, app usage, payments, communications, and search activity. Cloud computing and mobile broadband allowed systems to process this information instantly across platforms and devices. Advances in machine learning made it possible to predict likely user needs based not only on past behavior but also on immediate situational context.

Importantly, this form of personalization is highly functional rather than primarily expressive. The value comes from reducing friction and automating coordination problems that would otherwise require substantial user effort. Consumers do not manually configure these optimizations. Instead, systems continuously adapt in the background to improve efficiency, convenience, and responsiveness.

Advances in machine learning made it possible to predict likely user needs based not only on past behavior but also on immediate situational context.

Ride-hailing platforms illustrate this shift well. Uber, for example, uses contextual data—such as location, traffic conditions, route information, historical travel patterns, and real-time supply and demand—to coordinate rides, estimate arrival times, optimize matching, and improve reliability.[23] At the same time, Uber has publicly committed to not engage in individualized "personalized pricing" based on personal characteristics or device information. The company has stated that it does not use factors such as race, gender, battery level, phone model, or similar personal attributes to determine prices or promotions—and it provides users with transparency and opt-out mechanisms for personalized offers.[24] These commitments demonstrate how firms can deploy sophisticated personalization systems while still establishing guardrails around fairness, transparency, and consumer trust.

The same contextual logic increasingly spread beyond software interfaces into physical environments. Recognition systems became common across transportation, retail, hospitality, and building access. RFID badges, mobile credentials, and license-plate recognition systems

allowed physical infrastructure to identify and respond to users automatically. Parking garages could recognize registered vehicles and open gates without tickets or payment kiosks. Office buildings could grant seamless entry through contactless credentials. Hotels increasingly stored guest preferences and account histories to personalize check-in, room settings, and loyalty experiences. These systems represented the emergence of what some have called the "recognition economy," wherein environments automatically identify users and adapt services accordingly.[25]

Retail automation further expanded these capabilities. Amazon introduced Amazon Go stores in 2018 using "Just Walk Out" technology that combines cameras, sensors, computer vision, and machine learning to both detect which products shoppers remove from shelves and automatically process payment when customers leave the store.[26] Consumers no longer needed to scan items or wait in checkout lines. Personalization in this context was embedded directly into the transaction process itself.

At the same time, AI systems introduced a new layer of persistent contextual personalization. New AI assistants increasingly maintained conversational continuity, remembered preferences, and assisted with ongoing workflows. Systems could now retain information about user habits, communication styles, recurring tasks, scheduling preferences, and professional workflows across sessions.

Major AI platforms increasingly compete on contextual memory and continuity. Leading platforms have all introduced memory and persistent context capabilities for AI assistants. These systems can reference prior interactions, retain stated preferences, and automate recurring tasks across applications and devices. Consumers increasingly expect these systems to remember information rather than require every interaction to begin from scratch.

This marks an important shift in the history of personalization. Earlier systems primarily filtered information or recommended products. Now systems can increasingly perform delegated assistance functions: scheduling meetings, drafting emails, managing workflows, anticipating travel needs, organizing information, and automating repetitive tasks using persistent context and predictive reasoning.

PERSONALIZATION AS CONSUMER EMPOWERMENT

Critics of personalization often frame autonomy primarily as freedom from influence or data collection. But in modern digital environments, autonomy also depends on the ability to navigate complexity effectively. As information, services, and choices expand, consumers increasingly rely on personalized systems to filter options, reduce friction, and simplify decision-making. Personalization can therefore enhance consumer autonomy not by eliminating choice, but by making choice more manageable and actionable.

This distinction is increasingly important in digital markets characterized by abundance rather than scarcity. Consumers today face overwhelming volumes of information, products, media, communications, and services. Search engines index billions of webpages. Streaming platforms offer catalogs containing tens of thousands of titles. E-commerce marketplaces contain hundreds of millions of products. Navigation systems process continuously changing traffic conditions across entire transportation networks. In these environments, the central challenge is often not gaining access to choices, but rather identifying which choices are most relevant.

Cognitive scientist Herbert Simon anticipated this problem decades ago. He argued that a wealth of information creates a scarcity of attention. Human decision-making is constrained by bounded rationality: individuals possess limited time, attention, and cognitive processing capacity.[27] Consumers cannot realistically evaluate every available option in increasingly complex digital environments. Systems that help filter, prioritize, and organize information can therefore increase practical autonomy by reducing cognitive burden and improving decision quality.

