Carnegie Mellon University
HeinzCollege
INFORMATION SYSTEMS • PUBLIC POLICY • MANAGEMENT

Beyond the Wrist: A Comprehensive Analysis of Consumer Privacy and Data Protection in Smartwatches

Prepared by: Mburu Kagiri, Katelyn Wilson, Riley Wallace, Eric Kumara, and Aarushi Jaitly

Prepared for; The Government Accountability Office (GAO)

May 4th, 2026

Table of Contents

Executive Summary

Consumer smartwatches have become widely adopted wearable computing devices capable of continuously collecting and processing biometric, behavioral, and physiological data. Estimates suggest global adoption ranges between 454-640 million users as of 2025, a figure that has grown over 30 percent since 2023 (Ruby, 2022). Over 30 percent of Americans own a smartwatch and over 92 percent of smartwatch users rely on their devices for health and fitness tracking. This rapid adoption, combined with advances in data collection and inference capabilities, has introduced significant privacy and data protection risks, particularly given the sensitive nature of health-related information.

This report finds that existing U.S. legal and regulatory frameworks do not comprehensively address the full lifecycle of smartwatch-generated data. While certain protections exist under federal and state laws, gaps persist in areas including downstream data sharing, inference-based data use, and consumer control over how data is collected, processed, and retained. As a result, consumers may be exposed to risks that are not readily visible or meaningfully controlled under current notice-and-consent models.

The first component of this report, *The Health Data Ecosystem: A Holistic Analysis*, evaluates the current policy landscape across federal, state, and international frameworks, including voluntary standards such as the National Institute of Standards and Technology (NIST) Privacy and AI Risk Management Frameworks. It also examines the technical structure of smartwatch data collection, identifying key points of vulnerability across the data lifecycle, including the role of data brokers and third-party actors in enabling downstream uses that may contribute to consumer harm, such as profiling and inference-based decision making.

The second component, *Policy Proposals and Suggestions for Future Work*, presents four policy approaches designed to address identified gaps. These proposals are organized around a framework of four core objectives: (1) improving consumer comprehension of data practices, (2) strengthening actionable control over data use, (3) mitigating risks associated with secondary and downstream uses of data, and (4) reducing residual exposure within the broader data ecosystem. The proposals include the establishment of a private right of action, expanded transparency and accountability mechanisms, a standardized “nutrition label” disclosure model, and an assessment of pay-for-privacy models and their limitations in protecting consumer data.

Project Objective

This project examines the extent to which existing federal laws and regulatory frameworks protect consumer health and biometric data generated by smartwatches. Specifically, it identifies the roles and limitations of consumer protection statutes, health privacy laws, and emerging state-level privacy regimes in governing the collection, use, and sharing of smartwatch data. In addition, the project assesses how data practices within the smartwatch ecosystem, including downstream data sharing, aggregation, and inference, affect consumer privacy outcomes in areas where legal protections may be limited or unclear.

Based on this analysis, the project identifies gaps in oversight and enforcement and evaluates potential policy approaches to address these gaps. The objective is to provide the Government Accountability Office (GAO) with fact-based findings and policy considerations that can inform Congress and support the development of strategies to strengthen consumer privacy and data protection in the evolving wearable technology ecosystem

Project Methodology

The project team employed a multi-method research approach to assess privacy risks and regulatory gaps in the smartwatch data ecosystem. This approach included:

- Literature review of academic research, policy analyses, and industry reports to identify documented privacy risks, consumer behavior patterns, and emerging technological practices.
- Legal and regulatory analysis of federal, state, and international frameworks to evaluate the extent to which existing laws govern the collection, use, and sharing of smartwatch-generated data.
- Technical analysis of smartwatch data architectures and artificial intelligence (AI) inference systems to map how data is generated, processed, and transmitted across the data lifecycle.
- Semi-structured interviews with subject matter experts in academia and industry, including Carnegie Mellon's Professor Lorrie Cranor and representatives from the

Digital Medicine Society and the Electronic Frontier Foundation, to validate findings and provide contextual insight into real-world practices and policy considerations

Findings from these methods were synthesized using a data lifecycle framework to identify key points of data vulnerability. Evidence was triangulated across legal, technical, and empirical sources to develop a privacy harm matrix and to inform the evaluation of policy options.

This methodology is subject to certain limitations. The analysis relies on publicly available documentation and expert input, which may not fully capture proprietary industry practices or rapidly evolving technological developments. Additionally, while interviews provided valuable qualitative insights, they are not representative of all stakeholders in the smartwatch ecosystem.

Business Impact

This project provides the Government Accountability Office (GAO) with a structured assessment of privacy risks and regulatory gaps associated with smartwatch-generated health data, an area of increasing consumer adoption and limited federal oversight. By synthesizing legal, technical, and behavioral evidence, the report equips GAO with a comprehensive understanding of how existing frameworks apply to emerging wearable technologies and where those frameworks may be insufficient.

The findings are intended to support GAO's core mission of informing Congress and the public. Specifically, the analysis identifies areas where current laws and regulatory approaches may not fully address risks related to data collection, downstream sharing, and algorithmic inference. This positions GAO to more effectively evaluate federal efforts in this space and to communicate identified gaps in a manner that supports legislative consideration. In addition, the project's policy proposals provide GAO with a set of actionable considerations that can be used to assess potential approaches for strengthening consumer data protections. These proposals are designed to align with observed limitations in existing frameworks and may inform GAO recommendations regarding transparency, accountability, and consumer control mechanisms.

Operationally, the report establishes an analytical framework grounded in the data lifecycle and privacy outcome metrics that GAO can use in future work to evaluate similar technologies or data-intensive systems. This framework supports more consistent identification of risk points and facilitates comparison across sectors where health-related data is collected outside traditional healthcare settings. Overall, the project enhances GAO's ability to assess an evolving area of consumer technology, identify gaps in federal oversight, and provide Congress with evidence-based policy considerations to address emerging privacy risks.

The Health Data Ecosystem: A Holistic Analysis

This section provides an overview of the current health data ecosystem as it relates to consumer smartwatches. It describes applicable legal and policy frameworks at the federal, state, and international levels; identifies key characteristics of industry practices; and examines gaps in existing protections. It also reviews voluntary frameworks and summarizes observed challenges affecting consumers and regulated entities.

Existing Policy Frameworks

There are several legal and policy frameworks that operate at the state, federal, and international level that address aspects of consumer data privacy. These frameworks were not developed solely for consumer wearable devices and thus vary in scope, enforceability, and jurisdiction. As such, despite establishing meaningful baseline protections and enforcement authorities, they do not comprehensively govern the collection, processing, and use of smartwatch-generated data. This section is structured by looking at domestic enforceable frameworks (federal and state), domestic non-enforceable frameworks (federal and state), and the European Union's (E.U.) General Data Protection Regulation (GDPR).

Federal Level

At the federal level, several statutes and regulations address consumer privacy and data protection, including Section 5 of the Federal Trade Commission (FTC) Act, the Health Breach Notification Rule (HBNR), the Children's Online Privacy Protection Act (COPPA), and HIPAA. Through the review of these frameworks and laws, significant legal gaps are identified related to both the data subject and the service provider.

Section 5 of the FTC Act

Section 5 of the FTC Act (15 U.S.C. § 45) prohibits "unfair or deceptive acts or practices in or affecting commerce…that cause or are likely to cause reasonably foreseeable injury" (Federal Reserve, n.d.). The statute applies broadly and provides the agency with authority to take enforcement action against harmful business practices. Under this framework, a practice may be considered unfair if it causes or is likely to cause substantial injury to consumers that is not reasonably avoidable and is not outweighed by countervailing benefits. A practice may be considered deceptive if it involves a material representation, omission, or practice that is likely to mislead consumers acting reasonably under the circumstances.

Section 5 can best be understood as a general consumer protection authority rather than a comprehensive data privacy statute. It fails to establish prescriptive requirements governing data collection, processing, retention, or sharing. It also fails to offer guidance on expectations of data minimization, affirmative consent standards, breach notification requirements, risk assessments, or privacy-by-design measures. As a result, certain data practices including large-scale data collection, secondary use, and behavioral profiling, may fall outside the scope of enforcement if they do not meet the statutory thresholds for unfairness or deception. This possibility is further inflated by the inferential references about a consumer that manufactures and their partners might make, including in domains such as AI profiling or training. Lastly, reliance on disclosures as a basis for evaluating deception may limit the applicability of the statute in contexts where consumers do not fully understand complex data practices.

Health Insurance Portability and Accountability Act (HIPAA)

HIPAA establishes federal standards for the protection of sensitive patient health information. It comprises the Privacy Rule, which governs how Protected Health Information (PHI) is used and disclosed, and the Security Rule, which mandates technical, administrative, and physical safeguards for electronic PHI (ASPE, 1996). It also includes the HIPAA Breach Notification Rule (different from HBNR), requiring entities to notify individuals and the Department of Health and Human Services (HHS) of data exposures. HIPAA also grants individuals the right to access their medical records and request corrections.

However, HIPPA's applicability is limited by its definition of "covered entities," which includes doctors, hospitals, insurance companies, and their "business associates." Consumer wearable device manufactures and related application providers are generally not considered covered entities when they collect data directly from users outside of clinical settings. Additionally, HIPAA's data scope is constrained to static medical records rather than continuous biometric streams, further straining its applicability to wearables

HIPPA functions reactively as well, with the Office for Civil Rights (OCR) typically investigating instances of. Altogether, these limitations leave significant gaps in consumer protections by allowing large volumes of consumer-generated health and biometric data to fall outside HIPAA's regulatory scope.

There have been attempts to enhance health data privacy protections by extending regulations similar to HIPAA to consumer health applications and devices. One promising attempt is the Health Information Privacy Reform Act (HIPRA), a legislative proposal intended to extend HIPAA-like protections to non-medical entities such as wearable manufacturers (Congress, 2025). It expands the definition of health information, mandates affirmative consent before data can be sold, and grants the Department of Health and Human Services (HHS) the ability to oversee the restrictions of "commercial use of identifiable personal health information derived from a smartwatch or similar consumer device" (Congress, 2021). Such restrictions would prohibit companies from giving third-parties access to customers' PHI if the purpose is to generate profits. However, the Act does not apply if the consumer provides informed consent or if the PHI is necessary for specified business transactions. Finally, HIPRA introduces a federal Right to Deletion and Data Portability for consumer health data. Although these protections have the potential to protect consumers, HIPRA is yet to be passed into law as of April 2026.

Health Breach Notification Rule (HBNR)

The HBNR requires vendors of personal health records (PHRs), PHR-related entities, and certain third-party service providers to notify affected individuals, the FTC, and, in some cases, the media following a breach involving unsecured identifiable health information (Code of Federal Regulations, 2024). The rule is particularly relevant to consumer wearable technologies because such devices collect health data outside the traditional healthcare settings and thus fall outside

the scope of HIPAA. In these cases, HNBR may apply to broaden coverage where HIPPA falls short.

However, like Section 5 of the FTC Act, HNBR's impact *vis-a-vis* smartwatches is dwarfed by its reactive nature and its limitation to breach notifications. It fails to establish requirements for preventing breaches, does not regulate broader data practices, and offers no means for the FTC or consumers to take preventative action against risky data handling practices.

Children's Online Privacy Protection Act (COPPA)

The FTC-enforced COPPA applies to operators of websites and online services directed to children under age 13, or to operators with actual knowledge that they are collecting, using, and/or disclosing personal information from children under 13 (Code of Federal Regulations, 2025. The statute requires operators to provide notice to parents and obtain verifiable parental consent prior to collecting, using, or disclosing children's personal information. It also requires operators to maintain reasonable procedures to protect the confidentiality, security, and integrity of such information.

COPPA offers targeted protections for children but has limitations to consumer wearable devices. These devices are generally marketed to a broad audience and may not be classified as child-directed services. As a result, COPPA protections may not apply in cases where children use general-purpose wearable devices but the operator does not have actual knowledge of their age. This leads to another legal gap: because enforcement relies on services having actual knowledge of users' ages, platforms can structure their processes to avoid acquiring knowledge about the age of their users. In addition, the statute establishes general requirements for data protection but does not specify technical standards such as baseline encryption standards or detailed implementation criteria. Discretion is thus granted to operators and compliance is subject to interpretation by regulated entities, potentially resulting in uneven protections across platforms.

Finally, COPPA relies heavily on notice and verifiable parental consent as its primary safeguard. In a complex and evolving digital ecosystem involving data sharing, third-party analytics, and AI-driven inference, parents may not meaningfully understand the scope of downstream

processing. Even if consent is technically valid, it may not actually limit secondary uses or long-term data aggregation.

Together, all three of these FTC enforcement tools provide enforcement authority, but they are limited in scope. The FTC "has been the chief federal agency on privacy policy and enforcement since the 1970s" (Federal Trade Commission). However, the protections are largely triggered after something has already caused harm to the consumer. Consequently, reliance on the FTC as the primary privacy regulator results in reactive and limited enforcement mechanisms that fail to provide comprehensive, proactive protections for consumers.

Summary of Federal Frameworks

Collectively, federal frameworks provide enforcement authority and targeted protections but do not establish a comprehensive regulatory structure governing consumer wearable health data. Key limitations include reliance on reactive enforcement mechanisms, limited jurisdiction over non-covered entities, and the absence of prescriptive requirements for data practices across the data lifecycle.

State Level Privacy Frameworks

In the absence of a comprehensive federal privacy law, individual states have enacted their own data protection statutes. These laws vary in scope, definitions, enforcement mechanisms, and consumer rights, resulting in a fragmented regulatory landscape. This variation is indicated in Figure 1 below. The continued growth of privacy rights, in addition to being state constrained, remains nascent in its concerns for data specifically derived from wearables. The following sections discuss some of the most notable state privacy laws in the U.S. today, including those existing in California and Colorado.