Personalization plays a central role in reducing these search costs. Search engines personalize rankings to surface results more likely to match user intent. E-commerce platforms recommend products based on browsing history, prior purchases, and behavioral similarities across users. Streaming services personalize recommendations to help consumers navigate enormous content libraries. Navigation applications personalize routing based on current location, traffic conditions, historical travel patterns, and user preferences.

As information, services, and choices expand, consumers increasingly rely on personalized systems to filter options, reduce friction, and simplify decision-making.

Importantly, personalization often improves outcomes without reducing meaningful consumer control. Consumers remain free to override recommendations, conduct additional searches, or select alternative options. Personalized systems generally function as decision-support tools rather than mandatory constraints. A navigation app may recommend a faster route, but drivers can still choose another path. A streaming platform may recommend content, but users retain control over what they watch.

Personalization also enhances autonomy by reducing repetitive administrative tasks and routine friction. Many modern digital conveniences depend on systems remembering prior interactions and anticipating user needs. Autofill tools save addresses, payment credentials, and frequently entered information. Smart reply systems suggest likely responses based on conversational context. Saved payment methods streamline transactions. Predictive routing systems proactively adjust navigation based on changing traffic conditions.

These features reduce the cumulative cognitive and time burdens associated with digital interactions. Individually, each task may seem minor. Collectively, however, repetitive configuration and data entry impose significant friction costs. Personalization improves usability by allowing systems to retain continuity across interactions.

The value of these systems becomes particularly clear when they fail. Consumers become frustrated when devices forget passwords, lose preferences, fail to synchronize settings, or require the same information to be repeatedly re-entered across platforms. The expectation that systems should "remember" users has become deeply embedded in modern digital experiences because persistence and continuity improve efficiency.

This reflects a broader concept of practical autonomy. Consumers gain agency not simply from possessing formal choices, but from having tools that help them act effectively within complex environments. Systems that remember preferences, anticipate needs, and simplify coordination can increase users' ability to achieve desired outcomes with less effort and cognitive strain.

THE RECURRING "PERSONALIZATION PANIC" CYCLE

Debates surrounding personalization and data collection are often framed as unprecedented challenges created by new technology. In practice, however, most concerns follow a recurring historical pattern. New forms of personalization routinely generate early fears about surveillance, manipulation, social control, or the erosion of privacy. Over time, consumers adapt to many of these technologies as their practical benefits become clearer, social norms evolve, and governance frameworks mature.

This phenomenon is known as the "privacy panic cycle": a recurring tendency for emerging digital technologies to trigger predictions of widespread privacy harm that exceed what ultimately materializes in practice.[28] This does not mean all privacy concerns are unfounded. But the broader pattern is one in which public discourse frequently assumes worst-case outcomes before the long-term effects of technologies are fully understood.

Consider the debates over loyalty programs and RFID (Radio-Frequency Identification) technologies during the 2000s. Grocery loyalty cards, pharmacy rewards programs, and fuel discount systems all relied on collecting purchase histories and linking transactions to individual consumers. Critics argued that retailers were building intrusive behavioral profiles capable of revealing highly sensitive personal information. Despite these concerns, loyalty programs became one of the most widely adopted forms of retail personalization because consumers perceived clear economic value. Discounts, targeted coupons, rewards points, and convenience benefits outweighed privacy concerns for most users.

RFID technologies generated even more dramatic backlash. Advocacy organizations such as Consumers Against Supermarket Privacy Invasion and Numbering (CASPIAN) warned that RFID tags embedded in products could enable covert tracking and mass surveillance.[29] Activists frequently described RFID tags as "spy chips" and organized protests against retailers experimenting with RFID-enabled inventory systems.[30]

At the time, some critics suggested RFID technologies could evolve into ubiquitous systems for tracking consumers' physical movements and purchases without their knowledge. Yet many of the most expansive predictions did not materialize broadly. Instead, RFID systems became integrated into relatively specific and functional applications such as inventory management, toll collection, transit systems, hotel keycards, supply-chain logistics, and contactless payments.

Importantly, consumer resistance has since faded now that the practical benefits have become clear. Contactless cards, mobile wallets, automated toll systems, and building access technologies reduced transaction friction and improved convenience. The point is not that new technologies never create privacy risks, but rather that the historical record suggests that public discourse often veers toward maximalist assumptions about technological harms before longer-term patterns of adaptation, governance, and consumer behavior become clearer.