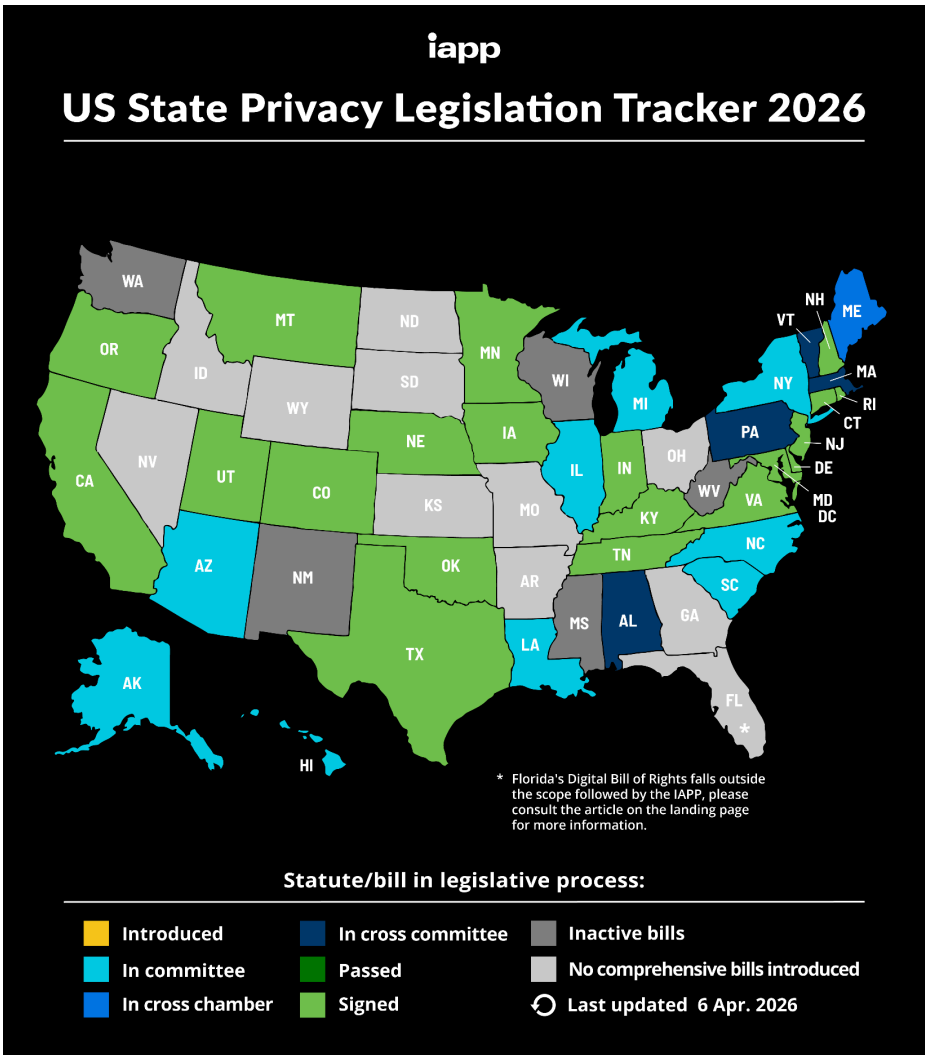

Figure 1: Fragmentation of U.S. State Privacy Laws (IAPP, 2026)

California Consumer Privacy ACT (CCPA)

The CCPA, enacted in 2020, establishes a set of consumer rights for California residents, including the right to access, delete, and opt out of the sale of personal information, as well as protections against discrimination for exercising these rights (State of California Department of Justice, 2024). Amendments to the law have expanded these rights to include correction of inaccurate information and limitations on the use of sensitive data. The law, offering a private right of action and enforced through California's Attorney General, applies to certain for-profit

entities that meet specified thresholds and defines personal information broadly to include identifiers, biometric data, geolocation data, and inferences derived from such data.

The CCPA is a continuously evolving legal framework that looks at the current digital ecosystem. A CCPA amendment enacted in early 2026 is the Drop and Delete Act (State of California, 2026). It aims to give consumers even more control over their personal information. Consumers are allowed to request that all data brokers delete their data all at the same time, which saves time and provides autonomy to data subjects. Nonetheless, the CCPA does not specifically address the unique characteristics of wearable health data. Furthermore, much of the CCPA relies on a notice and consent based framework, which places the burden on consumers to make requests and navigate disclosures to exercise data privacy rights.

Colorado Privacy Act (CPA)

Enacted in 2023, the CPA provides consumers with rights to access, delete, and correct their data. Additionally, citizens can opt out of the sale of their data and prevent its use for targeted advertising, and sometimes, profiling (Consumer Privacy Act). The CPA includes provisions related to sensitive data, including biometric and health-related information, and requires data protection assessments for high-risk processing activities including risk assessments and active consumer consent.

The CPA is particularly relevant to smartwatch technologies because it applies to sensitive data categories such as biometric identifiers and health-related information. The CPA regulates the outcomes from inferences made from consumer's personal data, but not the inferences themselves. For example, if a company is collecting smartwatch data to determine that a user is high-risk for a health condition, and this is used to deny them access to insurance or employment opportunity, this would be illegal as it is profiling with legal or significant effects. Under the CPA, the consumer can opt out of this profiling and potentially challenge the decision. However, if these inferences were used to show the user targeted ads, adjust the content they see, or if their profile was sold to advertisers, this use case would be legal. Although the CPA does offer strong consumer protections, there still exist limitations, especially as a framework not specifically designed to address the continuous, inference-driven data ecosystems of modern wearable technologies.

Fragmentation and Compliance Challenges

The variation among state privacy laws creates differences in consumer protections depending on geographic location, presenting operational challenges for organizations that must comply with multiple regulatory regimes (*What Good and Effective Data Privacy Accountability Looks Like: Mapping Organisations' Practices to the CIPL Accountability Framework*, 2020). Estimates indicate that, in the absence of a federal privacy framework, compliance with state-level requirements could impose substantial costs on businesses operating across jurisdictions. These costs include legal analysis, system modifications, and ongoing monitoring of regulatory changes.

International Framework: General Data Protection Regulation (GDPR)

The European Union's General Data Protection Regulation establishes a comprehensive enforceable framework governing the processing of personal data (GDPR, 2018). The regulation applies to organizations operating within the EU as well as those offering goods or services to EU residents. The regulation emphasizes substantive data protection principles, enforceable individual rights, and organizational accountability requirements. GDPR establishes core principles, including lawfulness, fairness, transparency, data minimization, accuracy, storage limitation, security, accountability, and legal bases for lawful processing (GDPR, 2018). It also imposes extra protections for children, demands functional de-intenfication of inference heavy data-driven systems, and expands the rights of data subjects. These include:

- The right to clear and concise language regarding their data rights
- The right to complain in certain circumstances
- The right to obtain and correct incorrect information
- The right to be forgotten
- The right to the restriction of processing
- The right to object
- The right to data portability.

Additionally, automated processing cannot be used to make a decision about a data subject.

The regulation imposes obligations on data controllers (under whom data brokers fall) and processors, including data collection limitations, requirements to conduct Data Protection Impact Assessments for high-risk processing activities, and compliance expectations when transferring data internationally.

Some limitations still exist within GDPR despite its comprehensive and generally proactive nature. While consent standards are heightened, GDPR still relies on notice-based mechanisms. In complex smartwatch ecosystems involving multiple processors and analytics partners, meaningful comprehension is difficult to achieve. Consent may be legally valid but functionally uninformed, particularly when downstream inference generation and cross-platform data aggregation are opaque.

Although GDPR provides clear and strong privacy rights for consumers, its vague language on data protection and notice-based consent platform still places some burden on the consumer to protect their online privacy. Despite these limitations, the strengths of the GDPR allow for drawing inspiration from one of the most comprehensive privacy frameworks currently in existence.

Voluntary Frameworks

The National Institute of Standards and Technology (NIST) has developed two privacy-relevant frameworks to act as guidance for organizations. The following will discuss the NIST Privacy Framework and the NIST AI Risk Management Framework. Both frameworks have the intention to guide companies to implement consumer privacy and data protection practices by establishing accountability and best practices.

NIST Privacy Framework

The NIST Privacy Framework is a risk-based tool designed to help organizations identify, assess, manage, and communicate privacy risks created through the processing of consumer data across the data lifecycle (NIST, 2025). This framework is modeled after the NIST Cybersecurity Framework and is designated to be adaptable across sectors, emphasizing communication between technical and organizational stakeholders including product design, vendor management, and enterprise risk processes. It is useful for addressing contextual privacy harms and downstream impacts on individuals that may not be captured by compliance-driven operations.

However, beyond its voluntary nature, this framework presents notable limitations for addressing consumer privacy for smartwatches. Much like other NIST frameworks, it avoids prescriptive requirements such as data minimization thresholds, retention limits, limitations on behavioral inference, or commercializing biometric data. Organizations can be in compliance with this framework while engaging in data aggregation, profiling, surveillance, and inference practices.

NIST AI Risk Management Framework

Areas where the AI Risk Management Framework (RMF) provides value within this context include lifecycle structure of AI in the functions of govern, map, measure, and manage. The govern function is focused on establishing AI policy, roles, accountability, documentation, and internal oversight of systems (NIST, 2023). This area within many industry organizations is weak at best at this point in time. The map function is focused on contextualizing the AI system in the scope of its purpose, stakeholders, data, operating environment, and its potential impacts. Measure is the primary source of assessment for AI risks. It calls for qualitative and quantitative methods to assess performance, but does not provide any explicit methods of measurement. Unfortunately, like the NIST Privacy Framework, this is fully voluntary and there is no quantifiable way to implement and measure the idealized functions of the AI RMF.

The voluntary and non-prescriptive nature of the AI RMF is further complicated by the fragmented data protection landscape in the U.S. Organizations attempting to implement responsible AI practices must navigate a patchwork of state-level privacy laws and sector-specific regulations, often requiring significant legal and compliance resources to continuously monitor and interpret evolving requirements. This creates substantial costs for businesses that aim to comply in good faith, while also straining regulatory bodies such as the FTC, which must oversee thousands of companies with inconsistent privacy practices under broad consumer protection authority. For consumers, this fragmentation results in a lack of transparency and standardization, making it nearly impossible to meaningfully compare products or understand how their data is being used across different platforms. Collectively, these dynamics undermine the effectiveness of voluntary frameworks like the AI RMF by placing disproportionate burdens on all stakeholders without ensuring consistent or enforceable privacy protections.

Existing Harms to Consumers

American consumer protection policies tend to be disclosure-heavy, meaning that the law sometimes requires companies to disclose to the consumer the data being collected and how it will be used. However, there are multiple issues with this method. The first issue identified is the complexity of End User License Agreements (EULAs), which often rely on technical jargon that users fail to comprehend. In these cases, users frequently accept terms without understanding the implications.

It is unfair to place the burden of reading these lengthy and confusing EULAs on the consumer. In 2001, it was estimated that reading privacy policies in their totality would cost the American Internet user about 201 hours in time per year. This is equal to about $3,534 annually per person. If every single American read "privacy policies word-for-word, we estimate the value of time lost as about $781 billion annually" (McDonald; Cranor, 2001). This demonstrates that reliance on notice-and-consent mechanisms is not only impractical but also structurally ineffective, as the sheer time burden makes meaningful informed consent unrealistic for the average consumer. Additionally, EULAs are extremely difficult to understand and interpret. Sometimes, even lawyers interpret EULAs differently from each other, leading to discrepancies even at the professional level. If experts come to differing conclusions about the state of a consumer's privacy, it is unfair to assume that a consumer can fight for their own privacy rights.

Overall, the current ecosystem relies on disclosure, prioritizing vendor protection over user privacy. The lack of a nation-wide protection leads to consumers not understanding their own privacy rights, making them even more vulnerable to potential harms.

Private Market Compliance Costs

Related to the existing patchwork of laws is the growing compliance burden arising from consumer health data regulations. The foundational challenge for manufacturers, application developers, and data intermediaries is jurisdictional ambiguity. HIPAA only applies to healthcare providers or one of their business associates. Consumer smartwatch data, which flows into third-party apps and cloud platforms instead of clinical systems is not governed by HIPAA, but rather by a patchwork of state privacy laws and lengthy terms of service agreements (Peremore, 2025). This gap does not relieve companies of compliance burdens, but may multiply

them. Firms operating across state lines must independently assess exposure under each applicable state framework which creates overhead costs for compliance operations.

It is estimated that in the absence of a federal privacy law, state privacy laws could impose out-of-state costs between $98 billion and $112 billion annually, with costs exceeding $1 trillion over a ten-year period. Small businesses would bear $20 to $23 billion of this burden annually (Castro et al., 2022). These figures reflect across industry generally, but the smartwatch sector is primarily exposed given its position between consumer technology and sensitive health data.

Technical Implications and Challenges

The following section analyzes each phase of the data lifecycle and identifies where the most consumer-relevant vulnerabilities lie within the smartwatch ecosystem. The data lifecycle is defined by six distinct stages: the *On-Device Sensing and Operating System (OS) Telemetry*, *Companion App Ingestion and Transformation*, *Cloud Storage and Access*, *AI/ML Model Training and Productization*, *Data Sharing with Third-Parties*, and finally, *Data Retention, Deletion, Portability, and Account Disclosure*.

Device Capabilities and Data Creation

Modern consumer smartwatches are often sophisticated, multi-sensor platforms that provide continuous and passive granular data collection across behavioral, physiological, kinematic, and environmental domains. Sensor capabilities include heart rate, blood oxygen saturation, body movement, activity classification, fall detection, emotional arousal, skin temperature, circadian rhythm analysis, location and GPS tracking, and microphones. These devices generate data streams that the device itself aggregates and transmits without requiring direct user engagement.

Data collected by these devices are not solely descriptive of user behavior. Data can be used to infer a large amount of information about any one person's health, presence of medical conditions, as well as their habits. Heart variability patterns can infer autonomic nervous system dysfunction or chronic stress. From menstrual tracking combined with resting heart rate and temperature data, pregnancy status and fertility windows can be discovered. Downstream inferences such as these expand the sensitivity of what appears to be raw biometric data and are frequently generated without explicit or well-communicated user disclosure or consent (IAPP, 2025).

Smartwatch sensors have been known to provide varied levels of accuracy across demographic groups as discussed in detail in the later section: *Protected Health Status*. When this data is then utilized in insurance underwriting, employment wellness screenings, clinical decision support, or any other high-stakes decision making, they may systematically disadvantage specific demographics and inadvertently compound existing inequities. Organizations have already been providing financial incentives to individuals who provide this health information from their wearables. United HealthCare currently offers incentives between $300-$1,000 per year based on health goals and activities (UnitedHealthCare, 2025). If readings are inaccurate, systemic price discrimination could occur even though its racial impact was unintended. The expansion of wearable devices and the data that they collect, such as with Meta's AI Smart Glasses, raises concerns. The capability to collect audio, videos, and photos when activated poses even larger challenges for consumers to adequately protect themselves from data exposure and misuse, even if they are not wearers of such devices.

The Data Lifecycle: Vulnerability Points

The path that consumer health data takes from the sensor to downstream commercial and analytical uses creates multiple vulnerability points with distinct security and privacy risks. The following will map each out of the identified seven stages in the technical lifecycle with their related technical vulnerabilities.

Stage 1 — On-Device Sensing and Operating System (OS) Telemetry: Raw sensor data is collected locally and stored temporarily on the wearable device before transmission to companion applications and services. Vulnerabilities within this stage include inference attacks related to ambient device signals and radio tracking. Inference attacks may infer information such as power consumption or Bluetooth broadcast patterns. Radio tracking may permit physical location monitoring through Bluetooth and Wi-Fi beacon detection. Additional risks include outdated firmware or unpatched software vulnerabilities. Consumer wearables tend to receive infrequent or delayed security updates, and older devices may not receive any patches.

Stage 2 — Companion App Ingestion & Transformation: Health data transferred from the device to companion applications creates multiple opportunities for attack. Man-in-the-Middle (MITM) attacks exploit insufficiently secured Bluetooth pairing protocols or unencrypted application program interface (API) calls to intercept data in transit. A documented implementation flaw in the WhisperPair protocol, which was originally developed for wireless

headphones, exemplified the feasibility of unauthorized device pairing in real-world deployments. Additionally, insecure third-party integrations such as health data API shared with fitness apps, nutrition trackers, or employer wellness platforms, create additional transmission vectors with varying security postures depending on the external party. Companion applications typically request sensor permissions that are beyond necessary for their stated functionality and app store security reviews have been documented to be inconsistently enforced (Man et al., 2023).

Stage 3 — Cloud Storage & Access: Once data is passed to cloud infrastructure, vulnerabilities primarily relate to misconfigured access controls, inadequate encryption at rest, weak identity and access management practices, and insecure third-party integrations. Real-world examples demonstrate the reality of these risks: In 2021, a database breach exposed approximately 61 million records from Apple and Fitbit users following a failure to encrypt cloud-stored health data (Man et al., 2023). In 2020, Garmin suffered a major ransomware attack that disrupted services and heightened concerns about data exposure. Additionally, the Identity Theft Resource Center documented 2,620 data compromise incidents in 2024. Of those, 537 involved the healthcare sector. This was the highest per-year total the healthcare industry had since tracking began, signaling a trend that these types of breaches will continue to increase (Identity Theft Resource Center, 2025). In 2024, more than 276 million patient records were compromised in healthcare data breaches (Winder, 2025).

Stage 4 — AI/ML Model Training and Productization: In the analytics stage, health data is aggregated with data from other sources and used to train predictive models. Once trained, they are processed into inference outputs such as health risk scores, behavioral profiles, and personalized recommendations. Technical vulnerabilities at this stage involve several AI-specific privacy attack classes:

1. Membership inference attacks are attacks that allow adversaries to discern whether a specific individual's data is used as part of a model's training set. Through this, an adversary may be able to reveal sensitive health statuses from model outputs alone.

2. Attribute and property inference attacks are attacks where sensitive personal attributes may be reconstructed from model parameters and outputs, even when raw data is not directly accessible.

3. Data reconstruction attacks in federated learning (where model weight updates instead of raw data transmissions) can allow for partial reconstruction of raw training data from gradient updates. Optimization-based reconstruction attacks have demonstrated to be effective against small-batch training data (Zhu et al., 2019), whereas linear layer leakage attacks may enable almost exact reconstruction of specific user data points of fully connected layer activations and input data (Zhao et al., 2023).

4. Model inversion attacks can reconstruct representative examples from training data in model parameters. Effectively, personal biometric profiles and health data can be derived through this means.

An additional concern is the dichotomy between explainable AI (XAI) and privacy. Techniques utilized to make AI systems interpretable like SHAP, LIME, or counterfactual explanations can leak private training data by providing adversaries differential access to model's inner workings. A 2024 study found that XAI explanations may enable member inference, attribute inference, model extraction, and reconstruction attacks. There is no current mitigation that fully resolves the tradeoffs between transparency and privacy (Nguyen et al., 2024).

Stage 5 — Sharing with Partners, Software Development Kits, & Data Brokers: The commercial data ecosystem of downstream use of consumer wearable data is one of the least regulated and most opaque. Software Development Kit (SDK) integrations in companion applications, and indirectly through real-time bidding (RTB). In RTB, apps broadcast user data attributes and health-adjacent behavioral signals to advertising networks. These networks distribute its requests to potential advertisers. Participants in the bidding process receive broadcast user data even if they do not place the winning bid, creating a mechanism by which data brokers can collect personal data profiles by participating in the bidding ecosystem without making advertising purchases. This proxy channel operates largely outside of consumer awareness and existing regulatory frameworks (Electronic Privacy Information Center, 2025).

The risks of this ecosystem are compounded by near-total absence of regulatory oversight governing how data brokers acquire, store, and secure the data they collect. The industry's opacity is not incidental, but structural. Databrokers are incentivized to obscure their data sources and the persistence of data access broker networks means the information (once sold and sold again) may circulate indefinitely with no meaningful path to removal (Reviglio, 2022). A January 2025 breach of data broker Unacast (attributed to a compromised AWS authentication

key) resulted in the exposure of approximately 30 million location data points, with the chain of acquisition of the data remaining unclear (Smalley, 2025). This incident illustrates the risk that consumer health and location data, once acquired by brokers through legitimate commercial channels, becomes a high-value target for external attackers with no direct relationship to the original data subject.

Individual consumer data has been estimated to have a market value of approximately $700 per person per year (Wolford, 2024). The data broker market, valued at approximately $349 billion in 2024, is projected to reach $612 billion by 2032, which emphasizes the significant commercial incentives driving data collection and secondary use within this ecosystem.

Stage 6 — Retention, Deletion, Portability, and Account Disclosure: The final stage of the data lifecycle presents vulnerabilities related to excessive retention, failure to purge data for deleted accounts, and the legal tension between deletion rights and litigation hold obligations. Important cases include the following:

1. FTC v. Amazon (2023): Amazon was found to have violated COPPA by retaining voice recordings of children collected by Alexa-enabled devices indefinitely. The resulting $25 million settlement represented less than 0.01 percent of Amazon's net income at the time. This highlights questions about the deterrent adequacy of per-violation penalties (Federal Trade Commission, 2023b).

2. Brown v. Google (2024): A class action settlement required Google to implement hard deletion of user data collected during Incognito mode browsing sessions, addressing a practice where "deleted" data was soft-deleted and retained in accessible storage (Navlakha, 2024).

3. New York Times Co. v. OpenAI (2025): A court-ordered litigation hold requiring OpenAI to preserve under conversation data that would otherwise have been subject to deletion created a direct conflict between a user's right to erase and a third party's right to discovery: a clear structural tension between privacy rights and legal process obligations (Kelly, 2025).

These cases relate to the concept of "deletion" in commercial data systems is not technically nor legally straightforward. Data may persist through backups, be retained by third party processors,

or be subject to litigation holds. It may also be deleted from storage, but remain reconstructible from trained AI model weights.

Identity Theft and Data Breaches

Identity theft and data breaches are an enormous problem for consumers and create a direct financial impact. Oftentimes, data breaches occur due to companies having poor cybersecurity and data protection practices. To underscore the importance of this issue, in 2024, Mozilla's Foundation report on bodily integrity in the digital age found that health-related cybersecurity breaches and ransomware attacks targeting wearable and body-centric data increased more than 4,000 percent between 2009 and 2023, and that the market for body-centric data is expected to exceed $500 billion by 2030 (Mozilla Foundation, 2024). The report found that nonconsensual data sharing by wearable companies is widespread and that consumers are largely unaware of how much of their data is being collected and shared. For individuals impacted by privacy breaches, the implications may be nominal or could wipe out their financial assets. There is no existing literature at this time on the direct costs to consumers as a result of privacy breaches. This is an area of opportunity for future research.

The technical vulnerabilities that this paper has discussed demonstrates the tangible consumer impact that poor cybersecurity practices can lead to. Oftentimes, there are few laws that mandate that companies utilize modern cybersecurity practices that protect consumer data against threats.

Merger and Acquisition Considerations

Data-driven acquisitions have become increasingly dominant in the market as organizations seek to leverage acquired data to enhance innovation and productivity. The process of data acquisition by one company and re-acquisition by another through business transformations are nuanced and carry implications for consumer privacy in regards to their data. Although intangible, collections of bankruptcy codes, accounting standards, and consumer protection laws treat data as a transferable business asset. Section 363 of the Bankruptcy Code allows a company in financial distress to sell its properties including customer databases to a buyer (11 U.S. Code § 363 - Use, Sale, or Lease of Property, 2019). Similarly, Accounting Standards 350 legally classifies acquired data as intangible assets recorded in the balance sheets. Once data ownership has been transferred, consumers are likely to face privacy drifts where their data is

applied to new secondary uses that diverge from the original intent of collection (Financial Accounting Standards Board, 2021).

The first exposure to consumers is modern privacy policy that usually contains a "Business Transfer" or "Merger & Acquisition" clause that enables the transfer of data ownership. For example, in Meta's privacy policy they state: "we sell or transfer all or part of our business to others, and in a number of cases, we will provide information about you to the new owner as part of such transactions, but only if permitted by law," (Meta, 2022) . Most privacy policies include these clauses by default, and, as discussed previously, the shortcomings of privacy policies include the discrepancy of consumer and company expectations which stems from policy fatigue and lack of consumer legal comprehension.

A major conflict arises when the new data owner intends to use the acquired data differently from the original company, such as a fitness app being acquired by an insurance firm to appropriate the health monitoring information to price insurance premiums. However, there are protections when data ownership transfer strays from the original purpose of the data listed in Section 5 of the FTC "Deceptive Practices" rule. Changes in data usage falls under material change to a legal contract, so companies must either 1) continue data operations using the old company's procedures without alterations or 2) obtain "opt-in" consent. The "opt-in" consent involves a clear alert that users must interact with before proceeding and an additional notice that explicitly states the changes. For example, in 2000, Toysmart was an online toy retailer that planned to sell its customer list as a standalone asset during its bankruptcy. Interestingly, their privacy policy promised that personal information would "never be shared with a third party", so the FTC intervened in the customer list sale stating that selling the data would be a deceptive practice because it contradicted an earlier policy to "never be shared". Ultimately, the sale was allowed given 1) the data must be sold as part of the entire business 2) the buyer must be in the same industry and 3) the buyer agreed to abide by the original privacy policy (Federal Trade Commission, 2000).

In 2005, relying on the FTC to intervene for every bankrupt company planning to sell PII was inefficient, so the Bankruptcy Code was amended to include 11 U.S.C. § 332 and § 363 which created a new role called the Consumer Privacy Ombudsman (CPO). CPOs are appointed to protect the privacy interest of customers who are not typically parties in a bankruptcy case and their primary responsibility is to suggest ways of customer data sales that minimizes harm to consumers while maximizing value from the assets (Harvard Law Review, 2025). Bankruptcy

courts refer back to the CPOs recommendation to determine if a sale can proceed, and the court may refuse to approve a sale until the CPO's conditions are met.

However, CPOs have shortcomings such as vague guidelines in the Bankruptcy Code on how CPOs should perform their roles and the qualifications required to appoint CPOs. Currently, the only qualification CPOs must have is to "be disinterested," meaning CPOs must not have an interest adverse to a party in the bankruptcy based on prior relationships with that organization or their connections. Another point is the inherent tension of CPOs and bankruptcy goals: increased overhead, spending, and time are counter to the company's goals of preserving the debtor's value and proceeding with the sale as quickly as possible. Consequently, debtors engage in an array of practices to avoid the appointment of a CPO such as adding terms containing consumer-protective language in data asset sales, submitting privacy reports by experts, and former CPOs at the time of filing. They may also include boilerplate language about consumer data in sales motions. In short, debtors do the work of a CPO before filing to avoid delays if the court were to appoint a CPO, but this removes the CPOs from the process in which Congress intended CPO participation (Harvard Law Review, 2025).

Conversely, when CPOs are appointed, the reports that are submitted to court are criticized for being boilerplate summaries of privacy policies and other legal authorities rather than providing tailored recommendations for the specific case. For example, a CPO report for Chrysler suggested the same privacy mitigations from the Toysmart bankruptcy which occurred in 2000. However, this is reflective of the time pressure CPOs face. The first CPO ever appointed reported he was appointed a mere 10 days before the hearing (which was later delayed), and it only gave him 16 days to "complete the investigation, analysis and drafting of the CPO's report" (Coordes, 2021). Referring back to the Chrysler example, Chrysler had one of the fastest bankruptcies in history, which could explain why the report did not develop additional mitigations beyond what was discussed in Toysmart.

Once a company is acquired, the parent company will inherit the pre-existing practices of the acquired company. Hence, data breaches are still a major risk vector if the parent company inherited existing vulnerabilities from the acquisition. For example, in 2018, Marriott discovered that the guest reservation database of Starwood, which Marriott acquired in 2016, had been compromised as early as 2014 (Miratech Staff, 2019). This breach exposed 339 million guests PII including passport numbers and credit card information which resulted in a $24 million fine by the UK's Information Commissioner Office.

Another potential vector is regulatory inheritance where the parent company inherits data collected illegally of the acquired company. In 2010, Disney acquired a social gaming company called Playdom who, shortly before, acquired a company called Acclaim. Acclaim had been illegally collecting personal information from children without parental consent which violated COPPA. Ultimately, even though Disney did not initiate these practices, it was responsible for the violation from transferring data ownership from Acclaim to Playdom to Disney, finally settling for $3 million which was the biggest COPPA violation at the time (Drye, 2011).

Ultimately, the treatment of data as a transferable business asset reveals a significant gap between corporation goals and protecting customer privacy. While existing guardrails such as the CPO and 'Deceptive Practices' rule have been implemented, they are frequently undermined by the economic pressures of bankruptcy and the shortcuts taken by those firms. With the cases of Marriott and Disney, it underscores that an acquisition is not merely the transfer of assets, but it is also an inheritance of liability - be it through hidden vulnerabilities or regulatory violations. As privacy drift becomes more prevalent, for consumers, the impact is perpetual exposure of sensitive information that stems from the balance sheet taking precedence over the promises made at the point of collection.

Consumer Risks and Harms

The following sections of this report have mapped the technical architecture of smartwatch data collection, the legal frameworks nominally governing it, and the commercial ecosystem through which consumer data flows once leaving the device. This section synthesizes those findings into a direct analysis of risks and harms consumers face as a result. These risks and harms arise from practices across the wearable technology industry and nature of data brokering.

A central theme through this analysis is that the primary risk is not necessarily the data the consumer may knowingly disclose, but rather what commercial and algorithmic systems can derive, infer, and act upon from that data, usually without consumer awareness or legal constraints. When raw data can be processed to infer health statuses like pregnancy, mental health, or sexual orientation, the harm potential of a wearable device extends far beyond what the product's original functions suggest.

Data Brokerages as a Source in Consumer Harms

While wearable devices are the origin point for health data collection, data brokers represent a distinct and underregulated layer of harm in the consumer privacy ecosystem. Data brokers are organizations that collect, aggregate, and sell personal information, often without any direct relationship with the individual whose data they hold. Unlike the wearable manufacturer, the data broker has no product obligation to the consumer, no terms of service the customer has agreed to, and often no obligation to notify the consumer that their data is being bought and sold.