In practice, consumers frequently adapt to new forms of personalization by developing new norms, expectations, and behavioral strategies. Users learn which services they trust, what tradeoffs they are willing to make, and how to manage privacy settings and platform choices over time. Markets also evolve in response to consumer preferences, regulatory pressure, and competitive dynamics.

This history is particularly relevant to current debates surrounding personalization and AI. Contemporary concerns about persistent AI assistants, contextual memory, and integrated personalization echo many earlier anxieties surrounding browser cookies, recommendation systems, and RFID. History suggests that policymakers should distinguish between plausible, evidence-based risks and speculative worst-case narratives.

This distinction is especially important in debates surrounding AI and privacy, wherein some critics argue that all forms of seemingly innocuous data should be treated as highly sensitive because any dataset could theoretically be combined with other information to reconstruct personal identities or reveal intimate details—a phenomenon referred to as the "mosaic effect."[31]

Consumers have repeatedly demonstrated willingness to adopt personalized technologies when benefits are clear, controls are meaningful, and systems operate within socially acceptable norms, which may change over time.

While reidentification risks are real in certain contexts, the strongest versions of these claims often overstate both technical feasibility and practical likelihood. Research on deidentification demonstrates that reidentification risks vary substantially depending on the type of data, available auxiliary datasets, access conditions, and institutional safeguards involved. Not all data creates equivalent risks, and not every dataset can realistically be transformed into comprehensive personal surveillance.[32] However, treating all data as equally sensitive can create counterproductive policy outcomes by discouraging beneficial uses of data, such as personalization, analytics, and fraud prevention that create substantial consumer value while posing relatively limited privacy risks when appropriately governed.

While public opinion surveys show that consumers express abstract concerns about data collection, actual market behavior reveals sustained adoption of personalized tools.[33] This tension is known as the "privacy paradox": individuals routinely choose—and even pay for subscriptions to—personalized navigation apps, streaming services, e-commerce platforms, and AI assistants because the practical benefits of reduced friction and improved relevance outweigh abstract privacy anxieties.

The broader historical lesson is not that privacy concerns should be ignored. Rather, it is that personalization debates benefit from proportionality, empirical grounding, and attention to actual consumer outcomes rather than speculative fears. Consumers have repeatedly demonstrated willingness to adopt personalized technologies when benefits are clear, controls are meaningful, and systems operate within socially acceptable norms, which may change over time.

GLOBAL APPROACHES TO GOVERNING PERSONALIZATION

Governments around the world have taken sharply different approaches to regulating data, tech companies, and digital services. These policy choices have increasingly shaped not only privacy practices but also the quality, functionality, and competitiveness of digital services themselves. The global experience suggests that overly rigid regulatory frameworks can reduce service quality, limit innovation, and make it more difficult for firms to deploy useful personalized technologies.

The European Union has created the world's most strict regulatory approach to data governance through frameworks such as the General Data Protection Regulation (GDPR), the Digital Services Act (DSA), and the AI Act. These laws were designed to strengthen consumer protections, increase transparency, and constrain perceived abuses by large digital platforms. However, they have also imposed substantial compliance burdens on firms deploying personalized digital services.

GDPR, in particular, has shaped how companies collect, process, and retain user data.[34] Strict consent requirements, purpose limitations, data minimization obligations, and broad compliance mandates have increased the legal and operational complexity of personalization systems. While large incumbent firms have generally been able to absorb these costs, smaller firms and start-ups often face disproportionate burdens due to limited legal, technical, and compliance resources.

A growing body of research suggests that GDPR produced significant economic trade-offs.[35] Studies have found reductions in venture capital investment among European technology start-ups following GDPR implementation, particularly among smaller data-driven firms.[36] Other analyses found increased market concentration advantages for large platforms with extensive first-party user relationships, as these firms were better positioned to navigate restrictive data governance requirements than were smaller competitors dependent on third-party advertising or analytics ecosystems.[37]

The global experience suggests that overly rigid regulatory frameworks can reduce service quality, limit innovation, and make it more difficult for firms to deploy useful personalized technologies.

The digital advertising market illustrates these dynamics clearly. Restrictions on data sharing and targeting capabilities reduced the effectiveness of many smaller advertising and publishing operations while reinforcing the position of dominant platforms with direct consumer relationships and large authenticated user bases. Compliance costs also diverted resources away from product development and innovation toward legal and regulatory administration.