Aggregation of consumer health data is central to harm caused by brokers. A broker may independently hold data points that appear innocuous such as a consumer's step count, zip code, retail purchasing history, and sleep disruption patterns. Individually, none of these reveal a protected health condition, but combined through predictive models, they may reliably infer pregnancy, chronic illness, or mental health status. The consumer never directly disclosed any of these conditions, and no single source would have revealed them. This data persists in the market regardless if a consumer deletes their wearable app or requests deletion from the device manufacturer.

The ecosystem databrokers operate in creates accountability gaps that make harm difficult to trace or challenge. When consumers are denied a loan, charged a higher insurance premium, or served advertisements that exploit a health vulnerability, the decision may have been informed by broker-derived health inferences several steps removed from the wearable device. There is no disclosure requirement, no audit trail for the consumer to follow, and no enforceable right to know what data informed the outcome (Brennan Center, 2026). It places consumers in the position of experiencing consequential harms without means of identifying their source or seeking redress.

Inferring Protected Health Status

Smartwatches collect immense amounts of data from the consumer, both from sensing the vitals of the patient and from the actual information that the patient voluntarily provides the company. Aggregated data allows the company to know or infer a patient's health status, including whether or not a woman is pregnant, disability status, and immunocompromization.

Oura Ring states that their rings are able to track ovulation through the "Physiology Method," which continuously takes a person's finger temperature readings to detect sustained changes in skin temperature, which occurs after ovulation due to the hormonal changes within a woman's menstrual cycle. Oura Ring boasts a 96.4 percent accuracy rate in detecting ovulations. Additionally, studies have shown that the Oura ring can detect pregnancies nine days before a woman receives a positive test result (Oura Team, 2021; Lomas, 2025). By combining multiple pieces of a person's health data, the Oura Ring is able to make health inferences. Related to this issue are menstrual tracking apps: despite widespread use of reproductive health apps, Mozilla Foundation found that 18 out of 25 apps failed basic privacy standards, with many collecting extensive sensitive data, sharing it with third parties, and lacking clear policies on law enforcement access (Mozilla, 2022). Additionally, 84 percent of period tracking apps share data with third parties (ORCHA, 2022). If a woman's pregnancy status was inferred and that information was poorly protected or sold, this information could be used by external parties, leading to harm to the consumer.

Along with accurate pregnancy predictions, wearable health technologies may be able to detect a person's disability status. Some technologies such as the Apple Watch are able to track a person's walking gait patterns. As such, it can determine a person's walking asymmetry, double support time, walking speed, step length, and steadiness. When these different aspects of walking are combined, it can potentially be used to make inferences about a person's overall health and disability status by Apple itself or a third-party that has the collected data (Apple, 2021).

Wearables health devices track blood oxygen levels, skin temperature, activities, sleep, and heart rate. These factors can lead the device to inferences such as illnesses or immunocompromization status by detecting changes from a person's baseline status. As a result, a case in Norway emerged where a wearable device suggested to a user that they were ill with a potential infection and were later found to have Lyme's Disease. Specifically, the device monitored the patient's "deviations from normal heart rate and oxygen levels" to make this diagnosis. There are also concerns that the devices may be able to detect higher levels of C-reactive protein, which are an immune system marker for inflammation. Oftentimes, they can detect "infection, autoimmune diseases, developing cardiovascular disease or even cancer." Additionally, wearable health devices likely already have the ability to detect diabetes and pre-diabetes by inferring whether or not a user has insulin resistance. In a study conducted at

Stanford, researchers made an algorithm using a combination of daily steps count and the difference between daytime and nighttime heart rate. With this information, the researchers were able to predict which individuals in the study were more likely to have insulin-resistance, thus inferring the potential for diabetes (Dusheck, 2018). Both of these examples can be seen as beneficial early detection mechanisms for consumers. However, when external parties are privy to this data and utilization of it, it may have downstream implications on the consumer themselves if utilized in decision making that directly impacts that consumer.

While there are helpful use cases, there are limitations to health wearables. For example, wearable health technologies may even be less accurate for people with darker skin pigmentations compared to lighter skin pigmentations. A study involving the Apple Watch Series 9 was done to detect the accuracy and variability of energy expenditure (EE) and heart rate (HR) between people of different skin tones. The study concluded that the Apple Watch Series 9 had varying EE and HR measures across intensity and skin pigmentation, with higher variability for measures of HR for individuals with darker skin pigmentation. This is likely attributed to wearable devices' use of photoplethysmography (PPG), which relies on light absorption and reflection through the skin. Melanin can interfere with light signals, making it more difficult to have accurate measurements for individuals with darker skin pigmentation.

Having lower accuracy for different skin tones is significant because biased data can lead to inaccurate inferences. If wearable data is used to infer health status, inaccuracies across skin pigmentation groups could disproportionately misrepresent the health of individuals with darker skin, possibly leading to unequal or discriminatory outcomes. Inaccurate health inferences may misinform financial, insurance, and employment-related decisions, disproportionately harming individuals with darker skin, which would exacerbate existing inequalities and discriminatory outcomes (Chase et al., 2024).

Inferring Pregnancy Status to Enforce Laws

The risk of a data breach of consumers' collected health information raises complex legal issues related to women's reproductive health. When Roe v. Wade was overturned in 2022, some privacy experts became concerned about menstruation tracking apps, such as Flo, being used to enforce abortion banning laws. In the case of Flo, experts were concerned that collected data could indicate when a woman has started her pregnancy, and whether or not that pregnancy has

reached full-term. If the app's collected data was to be subpoenaed or sold to a third party, there are concerns that it could be used to infer that a woman has had an abortion.

In 2021, the FTC settled with Flo Health after finding the app (used by more than 150 million users) shared menstrual cycle, pregnancy, and symptom data with Facebook and Google despite explicit promises that such data would remain confidential. The FTC documented that when a user entered pregnancy-related information, Flo transmitted App Events containing the word "pregnancy" to the analytics arms of third parties. The Wall Street Journal was able to intercept this unencrypted health data in transit (Fair, 2021; Torchinsky, 2022).

In the context of wearable health devices, this concern is more relevant than ever. Wearable health technologies are able to collect reproductive data, including whether or not a woman is ovulating, to make health inferences. If the data is sold, leaked, or subpoenaed, it could be devastating to a person that decided to exercise their reproductive freedoms. Having few regulations on wearable health technologies could lead to excessive government overreach related to women's reproductive health.

Dynamic Pricing, Insurance Discrimination, and Financial Access

Health data may be shared with insurers with the incentive to reduce premiums. Insurance companies may also use this health data to deny coverage or increase premiums for individuals with certain health conditions. Banks and financial institutions could also utilize these health records to deny loans of mortgages to individuals with specific illnesses to reduce default risk. In some cases, health data could be used to limit access to certain health services for extortion purposes (Bernard, 2024). The current regulatory landscape creates a significant vulnerability in this area. HIPAA covers health data held by traditional healthcare providers, but established earlier, the same physiological data collected by a consumer smartwatch falls outside that protection. This means that insurers and financial institutions are not formally prohibited from acquiring inferred health information from data brokers, even when the information was originally created on a wearable device worn in a non-clinical setting. A leaked symptom history can affect access to insurance, employment, or credit, and a disclosed mental health concern may resurface years later in a context of social stigma or professional discrimination. Once integrated into large-scale AI systems, health data may be copied, transferred across borders, combined with other datasets, and retained indefinitely (Vandenberghe, 2026).

Employment Discrimination

The sale of customers' data to advertisers has raised additional concerns about where that data ultimately ends up. One major concern is the sale of customer data to potential employers. As previously established, wearable health technologies are able to infer a myriad of information about a person's health, including pregnancy status, disabilities, immunocompromization, and illnesses. If employers were sold an individual's health data or the inferences made about that health data, the damage could be immense. It is possible that an employer would not hire someone based on their health status, which is clear employment discrimination. Although employment discrimination based on a person's health status is illegal, it could be difficult to prove in a court of law (Sivakumar et al., 2024). The Electronic Frontier Foundation notes that employers may access certain health-related information during background checks, typically with an applicant's written consent. However, wearable health data introduces a new and less transparent pathway. Rather than directly requesting health information, employers could obtain or infer sensitive health conditions through third-party data brokers or analytics, potentially bypassing traditional consent mechanisms and making discriminatory practices significantly more difficult to detect or prove (Electronic Frontier Foundation, 2026).

Financial Discrimination

Typically, an individual's creditworthiness is determined from their credit score, borrowing history, and repayment history. However, recent development in the financial sector has allowed for advanced credit predicting models that assess creditworthiness. Using these advanced models gives the financial industry the ability to use a person's alternative data to gain predictive power and subsequently, a competitive edge against their competitors.

The World Bank Group has expressed concerns about certain types of data, including health data, being used to determine a person's creditworthiness. The World Bank Group states that it is unethical and harmful for consumers' privacy to consider health data in a creditworthiness determination (2024). In fact, the European Data Protection Supervisor (EDPS) suggested that regulatory agencies (in this case, European legislation) should "clearly define the categories and sources of personal data that should and should not be used for the purpose of creditworthiness assessment." One of these suggested protected categories includes health data (Federal Trade Commission, 2014). The use of alternative data poses great risks to consumers. The GAO discusses that health data could lead to discrimination, violation of fair lending laws, and risks

for data privacy (Government Accountability Office, 2019a). When combined with the growing ecosystem of data brokers and wearable health technologies, these risks are further amplified. Wearable devices generate continuous streams of highly personal health data, which can be sold, aggregated, and incorporated into creditworthiness assessments without the consumer's knowledge or consent.

As a result, individuals may face adverse financial outcomes, including loan denials or higher interest rates from inferred or inaccurate assessments of their health. This creates a new form of financial discrimination that could be difficult to detect, regulate, or challenge, ultimately undermining both consumer trust and the fairness of the financial system.

Targeted Manipulation

Beyond direct opportunities for discrimination, consumer health data may enable a distinct category of harm: the use of intimate physiological and behavioral signals to craft psychologically targeted advertising and influencing campaigns. This is not limited to product advertisement, but can extend to political candidate advertisement as well.

Targeted advertising tailors content to specific individuals based on their characteristics, behaviors, and preferences, which raises significant ethical concerns in digital business environments regarding user privacy and data manipulation. This includes emotional targeting that may exploit vulnerabilities to influence purchasing suggestion and scarcity tactics that create a false urgency to drive impulse buying (Fiveable, 2025).

The concern is intensified when health data is incorporated into this targeting infrastructure. The FTC has explicitly warned about dark pattern tactics, which includes disguising advertisements to look like independent content, burying key terms, and tricking consumers into sharing their data. This represents a direct threat to consumer choice. The pervasive nature of these data collection techniques allow businesses to adapt ads to target specific demographics or even a particular consumer's interests (Federal Trade Commission, 2022). Dark patterns often take advantage of consumer's cognitive biases to steer their conduct or to delay or deny access to information they need to make informed decisions. Research suggests that dark patterns used in combination can have an even more harmful impact on consumers (Fair, 2022).

The BetterHelp enforcement action illustrates the real-world stakes of this particular dynamic. The FTC found that BetterHelp, an online therapy platform, had shared users' mental health

data including their seeking out therapy, to advertising partners including Facebook and Snapchat, despite explicit assurance to the contrary. Resulting ads were targeting users at moments of heightened psychological vulnerabilities. BetterHelp was only fined $7.8 million. The settlement affected approximately 800,000 consumers. BetterHelp's sharing attracted tens of thousands of new paying users and generated millions in revenue, meaning the fine meant little to nothing to the economic benefit derived from the violation. The FTC documented that BetterHelp used consumers' email addresses and therapy to instruct Facebook to identify "similar" users for ad targeting. An estimated 5.6 million individuals were served with advertisements via Criteo (a marketing platform) alone as a result of unauthorized sharing. (Federal Trade Commission, 2023a; Federal Trade Commission, 2024). This case demonstrates the intersection of sensitive health data and behavioral advertising infrastructure producing uniquely harmful outcomes that extend beyond typical privacy violations. Coupling instances like these with specific, derived health conditions from wearable devices could be used for highly targeted pharmaceutical or alternative medicine advertisements which could have major health ramifications for an individual.

These inferences may also be leveraged to exploit addictions or addictive behaviors. Research identifies addiction as both a harm in itself and as conducive to harming others, finding that it is often the result of a combination and interaction of dark patterns and personalization, especially for targeted advertising (Santos et al., 2025). Recommender systems, which connect addictive design with lack of self-control and self-determination. Children are particularly exposed to targeted advertising by systems that are designed to maximize engagement and time-on-platform, and health-adjacent behavioral data from wearables worn by or near children can be incorporated into these systems. The downstream effects on children's mental health, compulsive behaviors, and susceptibility to manipulative advertising represent a harm category that existing frameworks like COPPA are not equipped to address.

AI Inference, Data Persistence, and Limits of Consent

The most structurally distinctive harm arising from consumer wearable data is one that does not depend on a breach, leak, or deceptive practice. It arises from normal functioning of the commercial data ecosystem. As described in the technical lifecycle section of this document, wearable device data is routinely aggregated, processed, and utilized to train predictive models. The output of those models, such as health risk scores, behavioral profiles, and inferred

conditions, are utilized in targeted advertising, commercial profiling, and potentially third-party contexts without any direct disclosure of the inferential process to the data subject.

AI relies on vast quantities of data that travel from one context to another, draws inferences that were never present in the data, and can be used in unforeseen ways. In the consumer space that relies on health data, advertising and analytics companies purchase consumer data to run AI that allows them to target advertising to users based on derived characteristics they did not explicitly provide. Health insurers purchase data for a range of purposes, and reidentification is increasingly feasible: when a supposedly anonymized dataset is released, data from that set can be linked to another dataset with information regarding identified individuals. AI makes such reidentifications more plausible when certain characteristics are not present in a released dataset (Konnoth, 2024). A study published in JAMA Network Open found that machine learning algorithms could successfully reidentify deidentified physical activity data at rates approaching 95 percent for adults, even when the data had been aggregated and stripped of identifying information (Na et al., 2018).

Even if a consumer has never received a formal clinical diagnosis, a model may infer if that individual is at elevated risk for depression, attempting to conceive, or is managing a chronic condition. This can be used to deploy targeted advertising for diet supplements, fertility products, or speculative treatments. On another layer, these inferences may be utilized in political advertising, or leveraged to make decisions such as with loans or health insurance. The consumer will experience these effects of their data being used in this way without understanding that they were generated from inferences drawn from wearable health data. This is a notable breakdown in notice-and-consent models that existing legal frameworks rely upon: consumers cannot consent to uses they were not informed of, and they cannot opt out of inferences if they do not know they are being drawn.

Traditional consent frameworks are not designed for an inferential environment. The passive and pervasive nature of data collection, opacity of model inference, and risk of algorithmic discrimination call to question the adequacy of existing regulatory frameworks. Technical sophistication and clinical utility may obscure the normative decisions of automated health decision making. Scholars have argued that the regulation of inferences from derived data should receive at least as much legal attention as the regulation of inferences derived from data as the underlying data collection itself. Without that shift, the more granular and intrusive the

model, the gap continues to widen between the scope of actual data use and the protections available to the consumer (Radanliev, 2025).

Data persistence then magnifies these harms across time. As established earlier in the data lifecycle analysis, "deleted" data may persist through backups, third-party processors, and AI model weights. A consumer who revokes consent, closes an account, or requests deletion of their data is not able to ensure that inferences drawn from that data (or the models trained on it) are removed. This exposure created by persistent health data means that a single period of wearable device use can generate downstream commercial and algorithmic consequences years after the consumer has stopped using the device.

Consumers are often unaware that their devices are collecting their data and even their private conversations. Oftentimes, these conversations were then shared with third-parties apps for purposes beyond the original scope of the purpose of collection. For example, In January 2025, Apple settled a class action lawsuit for $95 million over allegations that Siri (including Apple Watch) activated without the "Hey Siri" trigger phrase and recorded private conversations, which were then shared with third-party contractors for review. The action covered owners of Siri-enabled devices between September 2014 and December 2024. An estimated 85.2 million eligible users, of whom approximately 97 percent never filed claims. One plaintiff alleged he received targeted ads for surgical treatments immediately after a private conversation with his physician (All About Lawyer, 2025). With the pervasive nature of health wearables in today's society, protecting individuals' privacy, especially related to their health conditions, becomes ever more important.

Consent and Consumer Disempowerment

Underlying each of the harm categories discussed above is a structural problem of consent and control. As established in the legal frameworks section, existing regulatory protections (Section 5 of the FTC Act, HIPAA, HBNR, and COPPA) each rely in different ways on notice-and-consent mechanisms as a primary safeguard. This reliance is ultimately inadequate for the wearable data ecosystem in ways that are not simply technical, but structural.

Scholars and commentators often argue that individuals do not care about their privacy and that users routinely trade privacy for convenience. However, this framing ignores the cognitive biases and design tactics that platforms utilize to manipulate users into disclosing information,

many times not optional for use of a specific service. This underscores how current laws allow manipulative and anti-consumer behavior to persist (Waldman, 2020).

The consent problem, which is that privacy policies are poor mechanisms to communicate with individuals about privacy (Cranor, 2025), in the wearable context is compounded by the technical opacity of inference creation. A consumer may understand from a general perspective that their smartwatch collects heart rate data. However, they are unlikely to understand that this data will be combined with location signals, behavioral patterns, and third-party commercial data to generate a predictive health profile that is then used to price services, target advertising, inform underwriting decisions, or be utilized in political campaigns. The data subject is not the data user, and the gap between what the consumer discloses and what the system derives is precisely where the most serious harms to human autonomy and dignity occur.

Further, this issue is also compounded by the sheer volume of documentation consumers encounter. Even if a consumer were motivated to read the privacy policies and EULAs governing every digital service they interact with annually, doing so would be functionally impossible. McDonald & Cranor demonstrated that fully reading privacy policies translates to 53.8 billion hours of time and an estimated opportunity cost between $559.7 billion to $1.1 trillion annually in lost productivity, and reading privacy policies carries a cost in time of 201 hours a year (2008). These figures were calculated before the mass proliferation of wearable devices and companion applications, which almost certainly underestimates the burden facing consumers today.

Future Concerns

Privacy risks posed by consumer wearables are not static. The emergence of AI-enabled smart glasses illustrate how the same data collection dynamics documented in the smartwatch context extend and intensify as wearable technology becomes more sophisticated and intimate. A primary example of this are the Meta Ray-Ban smart glasses.

Meta updated its privacy policy for its smart glasses in April, 2025, which removed the option to prevent voice recordings from being stored and requiring always-on AI features. Voice data is now retained for up to a year to enhance AI training. This raises significant concerns and is sparking public debate about whether user's consent and control over personal data have been compromised (Ferguson, 2025). In 2025, over seven million people bought these smart glasses and footage from those users is fed into a pipeline for review. Users are unable to opt out (Perez,

2026). Contractors review this footage, which includes extremely private content far beyond public scenes such as footage of bathrooms, bedrooms, and images of bank cards and people naked. The majority of buyers are reportedly unaware of where their recordings ultimately end up (Ebrahim, 2026).

Smart glasses also illustrate a harm that is not limited to the wearer. Individuals who are recorded by the glasses in private settings have not signed any contract with Meta and have not provided consent, creating an ethical and legal situation where third-party data collection obligations extend to bystanders who have no relationship to the platform. The capacity to record others passively, at scale, blurs the boundary between personal device data and environmental surveillance in ways consumer protection frameworks have not anticipated.

Mass Surveillance and the Sensor-Data-AI Pipeline

Taken individually, each component of the modern wearable ecosystem (device, sensor, companion app, cloud storage, AI models, and advertising networks) pose discrete but manageable privacy concerns. However, combined and projected forward, these form the infrastructure for mass surveillance capabilities that operate largely through ordinary commercial channels. The proliferation of devices capable of continuous physiological monitoring has enabled the capturing of vast volumes of personal health data and unprecedented scale of temporal data and inferences. Combining the nature of pervasive data collection, the opacity of which model inferences are made, and the risk of algorithmic discrimination calls to question the adequacy of existing regulatory frameworks (Radanliev, 2025).

Real-time sensor data when combined with audio, video, advertising identifiers, and data broker infrastructure creates the possibility of identifying or profiling individuals in public spaces (and potentially private spaces through use of wearables like Meta's smart glasses) in real time. Harvard students demonstrated in 2024 that Meta Ray-Ban glasses could be integrated with publicly available facial recognition tools to identify strangers on the street and surface their personal information. In October of 2025, a man in San Francisco was reported to have used Meta smart glasses to covertly record a woman at the University of San Francisco, which drew significant media attention. This is representative of a real-world manifestation of the surveillance potential that privacy researchers have warned about (ALM Corp, 2026). The

trajectory from passive health monitoring to ambient environmental surveillance represents a qualitative expansion of the wearable privacy problem that current law does not address.

AI Decision Making, Profiling, and Opaque Consequential Decisions

The final and perhaps most structural significant future concern revolves around the deepening integration of AI-generated health inferences into consequential decision making across commercial and institutional contexts. Advertisers are able to infer health status, political affiliation, and financial stress from ordinary behavior and utilize those data points for targeting, even without explicit, sensitive data fields. AI-driven targeting and optimization are integrated into every advertising channel with dynamic content and pricing tuned to real-time predicted user response, which amplifies existing concerns about discrimination, manipulation, and opaque decision making (Cheong, 2024).

As predictive models become more accurate and more pervasive, their generated health inferences are increasingly likely to not only be used in advertising, but also in pricing, underwriting, hiring, and access-to-service contexts. The consumer facing these decisions typically has no visibility into the chain from their wearable data to the model output to the consequential outcome. They do not know what was inferred, when, by whom, or on what basis. This is all currently absent of meaningful regulatory intervention. Consumers have no enforceable right to challenge, correct, or delete those inferences.

Consumer Harms Matrix

In this report's accompanying file, a constructed Consumer Harms matrix has been constructed based on the information above. This matrix details the harm to the consumer, the criticality, the type of provider conduct that enables that harm and what they do specifically, the affected population, laws applicable, and identifies gaps as they relate to frameworks.

Summary

Wearable health devices introduce a plethora of new issues that can lead to tangible harms for the consumer. There is no current active legislation that directly regulates the data protection practices, third-party data usage, and data collection practices of smartwatch companies. Additionally, these companies are not Covered Entities under HIPAA, meaning that consumers' data is offered very little protection. This paper has demonstrated the health inferences that

smartwatches are able to make about an individual. The technology can infer if a person is pregnant, their disability status, illnesses, and immunocompromizations. Aggregated data able to make predictive analyses leads to the potential for discriminatory issues and general privacy concerns, including employment and insurance discrimination and targeted advertisements.

It is essential that policymakers recognize that consumer smartwatches are not merely health devices, but a node in a broader data infrastructure that shapes individual behavior, economic outcomes, and democratic participation. The aggregation of biometric, behavioral, and inferred data into persistent personal profiles create leverage points that extend far beyond commercial transactions. They pervade political persuasion, community surveillance, and potential erosion of individual autonomy. Left unaddressed, these dynamics accelerate alongside the capabilities of the devices and the AI systems that process their outputs.

Proposed Policy Approaches to Enhance Consumer Privacy Protections

The following policy approaches are informed by analysis of smartwatch privacy policies, regulatory frameworks, and stakeholder input. This analysis identified several consistent patterns.

First, consumer-facing disclosures are difficult to interpret and compare. Privacy policies across major manufacturers vary in structure and terminology and are often written at a level exceeding general consumer comprehension. Research indicates that even when policies differ materially, those differences are not presented in a way that enables meaningful comparison or informed decision-making. As a result, the availability of information may not translate into practical usability.

Second, downstream data flows are often opaque. While manufacturers disclose categories of data sharing, the lack of standardized formats limits comparability, introduces ambiguity in data transferability usage, and constrains the ability of consumers, researchers, and regulators to assess data practices.

Third, existing enforcement mechanisms are primarily reactive. Under current authority, federal enforcement actions are typically initiated following complaints or investigations after potential consumer harm has occurred. This approach may limit oversight effectiveness given the scale and complexity of the data ecosystem.

In addition, consumer health data collected by smartwatches may fall outside comprehensive federal protections, creating a gap between the sensitivity of the data and the level of legal protection applied. As a result, firms may not face consistent economic incentives to adopt privacy-protective practices, and consumers may have limited ability to influence how their data are used.

A Privacy-Outcome Framework

In order to evaluate the efficacy of existing and proposed consumer privacy protections, this report establishes the following framework. It remains grounded in the provided definition of privacy, drawing from established consumer, government, and market-side compliance challenges, and in pursuit of improving consumer privacy. The four components of the framework include:

1. **Meaningful Comprehension**

This criterion begins by acknowledging the aforementioned push factors that strain proper consumer comprehension of privacy policies. These include reading opportunity costs, psychological reliance on cognitive heuristics in place of in-depth vetting, digital literacy deficits, variability of extracted meanings from privacy language, and a general consent fatigue arising from the abundance of policy documents. The resulting proposals are measured, consequently, based on their ability to address these consumer-side challenges or alleviate the downstream negative impacts that result.

2. **Actionable Control of Data Usage**

A natural extension of the first criterion, this criterion acknowledges the complexity of the data ecosystem and the participation of parties beyond the primary consumer-smartwatch provider relationship. In this layered series of transactions, control over usage, in the domain of privacy, necessitates transparent revelation by providers. This includes the corporate affiliates, service providers, and other partners with whom collected and derived data is made accessible to, and the conditions for such. Having actionable control of data allows for a traceable digital footprint, benefitting consumer privacy.

3. **Mitigation of Harmful Secondary Risks**

This criterion recognizes that the primary privacy risk in modern data ecosystems does not arise solely from collection, but from the expansion of meaning and use through secondary processing and inference. To this end, it evaluates whether privacy protections effectively limit, constrain, or govern the downstream use of collected data. In particular, this considers uses that extend beyond the original context of collection, including algorithmic inference, profiling, and cross-contextual data integration.

4. **Ability to Exist in the Data Ecosystem without Residual Exposure**

This criterion addresses the temporal dimension of privacy, recognizing that data often outlives the individual's relationship with the original collector. To evaluate this, the criterion looks at whether individuals can withdraw from data ecosystems and terminate the ongoing use, storage, and propagation of their data, including derived and shared data, without continued exposure or downstream persistence.

Policy Evaluation Criteria

Drawing from the analytical framework above, each proposal is evaluated against four criteria designed to capture not only its theoretical effectiveness in improving privacy outcomes, but also its practical viability within real-world institutional, market, and behavioral constraints.

1. **Institutional and Political Sustainability**

   This criterion evaluates whether the policy can be maintained over time given fiscal constraints, administrative capacity, and political turnover. This includes the durability of funding mechanisms, reliance on agency resources (e.g., FTC enforcement capacity), and the likelihood of continued political support across changing regulatory priorities. This encourages a consideration of sustainable adoption to generate lasting impact.

2. **Public Comprehension Capacity**

   This criterion assesses the extent to which the policy improves *collective* understanding of data practices, rather than relying on fully informed individual decision-making. Given well-documented limits to individual comprehension, this criterion emphasizes whether the policy enables intermediaries such as researchers, journalists, and advocacy groups to interpret, analyze, and disseminate information in ways that make privacy risks more legible at scale.

3. **Consumer Burden Minimization**

   This criterion measures how far the policy reduces the need for active, repeated, or technically complex decision-making by individual users. This includes reliance on opt-in/opt-out actions, exposure to consent fatigue, and the degree to which effective privacy protection depends on sustained user engagement. Policies that shift responsibility away from individuals toward structural protections score more favorably.

4. **Implementation Feasibility and Enforceability**

   This criterion evaluates the legal, technical, and administrative requirements needed to operationalize the policy. This includes whether new statutory authority is required, the complexity of compliance for firms, the clarity of enforcement mechanisms, and the

ability of regulators to verify and act on violations. Policies that depend on unverifiable disclosures or undefined institutional roles face lower feasibility.

Proposal 1: A Private Right of Action

The first proposal of this paper involves providing consumers with a Private Right of Action. In 2019, the GAO identified regulatory challenges faced by the FTC and Federal Communications Commission (FCC) in promulgating privacy-related trade regulation. As a result, they recommended expanding civic penalty authority to protect consumer privacy protections (Government Accountability Office, 2019b). Most notable was the enforcement burden on Section 5 of the FTC Act to encourage consumer privacy protections, which constantly faced industry pushback and judicial friction due to its narrow scope. Researchers Solove and Hartzog (2013) corroborate this observation, highlighting how despite significant enforcement actions and a robust privacy regulatory regime, this concentrated regulatory structure is structurally constrained in two ways: it is resource-limited relative to the scale of the data economy, and it is predominantly reactive, typically triggered by complaints, investigations, or ex post discovery of harm. As a result, the FTC's enforcement ability does not scale proportionally with the volume or complexity of data flows in the modern data economy. One outcome of this resource constraint is the strategic pursuit of cases based on success likelihood and where "companies have no viable defense." Choosing only specific cases will lead to consumer harm without retribution.

A Private Right of Action (PRA) would extend this enforcement framework by allowing individuals to bring civil claims against firms for violations of statutory privacy obligations (Oregon Consumer Justice, 2026). In the context of smartwatch ecosystems, such a mechanism would permit users and class representatives to seek remedies when biometric, behavioral, or inferred health data is collected, processed, or shared in ways that violate legal standards or materially diverge from disclosed practices. This parallel enforcement channel distributes monitoring capacity across private actors, in theory shifting the system from a low-frequency, high-cost enforcement model to a higher-frequency, decentralized one in which legal risk is continuously present across the ecosystem.