These effects increasingly extend beyond online advertising into AI systems and other personalized digital services. Many emerging AI assistants and workflow tools depend on persistent memory, contextual awareness, cross-service integration, and adaptive learning to function effectively. Regulatory frameworks that excessively restrict data retention, contextual inference, or cross-platform interoperability may reduce the usefulness of these systems by preventing them from maintaining continuity or adapting effectively to user needs.

The European experience with AI deployment already reflects some of these tensions. Companies have delayed or limited the rollout of certain AI features in Europe due to legal uncertainty surrounding data processing, model training, liability exposure, and compliance obligations.[38] Some personalized AI services available in the United States or Asia have launched later or in more limited forms within the European market because of regulatory concerns.[39]

By contrast, several Asian digital ecosystems have generally adopted more permissive approaches toward integrated personalization and digital infrastructure deployment. Countries such as Singapore, South Korea, and China have supported broad deployment of integrated payments, digital identity systems, smart-city infrastructure, and super-app ecosystems that rely heavily on

contextual personalization and cross-service coordination. These systems often enable highly seamless consumer experiences because services can share contextual information across transportation, payments, messaging, commerce, and authentication systems. Consumers can move through transportation networks, retail environments, and digital platforms with minimal repetitive configuration or transactional friction.

The point is not that all these regulatory systems are preferable or free from legitimate concerns. Some systems raise substantial issues surrounding state surveillance, competition policy, or civil liberties. But they also demonstrate that more permissive frameworks can accelerate deployment of highly integrated and functional personalization systems.

The United States has historically adopted a more sectoral and comparatively flexible approach. Rather than relying on a single omnibus privacy framework equivalent to GDPR, the United States regulates specific categories of sensitive information through sector-specific laws and consumer protection enforcement. This approach has generally allowed broader experimentation in personalized services, recommendation systems, digital advertising, and AI applications.

As a result, many of the world's leading personalized technology services have emerged from the United States, including modern search engines, streaming recommendation systems, ride-hailing platforms, cloud productivity tools, and AI assistants. A more flexible regulatory environment has helped firms experiment with personalization models, refine user experiences, and scale new digital services rapidly.

The divergence between regulatory models increasingly matters because personalization is becoming essential in the new wave of consumer-facing AI tools. AI assistants, for example, depend on their ability to retain information, adapt to users, and coordinate across services. Excessively restrictive regulatory frameworks would limit not only advertising models but also the development of useful AI capabilities and digital systems more broadly.

Large incumbents are often better positioned to absorb regulatory complexity than are start-ups or smaller competitors. Heavy-handed regulation can therefore unintentionally entrench dominant firms by raising barriers to entry for new personalized services that lack the scale or legal capacity to manage extensive compliance requirements.

The global experience suggests that effective governance requires balancing legitimate consumer protections with innovation flexibility. Transparency, accountability, and user control remain important. But frameworks that treat personalization itself as inherently suspect risk reducing service quality, slowing AI deployment, limiting consumer convenience, and weakening competitive dynamism in digital markets.

THE NEXT PHASE: AI AGENTS

The emergence of agentic AI systems, which can independently pursue multistep goals, reason through complex tasks, and execute actions on behalf of users across different applications, represents the next stage in the evolution of personalization. The underlying objective remains consistent with prior generations of personalization technologies: reducing friction, improving relevance, simplifying coordination, and helping users navigate increasingly complex digital environments. What is changing is the depth, continuity, and functional scope of personalization.

Earlier personalization systems primarily recommended products, ranked content, or optimized isolated interactions. AI systems increasingly maintain long-term context across tasks, applications, and workflows. These systems can remember user preferences, adapt to communication styles, retain project context, coordinate information across services, and automate recurring activities over time.

Major AI platforms are rapidly moving in this direction, having introduced persistent memory capabilities that allow AI systems to reference prior conversations, retain stated preferences, and provide more context-aware assistance. These systems increasingly operate less like isolated conversations with experts and more like familiar interactions with a close partner.

Technically, these systems combine several evolving capabilities. Long-context models can process larger amounts of information simultaneously, enabling AI systems to reference broader conversational histories and supporting documents. Persistent memory systems allow models to retain structured information about user preferences and recurring tasks across sessions. Retrieval systems enable assistants to access relevant information dynamically from connected applications, documents, calendars, and enterprise systems. Preference modeling allows systems to adapt outputs based on observed user behavior and stated priorities over time. Together, these capabilities move personalization from recommendation toward assistance and workflow coordination. AI systems increasingly help users perform tasks rather than merely suggesting content or products.