Market Misalignment & Structural Enforcement Limits

A central limitation of current enforcement frameworks is the misalignment between regulatory penalties and the economic value generated through consumer data. While agencies can impose fines and negotiate settlements, these penalties are episodic and may be insufficient to deter large-scale violations in data-intensive markets. This structural imbalance is visible in contemporary privacy enforcement regimes. For instance, since the implementation of the GDPR in 2018, aggregate fines have reached approximately €5.88 billion by the end of 2024 (Cannovale, 2025), alongside more than 2,200 enforcement actions and a reported increase in incidents of approximately 22 percent, or roughly 443 cases per day in 2025 (Broms & Lundin, 2025). Despite this volume of enforcement activity, regulators remain structurally outpaced by the scale and complexity of firms engaged in cross-border data collection and processing.

Equally pressing, even in regimes with comparatively aggressive penalty structures, enforcement does not reliably alter firm incentives when violations remain profitable at scale. High-profile penalties illustrate this constraint. In 2023, the Irish Data Protection Commission issued a €1.2 billion fine against Meta for unlawful data transfers to the United States, followed by an additional €390 million penalty for violations related to personalized advertising practices (Bracy, 2024; Chan, 2023). Collectively, these enforcement actions amounted to approximately $1.7 billion in regulatory sanctions. However, this figure represents a marginal proportion of Meta's overall financial performance, equating to approximately 0.96 percent of annual revenue and 3.32 percent of annual profit based on its 2023 SEC filings (SEC, 2023). In practice, revenue growth in its advertising business significantly outpaced the financial impact of these penalties, indicating that non-compliance remains economically absorbable within existing business models. The misalignment between meaningful fines and data profitability may limit how much financial penalties alone alter firm behavior.

State PRA Models: BIPA & CPRA

State-level privacy regimes illustrate two distinct models of private enforcement: a broad procedural rights model and a narrow harm-based model. The Illinois Biometric Information Privacy Act (BIPA) represents the more expansive approach. It authorizes private rights of action for any statutory violation, including failures to provide notice, obtain informed consent, or maintain a retention policy. Liability does not require proof of material harm; instead, standing is triggered by the violation of procedural privacy rights. Statutory damages are set at $1,000 for negligent violations and $5,000 for intentional or reckless violations, applied on a

per-violation basis (Rochford, 2023). This structure enables high-frequency litigation and may increase deterrence, but it also creates exposure to substantial aggregate liability.

By contrast, the California Privacy Rights Act (CPRA) adopts a narrower harm-contingent model. Codified in Section 16 of the proposition, the private right of action is restricted to specific data breaches involving personally identifiable information due to inadequate security safeguards (Text of Proposed Laws: Proposition 24, n.d.). Claims require demonstration of harm, particularly as a consequence of failed security protocols. Statutory damages range from $100 to $750 per consumer per incident, with judicial discretion based on severity and intent. Broader enforcement authority is retained by the California Privacy Protection Agency, with private litigation functioning as a more limited backstop.

These models reflect different enforcement approaches: BIPA emphasizes continuous private enforcement, while CPRA prioritizes regulator-led oversight with targeted private remedies.

Enforcement Design Tradeoffs in PRA Systems

Empirical experience from BIPA and CCPA highlights key design tradeoffs in structuring private enforcement. BIPA's broad scope and per-violation damages structure has been associated with strong deterrent effects, particularly for large firms. However, the framework has also generated concerns regarding litigation volume, proportionality, and uneven effects across firms based on capital. For instance, following the Illinois Supreme Court's decision in *Cothron v. White Castle System*, courts interpreted violations as accruing per scan, significantly increasing potential damages exposure and prompting subsequent legislative proposals to recalibrate liability and distinguish between procedural noncompliance and substantive misuse (Smigielski & Kantrow, 2022; Doran et al., 2024).

CCPA/CPRA, in contrast, incorporates procedural safeguards that limit litigation exposure, including a 30-day cure period under the agency's discretion (California Consumer Privacy Act of 2018, 2018) and phased-in requirements for independent cybersecurity audits for firms that process consumers' personal information in a manner that presents a significant risk to consumers' security (Faircloth et al., 2026). At the risk of constraining the breadth of private enforcement, these provisions likely reduce litigation volume and provide clearer compliance benchmarks.

Implications and Design Considerations for a Federal Private Right of Action

A federal PRA would likely need to balance these approaches by clearly defining actionable violations while avoiding excessive or indeterminate liability exposure. In the context of smartwatch and wearable health data ecosystems, enforceable violations may be most effective when limited to observable and verifiable conduct. This includes:

- Collection of biometric or health data without affirmative, informed consent
- Failure to execute user-directed deletion requests within a defined timeframe
- Secondary use of data beyond disclosed or authorized purposes
- Material misrepresentation of data practices in standardized disclosures
- Failure to implement reasonable security safeguards resulting in unauthorized access

Within this scope, a PRA could extend enforcement capacity while maintaining clarity around actionable conduct. Drawing from BIPA and CPRA, design considerations may also include:

- *A phased-in implementation*: Implemented in California, a tiered adoption process for firms based on revenue size can provide smaller firms with sufficient time to comply without taking on excessive financial costs.
- *Standing*: Given the sensitivity of biometric and physiological data, standing could be based on statutory violations rather than requiring proof of financial harm for core privacy breaches.
- *Tiered Liability*: Damage structures could be calibrated to reflect the severity of violations, for example, by distinguishing between technical, negligent, and knowing or reckless conduct. Such an approach may better align penalties with culpability and reduce the likelihood that minor procedural violations are penalized at levels comparable to substantive and purposeful misuse of sensitive data.
- *Litigation Limits:* Statutory damages could be subject to caps or other limiting mechanisms. Evidence from state-level regimes, particularly under BIPA, suggests that uncapped, per-violation damages can result in substantial aggregate liability, raising concerns regarding proportionality and financial exposure. Incorporating limits on damages may mitigate these risks while preserving deterrent effects for higher-severity violations. At the same time, broader private enforcement frameworks have been

associated with increased litigation activity. Depending on how liability is structured, this could affect both judicial capacity and firm behavior. For example, firms may allocate additional resources toward legal compliance and defense, although the extent to which this occurs, relative to investments in privacy-protective practices, is uncertain.

- *Contractual Constraints on Enforcement*: The effectiveness of a PRA also depends on the extent to which consumers can access courts in practice. Many technology firms, including smartwatch manufacturers such as Fitbit, incorporate mandatory arbitration clauses and class action waivers into their terms of service, requiring disputes to be resolved through individual arbitration rather than collective litigation (Monteith et al., 2023). These provisions limit the ability to aggregate claims, which is often necessary to make low-value but widespread harms economically viable to litigate. As a result, a PRA that does not address such contractual constraints may be formally available but practically limited in scope, particularly for diffuse harms such as unauthorized secondary use of biometric or behavioral data. Legislative design may therefore need to consider whether, and under what conditions, mandatory arbitration should be restricted or preempted for statutory privacy violations.

- *Safe harbors*: Firms demonstrating compliance through certified privacy or cybersecurity programs could receive limited protections from liability for technical violations. This positive incentive structure would encourage proactive audits by firms without requiring additional oversight by any particular federal agency.

- *Preemption*: Federal legislation establishing a PRA would need to address its interaction with existing state privacy regimes. In states such as California, dedicated regulatory bodies administer and enforce consumer privacy protections, including limited private rights of action. In these contexts, federal preemption could displace or alter established enforcement structures. A non-preemptive approach, in contrast, could allow state-level regimes to continue operating alongside a federal baseline. Under this structure, states may retain flexibility to implement and refine different enforcement models, including variations in standing requirements, damage structures, and procedural safeguards. Prior analyses of federal preemption frameworks have noted that preserving state authority in this manner can enable states to function as "laboratories" for policy development, including the adoption of more stringent notice requirements or enforcement mechanisms than those required under federal law (15 U.S.C. 1681t: How Federal

> Preemption Affects State Laws, 2025). In the context of privacy enforcement, such variation may provide an opportunity to observe the effects of different PRA designs over time, including their impact on litigation volume, deterrence, and compliance behavior. These observations could inform future federal adjustments without requiring comprehensive statutory redesign. However, a non-preemptive framework may also introduce complexity for firms operating across jurisdictions, particularly where state requirements diverge.

These considerations suggest that tiered liability structures involve a tradeoff between deterrence and predictability. Designs that expand access to damages may strengthen enforcement incentives but also introduce variability in litigation outcomes, while more constrained approaches may improve stability but limit the reach of private enforcement.

Evaluation Against Policy Criteria

PRA scores relatively strongly on **consumer burden minimization**, as it shifts a portion of enforcement responsibility away from individual users and toward institutional actors capable of operating at scale. Rather than requiring consumers to continuously monitor and respond to privacy disclosures, a PRA enables third parties to identify potential violations and pursue remedies. Additionally, reasonable fines for violations reinforce this shift by incentivizing organizations to proactively protect consumers' biometric data, rather than relying on individuals to detect and respond to harms.

It may also enhance **public comprehension capacity**, not by directly improving individual understanding, but by contributing to the production of publicly accessible legal records. Litigation filings, judicial decisions, and settlements can generate information that researchers, journalists, and advocacy organizations may use to interpret and communicate privacy risks to broader audiences.

However, its performance on **implementation feasibility and enforceability** is more constrained. Establishing a PRA would likely require federal legislation or modification of existing statutory frameworks, particularly to define actionable violations and address contractual mechanisms such as arbitration clauses. In addition, enforcement outcomes may depend on the incentives and resources of private litigants, which could result in uneven coverage across different types of violations.

Finally, **institutional and political sustainability** remains uncertain. Although a PRA does not require ongoing appropriations in the same manner as agency enforcement, it has historically generated debate regarding litigation volume, compliance costs for firms, and potential effects on innovation. Its long-term stability would therefore depend on legislative design choices that balance deterrence with predictability and administrative manageability.

Proposal 2: Expanded Transparency Practices

Current transparency approaches in the smartwatch data ecosystem rightfully focus on disclosure, but neglect comprehensibility, timeliness, and verifiability. This weakens their ability to be effective governance mechanisms, functioning instead as procedural formalities. This proposal bridges this accountability gap by introducing three mechanisms designed to address distinct structural failures:

1. Ineffective consumer comprehension at the point of collection
2. Misalignment between consent and data sensitivity over time
3. The absence of a verifiable, system-wide record of data practices.

Together, these mechanisms aim to convert transparency into an enforceable and auditable framework.

Failures of Current Transparency Frameworks

Empirical evidence indicates that current notice-and-consent models do not produce informed understanding. Fabian et al. (2017) analyzed approximately 50,000 privacy policies and found that most require at least a high school or college-level reading ability. Proctor, Ali, and Vu (2008) similarly found that policies are written at reading levels corresponding to approximately 13 years of education. Lastly, Ibdah et al. (2021) reported that 89 percent of participants experienced difficulty understanding privacy policies, while 55% derived incorrect conclusions from policy content. These findings indicate a clear gap between perceived and actual comprehension.

More significantly, comprehension failures are concentrated in the specific categories of disclosure most relevant to consumer decision making. Reidenberg et al. (2015) found that terms such as “service providers,” “business partners,” and “affiliates” are consistently misunderstood

across consumer populations, despite being legally sufficient. These categories describe the primary mechanisms through which data is shared, yet they are the least accurately interpreted. Evidence specific to wearable technologies reinforces these findings, as 76 percent of smartwatch manufacturers were found to exhibit high-risk transparency practices, the highest failure rate among evaluated dimensions (Van Deursen et al.,2025). These findings indicate that the issue extends beyond general policy complexity and into the specific structure and content of disclosures in the wearable ecosystem.

These findings establish a specific design problem that is distinct from comprehension failure described earlier in this report. The issue is not only that policies are long and complex, but that the specific language used to disclose sharing relationships is often unclear; this is the most consequential content for consumer decision making and the most systematically misunderstood. Structural redesign of that disclosure layer is necessary, rather than just simplification.

Two additional structural failures further limit the effectiveness of transparency. First is the static nature of disclosures. Privacy policies are typically presented at device activation and are not updated in a manner that ensures users are aware of material changes to data practices. When changes do occur, they are often communicated in the same format as the original policy, requiring consumers to recall prior terms, identify what has changed, and evaluate whether to continue using the product or accept the revised conditions. This process assumes a level of comprehension and continuity that most users do not possess, limiting the practical ability to make informed decisions. Second, transparency obligations do not extend to downstream data recipients. Analytics providers, advertising networks, and data brokers operate without independent disclosure requirements, limiting visibility into how data is used after initial collection. As a result, enforcement actions, such as those identified in prior FTC cases, are often reactive and dependent on external complaints rather than systematic monitoring. As illustrated in the BetterHelp case referenced in the *Targeted Manipulation* section, this gap allows harmful practices to persist until impacts accumulate at scale, with no structured record available to detect discrepancies between stated and actual behavior. The legal standard for disclosure remains adequacy of notice rather than verifiability of practice, which is an asymmetry that undermines accountability.

Limitations of Existing Legal Frameworks

Current U.S. legal frameworks do not address these structural transparency failures. Section 5 of the FTC Act enables enforcement against deceptive practices but does not require firms to produce standardized, verifiable records of data practices. Enforcement is therefore reactive and limited in scale. Prior findings indicate that the FTC relies heavily on complaint-driven investigations, which constrains its ability to monitor systemic issues.

Comparative frameworks demonstrate alternative approaches but also highlight gaps in U.S. law. The GDPR requires disclosure of processing purposes, legal bases, downstream recipients, and retention periods, and mandates records of processing activities. However, analysis of smartwatch privacy policies shows that these elements are often omitted, suffer from definitional variation or transfer location opacity, or are described in such broad terms as to limit their practical utility (Chan & Toh, 2025).

Domestic transparency initiatives such as Vermont's data broker registry and California's Delete Request and Opt-out Platform (DROP) have established useful precedents for public disclosure systems. Vermont's regisry confirms the value of the concept and its central structural limitation: the Vermont Attorney General has no mechanism to identify non-registrants, and coverage is incomplete because the model depends on self-registration without active monitoring authority (Vermont Office of Attorney General, 2018). California's DROP platform under the Delete Act presents the same limitation at a larger scale (California Privacy Protection Agency, 2024). Both precedents define what federal mechanism must do differently: mandatory filing obligations must be paired with active compliance monitoring, or coverage will be structurally incomplete.

Design Requirements for a Workable Transparency System

A workable transparency system for the wearable data ecosystem must accomplish three things current frameworks do not. These requirements flow directly from the analysis above and correspond to three stakeholder groups: consumers, civil society actors, and regulators. Each of these stakeholders have different requirements of transparency to fulfill their respective oversight functions.