This distinction is important because much of the debate surrounding AI and privacy still reflects older assumptions associated with advertising and behavioral targeting. Historically, personalization was often discussed primarily in the context of marketing optimization: recommending products, ranking ads, or maximizing engagement. Increasingly, however, AI personalization is becoming assistive rather than merely promotional.

Consumers use AI systems to summarize meetings, draft documents, organize schedules, automate repetitive workflows, manage communications, conduct research, and coordinate information across applications. The value proposition is less about influencing consumer purchasing behavior and more about reducing cognitive load and administrative friction. These systems become more useful over time precisely because they retain contextual understanding and continuity.

Historically, personalization was often discussed primarily in the context of marketing optimization: recommending products, ranking ads, or maximizing engagement. Increasingly, however, AI personalization is becoming assistive rather than merely promotional.

Persistent context is especially valuable in complex workflow environments in which repetitive reconfiguration creates substantial inefficiencies. Workers do not want to repeatedly explain organizational structures, formatting preferences, project histories, or procedural requirements during every interaction with an AI system. Similarly, consumers often prefer systems that remember travel preferences, dietary restrictions, scheduling habits, accessibility needs, or recurring purchasing patterns without requiring constant manual reentry.

This does not mean persistent AI systems should operate without constraints. Long-term memory and cross-service integration create legitimate governance challenges surrounding transparency, data security, retention limits, user control, and institutional trust. But these governance challenges should be understood within the broader historical trajectory of personalization. The evolution from stateless systems to persistent computing, behavioral recommendations, contextual services, and AI assistants reflects a continuous effort to reduce friction and improve continuity in increasingly complex digital environments.

GOVERNING HARM WITHOUT STOPPING PROGRESS

As personalized digital systems become more sophisticated and integrated into everyday life, policymakers face an increasingly important challenge: how to address legitimate risks without undermining technologies that deliver substantial consumer and economic benefits. The appropriate policy objective is not restricting the development and deployment of AI, restricting personalized services, or preventing systems from learning, remembering, or adapting personal information. Modern digital services increasingly depend on these capabilities to function effectively. Instead, governance frameworks should focus on preventing harmful outcomes while preserving innovation, usability, and consumer value.

This distinction is particularly important as AI systems evolve from passive recommendation tools into active assistants capable of maintaining memory, coordinating workflows, and automating tasks. Policymakers should distinguish between personalization itself and specific harmful uses of systems with personalization capabilities.

Certain risks clearly warrant regulatory attention. Bad actors could use personalized systems to facilitate fraud, engage in deceptive practices, discriminate against users, or manipulate individuals. Systems that use personalization to exploit vulnerable users, obscure material information, or create unfair market outcomes can create legitimate consumer harm. Similarly, highly personalized systems may generate significant risks if they mishandle sensitive data, produce unsafe outputs, or fail in ways that materially harm users.

Governance frameworks should focus on preventing harmful outcomes while preserving innovation, usability, and consumer value.

Transparency therefore remains an important principle. Consumers should generally be able to learn about how companies personalize their services, what categories of information they retain or infer, and how this personalization affects significant decisions or interactions. Transparency allows consumers to make informed choices about how they engage with these services.

Accountability also becomes increasingly important as AI systems gain greater autonomy and operational responsibility. Traditional recommendation systems primarily suggested information or products. Emerging AI assistants may schedule appointments, manage communications, automate transactions, or coordinate workflows across services. As AI agents take more consequential actions on behalf of users and obtain greater access to users' digital lives, policymakers will need clearer frameworks for assigning responsibility when mistakes occur.

Liability questions will become particularly important in this context. Much like a fallible human assistant, personalized AI systems may occasionally produce harmful recommendations, execute

incorrect instructions, expose sensitive information, or make operational mistakes that create financial or reputational harms. Governance frameworks should clarify how accountability is allocated among developers, deployers, enterprise users, and consumers depending on the nature of the failure, the degree of system autonomy involved, and the ability of that stakeholder to reasonably prevent harm.

Importantly, these liability discussions should focus on harmful outcomes and operational responsibility rather than treating memory or learning capabilities themselves as inherently problematic. Systems that retain contextual information are not uniquely dangerous simply because they remember user preferences or prior interactions. In many cases, memory and continuity improve reliability, reduce user error, and make systems more useful.

Indeed, using information to fulfill user requests more effectively is often the central purpose of personalization. A navigation system that remembers frequent destinations, a fraud prevention system that recognizes unusual account activity, or an AI assistant that recalls workflow preferences is functioning as intended. Restricting systems from learning from prior interactions too aggressively could substantially reduce functionality and usability.

This distinction is especially vital when evaluating policy proposals that mandate constant permission prompts, repeated reauthorizations, or highly fragmented consent structures for routine personalization. While user control remains essential, excessive permission friction paradoxically reduces practical consumer autonomy by introducing fragmented workflows that make digital tools cumbersome, repetitive, and frustrating to navigate. Consumers prioritize continuity and persistence precisely because they eliminate these administrative burdens. Indeed, in some cases, policymakers may want to modify existing requirements to prevent forcing users to respond to excessive consent requests.

Similarly, policymakers should avoid treating all forms of inference as inherently problematic. Many useful digital services depend on systems being able to infer likely preferences, detect anomalies, prioritize information, or anticipate user needs. Fraud-detection systems infer suspicious behavior patterns. Recommendation systems infer likely consumer interests. Accessibility tools infer user needs based on interaction patterns. AI assistants infer workflow preferences to improve task coordination. Inference itself is therefore not the problem. The relevant policy question is whether systems use inferred information in ways that create material harms, unfair discrimination, deception, coercion, or unacceptable privacy risks.

The central policy principle is straightforward: regulate harmful outcomes, not memory, learning, or personalization itself.

This distinction also applies to debates surrounding the so-called “mosaic effect,” which holds that combining seemingly innocuous data points can potentially reveal sensitive information. While this concern is valid in some contexts, proponents of stricter regulation have concluded that virtually all data collection and personalization should be treated as inherently dangerous because any data could theoretically contribute to future inference risks.

Such an approach collapses meaningful distinctions between different categories of data, uses, and risk levels. Not all personalization creates equivalent harms, and not all data processing

creates meaningful threats to privacy or autonomy. A proportional governance framework should therefore evaluate personalization practices based on actual risks, contexts, and outcomes rather than assume that all learning, memory, or inference capabilities are inherently suspect.

Several policy principles can help guide this approach.

First, transparency should be a priority. Users should be able to find when personalization (or a lack of personalization) is occurring and possess meaningful visibility into important data practices and automated decision-making processes.

Second, consumers should retain meaningful control. Consumers are not all the same. Users should be able to modify preferences, manage memory features, delete stored information, and opt out of certain forms of personalization wherever feasible and appropriate.

Third, interoperability and data portability should be encouraged wherever appropriate. Consumers should not become unnecessarily locked into particular ecosystems because personalization histories or contextual profiles cannot transfer across services.

Fourth, accountability mechanisms should focus on harmful outcomes and operational responsibility. Systems that create demonstrable harms should face appropriate legal and regulatory scrutiny regardless of whether those harms arise from digital systems.

Fifth, policymakers should preserve innovation flexibility. Personalization technologies evolve rapidly, and overly rigid ex ante restrictions may unintentionally reduce competition, slow beneficial innovation, and entrench incumbent firms better able to absorb compliance burdens.

Finally, governance frameworks should maintain proportionality. Regulation should distinguish between low-risk personalization that improves usability and higher-risk practices involving sensitive decisions, manipulation, discrimination, or material consumer harms.

The central policy principle is straightforward: regulate harmful outcomes, not memory, learning, or personalization itself. The ability of systems to adapt, retain context, and improve through interaction is increasingly fundamental to modern digital functionality. Policymakers should therefore focus on ensuring that these capabilities operate within trustworthy and accountable frameworks rather than attempting to prevent digital systems from becoming more context aware and useful.

About the Author

Daniel Castro is president of ITIF and director of ITIF's Center for Data Innovation. He has a B.S. in foreign service from Georgetown University and an M.S. in information security technology and management from Carnegie Mellon University.

About ITIF

The Information Technology and Innovation Foundation (ITIF) is an independent 501(c)(3) nonprofit, nonpartisan research and educational institute that has been recognized repeatedly as the world's leading think tank for science and technology policy. Its mission is to formulate, evaluate, and promote policy solutions that accelerate innovation and boost productivity to spur growth, opportunity, and progress. For more information, visit itif.org/about.