Requirement 1: Consumer-facing disclosure must be legible at the point of collection and dynamic as practices change. The Reidenberg et al. (2015) finding establishes the specific design constraint: plain language is insufficient if it retains vague labels without proper explanation. The short-form disclosure layer must characterize sharing relationships with genuine specificity such as named categories of recipients, stated processing purpose, and conditions for further sharing. Using terms like "service providers" does not satisfy this requirement and leads to consumer misunderstandings.

Requirement 2: Consent must be calibrated to data sensitivity and delivered at the moment elevated-sensitivity data events occur, not only at enrollment. Gupta et al. (2023), in a nationally representative conjoint study of 3,539 U.S. adults, found that the purpose of data use was the single most important individual factor in health data sharing willingness, accounting for 29.9 percent of variance in the conjoint analysis. When all four privacy protections were examined (consent, transparency, oversight, and deletion) were considered together, they collectively accounted for 51.5 percent of variance in willingness to share. Additionally, Schaub, Balebako, Durity, and Cranor (2015) found that notices delivered at the moment of data collection are significantly more effective than enrollment-time notices. The design implication is that consumers make materially different decisions when they know specifically what is shared, with whom, and for what purpose the moment a specific data event occurs.

Requirement 3: A publicly accessible, machine-readable, structured record of data practices must exist. But no consumer-facing mechanism can by itself create an auditable institutional record. Without a structured public record, regulators must rely on complaints to trigger investigation, the PRA is difficult to prosecute because there is no baseline representation against which deviation can be measured, and consumers cannot verify whether what a short-form disclosure says matches what a manufacturer has formally committed to. THe technical foundation for this mechanism already exists: Buschhaus et al. (2025), working through the German government-funded InviDas project, developed a reusable conceptual model for smartwatch privacy policies validated with privacy law practitioners, demonstrating that encoding disclosure categories in machine-readable format enables automated comparison across manufacturers without requiring engagement with natural-language documents. The validated schema exists, but the federal mandate to file in that format does not.

Mechanism 1: Layered Policy Design with Dynamic Material Change Notification

This mechanism addresses Requirement 1. A two-tier disclosure structure is established by statute. The first tier is a mandatory short-form summary presented at device activation before setup proceeds, written at or below the 8th grade reading level that the Digital Medicine Society identifies as the minimum standard for health-related disclosures (Digital Medicine Society, 2026). It must cover four categories: the specific types of data collected, the identified or specific categories of downstream recipients and the purposes for which each processes the data, the consumer's rights and how to exercise them, and applicable retention periods. The full legal policy remains accessible as the underlying document. This layered approach is endorsed by the European Data Protection Board under GDPR Article 13 (GDPR, 2018) as balancing legal completeness with practical accessibility.

Reidenberg et al. (2015) findings are the binding document constraint: the short-form layer must not use categorical placeholders. Named categories of recipients, stated processing purposes, and conditions for further sharing are required fields in the short-form layer.

When a manufacturer makes a material change to data practices, such as new data-sharing relationship, new collection purpose for data, or change in retention periods, it must notify existing device users through a direct, plain-language alert distinct from a general terms-of-service update. The alert must specify what has changed, why, and what options the consumer has in response. The definition of "material change" must be established through regulatory rulemaking rather than left to manufacturer discretion: manufacturers have a structural incentive to define the category narrowly to minimize friction.

A limitation of this mechanism is that it does not guarantee that consumers will act on the information they receive. As Acquisti, Brandimarte, and Loewenstein (2015) established and in the Consent and Consumer Disempowerment section of this report confirms, there is a gap between stated privacy preferences and actual behavior. This mechanism creates the informational conditions for meaningful comprehension. It cannot compel corrective behavior.

Mechanism 2: Point-of-Risk Dynamic Consent

This mechanism addresses Requirement 2. A consumer who activates a smartwatch consents to general data collection at a moment when they cannot meaningfully anticipate what the device will subsequently detect. An ECG reading identifying an irregular cardiac rhythm, a

temperature and heart rate pattern from which fertility status can be inferred, a sleep profile consistent with a clinical condition are events of substantially elevated sensitivity that general enrollment consent does not adequately cover. The NIST Privacy Framework 1.1 identifies this mismatch between enrollment-time consent and event-time sensitivity as a structural gap in existing governance for data-intensive systems (NIST, 2025).
A tiered sensitivity classification, established through regulatory rulemaking, governs when targeted consent prompts are triggered:

Tier 1 covers routine biometric monitoring. This includes step count, general heart rate, caloric estimates, and proceeds under background consent without interruption. This is the core fitness tracking function most consumers understand they are authorized at activation.

Tier 2 covers location data, behavioral inferences from activity patterns, and transmission to third-party analytics partners for purposes beyond core service delivery. This tier generates periodic plain-language disclosure notices at the moment data is transmitted. Harkening back to Schaub et al. (2015), notices delivered at the point of data collection are more effective because they provide decision-relevant context at an actual decision point.

Tier 3 requires affirming consent before transmission. This covers data events from which from which the most consequential downstream harms documented in this report flow: cardiac anomaly detection, temperature and heart rate patterns from which fertility status or pregnancy probability may be inferred, oxygen saturation changes consistent with a clinical condition, and sleep profiles from which mental health indicators can be inferred. The Gupta et al. (2023) finding that states the purpose of use is the most important consideration for health data sharing variance provides empirical support for this tier and that consumers will make materially different sharing decisions when they know what is being shared, with whom, and for what purpose.

The principal implementation risk is notification fatigue. Schaub et al. (2015) document that over-frequent interruptions erode the protective function of consent as reflexive approval rates rise. This is why Tier 3 must be narrowly defined through rulemaking rather than manufacturer discretion. Manufacturers have a structural incentive to classify as many events as possible at lower tiers to minimize the friction. Regulatory tier definitions must be revisited on a defined schedule as inference technology advances. One limitation is that this mechanism creates a documented consent form at the highest-sensitivity tier, which has evidentiary value for private

litigation under the PRA proposal. It cannot prevent downstream misuse once consent is given, and it does not address data events that occur below the Tier 3 threshold.

Mechanism 3: Machine-Readable Disclosure Registry

This addresses Requirement 3. Congress establishes a mandatory public disclosure registry, administered by a designated federal authority with rulemaking and enforcement power. Manufacturers are required to file structured disclosures in a standardized, machine-readable format at product launch and update filings when practices change materially. Each filing must specify for each category of data collected: the legal basis for collection; the identity or specific category (including any data brokers, data aggregators, or reseller platforms to whom data is transmitted directly or though intermediary SDK or pipeline relationships) and the purpose for which that recipient processes the data; the application retention period; any use of the data for AI inference or profiling; and a deletion propagation specification identifying which downstream parties must honor deletion requests and within what timeframe. Cross-border transfer terms, destination countries, and the legal mechanism relied upon for each transfer must be required fields. Chan and Toh (2025) found these are precisely the elements that current smartwatch privacy policies omit.

The technical feasibility of this mechanism has been demonstrated by Buschhaus et al. (2025), whose InviDas platform shows that encoding standardized disclosure categories in machine-readable form enables automated comparisons across manufacturers.

For regulators, this registry enables automated compliance scanning identifying discrepancies between filed representations and observed practices without devoting investigative resources to individual firms. Violations become pattern-driven instead of compliant-driven. For civil society actors and plaintiffs' attorneys, each filing becomes a data point in a comparative body of evidence that can be scanned at market scale. For consumers, the registry would make it possible to verify whether what a short-form label or consent prompt represents matches what the manufacturer has formally committed to in a public record. This mechanism makes the other two enforceable: without an independent structured record, neither the short-form disclosure nor the event-specific consent prompt can be verified against actual practice. The DROP under the Delete Act provides the closest domestic institutional model for the oversight function this registry requires.

This mechanism has limitations. A registry built on manufacturer filings captures the original collection, but leaves meaningful portions of the downstream chain unverified. Software Development Kit (SDK)-level data flows, real-time bidding pipelines, and broker-to-broker transfers (described at length in Stage 5 technical section) are often probabilistic. Manufacturers may not have complete visibility into which downstream parties ultimately receive data transmitted through these channels. Vermont's registry demonstrates this limitation through its structural incompleteness without monitoring authority. A registry extended to processors and brokers would be more complete, but requires a standing federal agency with the authority to mandate filings, verify accuracy, and investigate discrepancies. These are open design questions that registry legislation must address.

Without a federal body holding the mandate to develop, standardize, and enforce these disclosure requirements, the registry remains institutionally unanchored.

How these Mechanisms Work Together

The three mechanisms are sequenced. The registry is the accountability layer that makes Mechanisms 1 and 2 enforceable. A consumer or enforcement actor can only verify that a layered disclosure or Tier 3 consent prompt accurately represents a firm's practices if there is an independent structured record to compare against. Without the registry, Mechanisms 1 and 2 are informational improvements. WIth it, they become enforceable representations

In context of the Privacy-Outcome Framework established earlier:

Meaningful Comprehension is addressed primarily by Mechanism 1, which delivers plain-language disclosure in the specific content area (sharing relationships) where comprehension failures are empirically most severe. Mechanism 3 is the verification layer that makes comprehension-based claims auditable.

Actionable Control of Data Usage is addressed primarily by Mechanism 2, which ensures that consumers can make contextually informed decisions about sharing their most sensitive data at the moment those specific data events occur. Mechanism 1's material change requirement also supports this by ensuring changes to practices trigger a fresh consumer decision.

Mitigation of Harmful Secondary Risks is primarily addressed by Mechanism 3, which creates a structured, auditable record of the downstream data relationships through which the most

consequential harms document in this report (health status inference, insurance underwriting, and behavioral profiling).

Ability to Exist in the Data Ecosystem Without Residual Exposure is addressed through the deletion-propagation requirement in Mechanism 3: manufacturer filings must specify which downstream parties bear a legal obligation to honor deletion requests and within what timeframe. This does not resolve the problem of inferences that survive deletion in AI model weights (a gap identified in the AI Inference section of this report), but creates a documented legal obligation and verifiable baseline against which compliance failures may be measured.

Proposal 3: Nutrition Label Model

As discussed in the preceding section, existing privacy policies are legally sufficient yet functionally ineffective for both consumers and oversight actors. A standardized privacy "nutrition label" model offers a complementary intervention by restructuring disclosure into a consistent, comparable, and interpretable format. Rather than expanding the volume of information available, this approach focuses on organizing high-salience data practices into a uniform schema that reduces cognitive burden and enables side-by-side comparison across devices. In doing so, it seeks to convert disclosure from a legal formality into a usable decision making and accountability tool, while creating a common interface through which consumers, regulators, and third parties can evaluate privacy practices at scale.

Historical Precedence and Inspiration

A government-regulated system of standardized labeling has precedent in several federal programs. Standardized labeling is seen in the ENERGY STAR, administered by the U.S. Environmental Protection Agency in coordination with the U.S. Department of Energy, and the Nutrition Facts label administered by the U.S. Food and Drug Administration and the U.S. Department of Agriculture. These programs were designed, in part, to improve consumers' ability to evaluate products by standardizing the presentation of key attributes, thereby facilitating product comparisons and influencing market incentives.

Established in 1992, ENERGY STAR is a voluntary public-private partnership that identifies and promotes energy-efficient products, buildings, and industrial practices to align economic growth with environmental protection considerations (Energy Star, 2024). The program relies on performance-based criteria and third-party certification through accredited

laboratories and certification bodies, allowing manufacturers flexibility in how efficiency standards are achieved. Its success has been attributed to a simple and recognizable labeling format, alignment with direct financial incentives through reduced energy costs, and periodic updates to performance criteria that encourage innovation. Importantly, research suggests that its effectiveness derives less from consumers' detailed understanding of energy metrics and more from the credibility and salience of a government-backed signal.

The Nutrition Facts label, adopted following the Nutrition Labeling and Education Act (NLEA) of 1990, similarly standardized complex information to address consumer confusion, rein in industry packaging variation, and improve comparability (Kim at al., 2000). Unlike the ENERGY STAR program, research evaluating the label's impact on consumer behavior and health outcomes has yielded mixed results. On the one hand, studies have found that a significant majority of adults, particularly those concerned about health risks, reported reading the Nutrition Facts label when purchasing packaged foods, that nutritional label use reduced intakes of total and saturated fat, cholesterol, and sodium, and that information about sugar and fiber has been used consistently to inform consumers choices (Bleich & Wolfson, 2015; Kim et al., 2000; Todd and Variyam 2008). On the other hand, comprehension has been found to vary across populations, with lower levels of nutrition literacy associated with reduced use and understanding of label information (Persoskie et al., 2017). Moreover, even when understood, nutritional information does not consistently translate into sustained changes in purchasing behavior or overall diet quality, as factors such as price, access, taste preferences, and convenience continue to play a significant role in food selection (Drichoutis et al., 2009). Nonetheless, on the industry side, the disclosure requirements have been associated with some degree of product reformulation, as manufacturers have adjusted nutrient profiles particularly for nutrients such as fat, sodium, and, in later regulatory updates, trans fat, in response to increased transparency and evolving regulatory requirements (Rahkovsky et al., 2012).

More recent efforts, including broadband consumer labels (Federal Communications Commission, 2022) and the Cyber Trust Mark (Federal Communications Commission, 2024) developed by the FTC, reinforce both the potential and the limitations of standardized disclosure. These formatted labels improve consumer comparability and pricing transparency, but their effectiveness is constrained by technical complexity and variation in user comprehension (Choy et al., 2024). Taken together, these programs highlight several consistent features of effective disclosure systems:

1. **Format and content standardization is foundational.**

The presentation of information in a consistent, recognizable structure across products and manufacturers reduces cognitive burden by allowing consumers to develop familiarity with where specific information is located and how it is expressed. Standardization also enables direct comparison across competing products, which is essential for markets to reward higher-performing or safer options.

2. **Simplification is paired with credibility.**

The attached authority of a trusted government-backed framework functions as a heuristic signal that reduces the need for consumers to interpret detailed efficiency metrics, suggesting that comprehension gains are driven both by trust and signal clarity as by the underlying data itself.

3. **Layered information architectures balance accessibility and completeness.**

The Cyber Trust Mark and ENERGY STAR programs combine a high-level, easily interpretable front-facing label with access to more detailed information. This layered approach accommodates both low- and high-engagement users without overwhelming either group.

4. **Alignment with incentives amplifies effectiveness.**

These programs are most effective when paired with tangible incentives, such as reduced energy costs in the case of ENERGY STAR (2024) or potential liability and reputational considerations in cybersecurity labeling. On the supply side, standardized disclosure can influence firm behavior by making product attributes more visible and comparable, thereby encouraging competition on those dimensions.

5. **Government oversight and periodic updating sustain relevance.**

All three programs rely on ongoing regulatory or administrative involvement to define criteria, verify compliance, and update standards over time. This ensures that labels remain meaningful as technologies and market practices evolve, and mitigates the risk of stagnation or strategic manipulation by firms.