ENDNOTES

1. Alex Ambrose, "The Promise of Wearable AI: Opportunities Across Emergency Response" (ITIF, April 15, 2026), https://itif.org/publications/2026/04/15/the-promise-of-wearable-ai-opportunities-across-emergency-response/.
2. "ChatGPT can now reference all past conversations," OpenAI Developer Community, April 10, 2025, https://community.openai.com/t/chatgpt-can-now-reference-all-past-conversations-april-10-2025/1229453.
3. Josh Woodward, "Gemini introduces personal intelligence," *Google Blog*, January 4, 2026, https://blog.google/innovation-and-ai/products/gemini-app/personal-intelligence/.
4. "Introducing Copilot Memory: A More Productive and Personalized AI for the Way You Work," *Microsoft 365 Copilot Blog*, July 14, 2025, https://techcommunity.microsoft.com/blog/microsoft365copilotblog/introducing-copilot-memory-a-more-productive-and-personalized-ai-for-the-way-you/4432059.
5. Meredith Whittaker and Udbhav Tiwari, "AI Agent, AI Spy," 39th Chaos Communication Congress, December 2025, https://media.ccc.de/v/39c3-ai-agent-ai-spy.
6. Daniel J. Solove, "Artificial Intelligence and Privacy," *Florida Law Review*, 77, no. 1 (2025), https://www.floridalawreview.com/article/129976-artificial-intelligence-and-privacy.
7. Miranda Bogen and Ruchika Joshi, "What AI 'remembers' about you is privacy's next frontier," *MIT Technology Review*, January 28, 2026, https://www.technologyreview.com/2026/01/28/1131835/what-ai-remembers-about-you-is-privacys-next-frontier/.
8. "CX Trends 2026: Leading in the AI Era," Zendesk, 2026, https://cxtrends.zendesk.com/.
9. John C. McCallum, "Memory Prices 1957 – 2024," (with minor processing by Our World in Data), https://ourworldindata.org/grapher/historical-cost-of-computer-memory-and-storage.
10. Meg Leta Jones, "Cookies: A Legacy of Controversy." *Internet Histories* 4, no. 1 (2020): 87–104. https://doi.org/10.1080/24701475.2020.1725852.
11. Soyeon Kwon, Sejin Ha, and Chris Kowal, "How online self-customization creates identification: Antecedents and consequences of consumer-customized product identification and the role of product involvement," *Computers in Human Behavior*, October 2017, https://www.sciencedirect.com/science/article/abs/pii/S0747563217302947.

12. "Digital Broadband Content: Mobile Content," OECD Digital Economy Papers, No. 97, OECD Publishing, Paris (2005), https://www.oecd.org/content/dam/oecd/en/publications/reports/2005/05/digital-broadband-content_g17a1be2/232201337160.pdf.
13. "Web Users Create Their Own Yahoo!" Yahoo! Press Room, July 15, 1996, https://web.archive.org/web/20080331063253/http://yhoo.client.shareholder.com/press/ReleaseDetail.cfm?ReleaseID=173417.
14. James Niccolai, "AOL revamps Web site," *ComputerWorld*, January 13, 2000, https://www.computerworld.com/article/1367806/aol-revamps-web-site.html.
15. J. Ben Schafer et al., "Collaborative Filtering Recommender Systems," In *The Adaptive Web*, ed. Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl, Lecture Notes in Computer Science, 4321 (Berlin and Heidelberg: Springer, 2007), https://doi.org/10.1007/978-3-540-72079-9_9.
16. Xiaoyuan Su and Taghi M. Khoshgoftaar, "A Survey of Collaborative Filtering Techniques," Advances in Artificial Intelligence, October 27, 2009, https://doi.org/10.1155/2009/421425.
17. Brent Smith and Greg Linden, "Two Decades of Recommender Systems at Amazon.com," *IEEE Internet Computing*, May – June 2017, https://doi.org/10.1109/MIC.2017.72; Ian MacKenzie, Chris Meyer, and Steve Noble, "How Retailers Can Keep Up With Consumers," McKinsey & Company, October 1, 2013, https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers/.
18. Brent Smith and Greg Linden, "Two Decades of Recommender Systems at Amazon.com," IEEE Internet Computing, May – June 2017, https://doi.org/10.1109/MIC.2017.72.
19. Matt Hines, "Google takes searching personally," *CNET*, March 29, 2004, https://www.cnet.com/tech/services-and-software/google-takes-searching-personally/.
20. Mark Tonkelowitz, "Interesting News, Any Time You Visit," *The Facebook Blog*, September 20, 2011, https://web.archive.org/web/20110925211838/http://blog.facebook.com/blog.php?post=10150286921207131.
21. "Spotify Users Have Spent Over 2.3 Billion Hours Streaming Discover Weekly Playlists Since 2015," *For The Record*, July 19, 2020, https://newsroom.spotify.com/2020-07-09/spotify-users-have-spent-over-2-3-billion-hours-streaming-discover-weekly-playlists-since-2015/.
22. Ibid.
23. "Algorithmic Transparency Report 2026, United States," Uber, 2026, https://tb-static.uber.com/prod/uber-static/uber-sites/_pdf/ai-on-uber/US-Algorithmic-Transparency-2026.pdf.
24. Ibid.
25. "Building the Recognition Economy," *Metropolis*, December 4, 2025, https://www.metropolis.io/blog/building-the-recognition-economy.
26. Tian Lan and Chris Broaddus, "Enhancing Just Walk Out technology with multi-modal AI," *AWS Blogs*, September 23, 2024, https://aws.amazon.com/blogs/machine-learning/enhancing-just-walk-out-technology-with-multi-modal-ai/.
27. Herbert A. Simon, "Bounded Rationality," In *Utility and Probability*, ed. John Eatwell, Murray Milgate, and Peter Newman (The New Palgrave, Palgrave Macmillan, London. 1990), https://doi.org/10.1007/978-1-349-20568-4_5.
28. Daniel Castro and Alan McQuinn, "The Privacy Panic Cycle: A Guide to Public Fears About New Technologies" (ITIF, 2015), https://itif.org/publications/2015/09/10/privacy-panic-cycle-guide-public-fears-about-new-technologies/.
29. Michael Beckel, "Total Surveillance," *Mother Jones*, December 6, 2005, https://www.motherjones.com/politics/2005/12/total-surveillance/.

30. Simson Garfinkel, Ari Juels, and Ravi Pappu, "RFID Privacy: An Overview of Problems and Proposed Solutions," IEEE Security & Privacy, 3, no. 3, June 30, 2005, https://doi.org/10.1109/MSP.2005.78.
31. Daniel J. Solove, "Data Is What Data Does: Regulating Use, Harm, and Risk Instead of Sensitive Data," *SSRN Electronic Journal*, 2023, https://doi.org/10.2139/ssrn.4322198.
32. Ann Cavoukian and Daniel Castro, "Setting the Record Straight: De-Identification Does Work" (ITIF, June 16, 2014), https://itif.org/publications/2014/06/16/setting-record-straight-de-identification-does-work/.
33. "How Americans View Data Privacy," Pew Research Center, October 18, 2023, https://www.pewresearch.org/internet/2023/10/18/how-americans-view-data-privacy/.
34. Regulation (EU) 2016/679 of the European Parliament and of the Council of April 27, 2016, on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), Official Journal of the European Union L 119 (May 4, 2016): 1–88, https://eur-lex.europa.eu/eli/reg/2016/679/oj.
35. Chinchih Chen, Carl Benedikt Frey, and Giorgio Presidente, "Privacy Regulation and Firm Performance: Estimating the GDPR Effect Globally," The Oxford Martin Working Paper Series on Technological and Economic Change, January 6, 2022, https://oms-www.files.svdcdn.com/production/downloads/Privacy-Regulation-and-Firm-Performance-Giorgio-WP-Upload-2022-1.pdf.
36. Eline Chivot and Daniel Castro, "What the Evidence Shows About the Impact of the GDPR After One Year" (Center for Data Innovation), June 17, 2019, https://datainnovation.org/2019/06/what-the-evidence-shows-about-the-impact-of-the-gdpr-after-one-year/.
37. Guy Aridor et al., "Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency," Law & Economics Center at George Mason University Scalia Law School, April 28, 2025, http://dx.doi.org/10.2139/ssrn.4698374.
38. Foo Yun Chee, "Meta pauses AI models launch in Europe due to Irish request," *Reuters*, June 14, 2024, https://www.reuters.com/technology/artificial-intelligence/meta-will-not-launch-meta-ai-europe-now-2024-06-14/.
39. Rohan Goswami, "Apple Intelligence won't launch in EU in 2024 due to antitrust regulation, company says," *CNBC*, June 21, 2024, https://www.cnbc.com/2024/06/21/apple-ai-europe-dma-macos.html.