These findings suggest that disclosure systems are most effective not when consumers fully understand underlying technical details, but when they can rely on clear, consistent, and

credible signals to guide decision making (McDonald et al., 2009). This insight is directly applicable to privacy disclosures for smartwatches, where technical complexity and information asymmetries are particularly pronounced.

Consumer Data Privacy and Security Label Proposal

This report proposes a similar standardized consumer data privacy and security label for smartwatches and similar wearable devices. Prior research on privacy labels provides both validation and design constraints. Kelley et al. (2009) found that a label implementing visual aides like a structured grid, colors, simplified symbols, a legend, and placement formatting improved users' ability to interpret and compare privacy practices relative to natural-language policies. The information presented included the types of personal data collected, how data is used, whether data was shared with third-parties, and user control options and retention practices. As Figure 2 indicates, the researchers adopted a consistent labeling format including data categories in rows regardless of whether the company collected them or not, in response to user confusion over comparability. Study groups were found to be more accurately able to discern privacy policies relative to natural language policies, reported better navigation, had an easier time understanding complex language, could better compare company policies and make personalized decisions better, and experienced less information overload. The study participants also voiced higher levels of enjoyability in reading the privacy policies in this format.

Figure 2: Early Iteration of a Proposed Privacy Nutrition Label, (Kelley et al., 2009)

A recurring challenge with this single format was balancing information loss and overload. Input from privacy and experts from industry, academia, government, and nongovernmental organizations found the information presented to be insufficient to allow for effective scrutiny by consumers and consumer advocacy groups (Emami-Naeini et al., 2021). This generalization undermined the initial goal of allowing consumers to better differentiate privacy policies and hold IoT companies accountable to the adoption of privacy enhancing technologies.

To address these challenges, this report proposes a two-layer architecture for structured layering and definitional support, in effect balancing accessibility with completeness, while enabling both consumer use and external scrutiny. Emami-Naeini et al. (2021) designs the primary layer as "the concise format of the label that can be printed and attached to a product

package. The secondary layer is an online format of the label and can be accessed through the primary layer, by scanning a QR code or typing in the uniform resource locator (URL) that is placed at the bottom of the primary layer." This layered approach had the advantage of appealing to varying degrees of consumer comprehension. As Figure 3 indicates, the first layer, much like the 2009 iteration, provides immediate consumers with a snapshot understanding of the type of data collected, its storage process, with whom it is shared, and to whom it is sold.

Figure 3. Primary Layer of a Proposed Data Privacy Label (Emami-Naeini et al., 2021)

The second layer (Figure 4), accessible through an embedded link in the first layer, provides expanded, structured, and interactive information for users and oversight actors seeking a deeper understanding. With regard to actionable information, this layer empowers consumers,

consumer advocacy groups, tech journalists, and other stakeholders to more closely evaluate tech policies by including information such as security protocols, data retention periods, data storage facilities, privacy framework compliance, data linkage and inference strategies, and contact information.

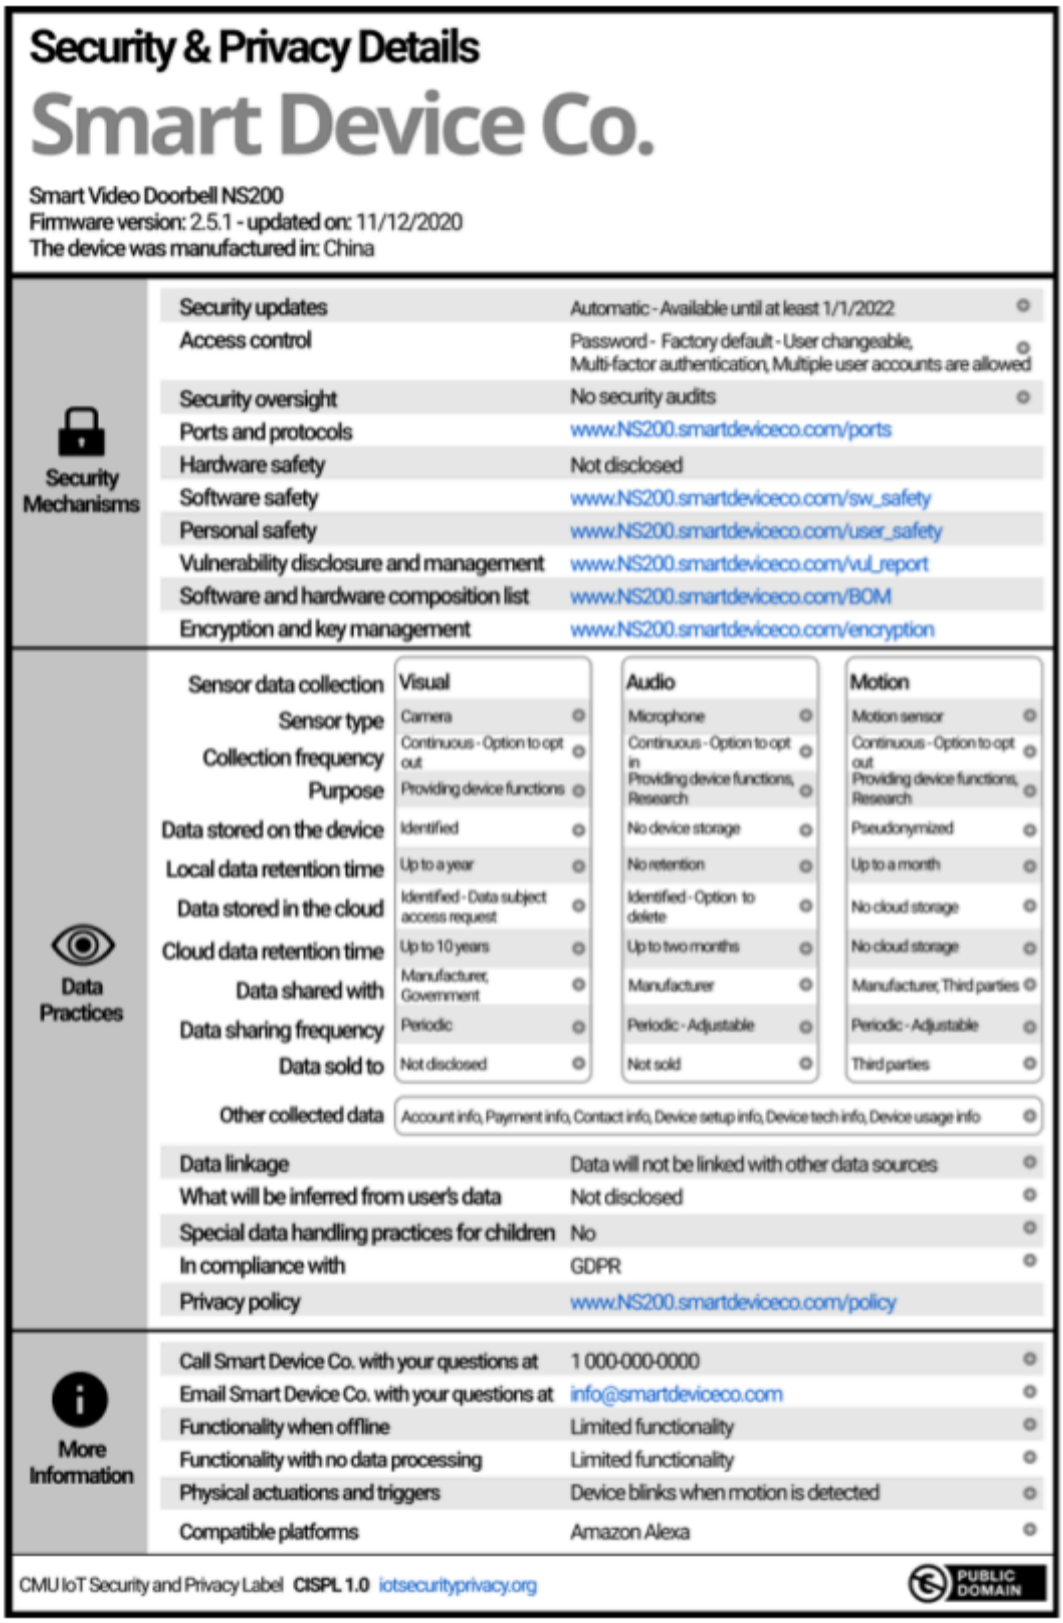

Figure 4: Secondary Layer of a Proposed Data Privacy Label (Emami-Naeini et al., 2021)

The authors find that the secondary layer significantly improves users' ability to evaluate devices, particularly by increasing sensitivity to high-risk practices such as third-party data sharing and extended cloud retention periods. These features were shown to meaningfully influence perceptions of device risk, suggesting that more detailed, structured disclosures can shape consumer judgment when presented in an accessible format. They attribute these improvements to several design elements. First, embedded links allow users to access the full natural-language policy when additional detail is required, preserving completeness without overloading the primary interface. Second, interactive features such as hover-based definitions reduce ambiguity around technical terminology and support incremental learning. Third, the inclusion of recommended privacy and security practices provides normative context, helping users interpret whether a given practice is typical or potentially problematic. Finally, optional provider-supplied considerations offer space for additional clarification without disrupting standardization.

Taken together, a minimum label schema should include the following high-salience features in its primary layer, presented in a consistent grid format:

1. Data types collected: Categories such as biometric, location, behavioral, and device usage data.
2. Data sharing practices: Whether data is shared with affiliates, third parties, or sold.
3. Data retention: Standardized ranges (e.g., <30 days, 30–180 days, >1 year).
4. Security practices: Presence of baseline protections (e.g., encryption, regular security updates).
5. User controls: Availability of data access, deletion, and opt-out mechanisms.

The secondary layer should provide detailed and interactive information for users seeking deeper understanding such as:

1. Data inference and profiling practices.
2. Storage location (on-device vs. cloud).
3. Specific security protocols and certifications.
4. Data retention policies by category.
5. Third-party data recipients and purposes.
6. Compliance with recognized frameworks (e.g., NIST standards).
7. Contact information and links to full privacy policies.

8. Embedded definition links for technical terms.

Governance and Enforcement

Effective implementation of a standardized privacy labeling system requires clearly defined institutional roles, credible verification mechanisms, and an adoption strategy that balances feasibility with long-term regulatory consistency.

Standard Setting

The technical standards for label content, format, and terminology proposed above, drawn from privacy researchers like Emami-Naeini et al. (2020), should be developed in coordination with NIST standards. Leveraging existing resources such as the NIST Privacy Framework and IoT security baselines would ensure technical rigor and adaptability. Standard-setting should define data categories, establish uniform terminology, specify formatting requirements, and include mechanisms for periodic updates.

Regulatory Enforcement

The FTC should serve as the primary enforcement authority, in part due to its existing powers to investigate and prosecute inaccurate, incomplete, or misleading privacy labels as violations. Enforcement responsibilities would include monitoring compliance with labeling requirements across manufacturers and platforms, investigating discrepancies between disclosed and actual data practices, and providing interpretive guidance to clarify expectations for industry. This approach would also increase administrative feasibility since it would build upon existing FTC authority rather than requiring the creation of a new enforcement body. While the FTC would serve as the primary enforcement body, coordination with the FCC may be appropriate for devices operating within communications ecosystems, particularly where smartwatch labeling overlaps with connectivity standards or IoT security initiatives such as the Cyber Trust Mark.

Certification and Verification

A system of third-party certification, modeled after ENERGY STAR, could support scalability and credibility. Independent, accredited auditors could verify that disclosed practices align with actual behavior, expanding oversight capacity beyond the federal government, and adding reputational incentives for firms to maintain compliance. This independent verification process could also strengthen consumer trust in the label.

Adoption Pathway

A phased approach may improve institutional and political sustainability. Initial voluntary adoption could help establish credibility and allow refinement based on real-world feedback. This could then expand into mandatory standardization to ensure full comparability and prevent selective participation by firms with weaker privacy practices. This phased approach mirrors the trajectory of programs like ENERGY STAR where voluntary adoption established credibility before broader normalization.

Benefits of Standardized Privacy Labels

A standardized privacy label offers several advantages within the report's evaluation framework.

1. *Improved public comprehension capacity*: Research shows that consumers are more likely to engage with disclosures when information is presented in structured, simplified formats with predictable organization. A standardized label reduces navigational complexity and cognitive load by presenting key privacy attributes in a consistent visual layout. This also supports consumer burden minimization by allowing users to develop familiarity with where specific information is located and how it is represented. This reduces the effort required to interpret disclosures over time and supports faster, more intuitive evaluations. This correlates with improved accuracy in identifying data practices and lowered perceived information overload relative to traditional policies.

2. *Product comparability:* The minimization of information heterogeneity enables consumers to differentiate products based on privacy practices in addition to other factors.

3. *Enhanced market transparency and firm accountability*: By requiring disclosure of specific categories of data practices, labels constrain the ability of firms to obscure or selectively present information, especially sensitive information such as third-party data sharing and retention periods. Over time, this may incentivize firms to compete on privacy-enhancing features, particularly in markets where core product differentiation is limited. Even in the absence of regulation mandating its adoption, a government-backed standardized privacy label can create a market incentive for self enrollment, much like the ENERGY STAR program, helping to achieve similar aims.

4. *Aligned policy and regulatory oversight:* A consistent disclosure framework provides regulators with a clearer basis for monitoring compliance, identifying misleading claims, and enforcing transparency requirements. This aligns with emerging legislative proposals, such as the American Privacy Act, which emphasize data minimization, transparency, and user control. A labeling system can operationalize these principles in a format that is directly accessible to consumers while remaining auditable by regulators.

Implementation Considerations and Tradeoffs

The effectiveness of standardized privacy labels depends on careful design and institutional support, leading to several tradeoffs. First, there is an inherent tension between simplification and completeness. Labels that are too simplified risk omitting important nuances, while overly detailed labels may overwhelm users. The two-layer architecture largely mitigates but does not eliminate this possibility.

Second, there is a weak correlation between improved access to information and comprehension or behavioral change (Kelley et al., 2009). At the same time, labels work best for stakeholders who are aware of the value of such information. As such, label design must incorporate not only standardized content but also standardized definitions and contextual cues. Interactive elements in the secondary layer, such as hover explanations, can partially address this issue, but their effectiveness depends on user engagement.

Third, standardized systems may create incentives for firms to optimize for the label rather than underlying privacy outcomes. Practices may be structured to appear favorable within the labeling framework without materially reducing risk. Having certification programs and third-party vetting partners can help address this, although this introduces scope complexity for the agency tasked with overseeing this process.

Lastly, in connection with the previous consideration, implementing a standardized labeling system requires coordination across multiple stakeholders, including regulators, manufacturers, standards bodies, and third-party certifiers. Expanding the role of federal agencies such as the FTC or FCC to set specification demands may present administrative and political challenges. In addition, ensuring consistent enforcement across a rapidly evolving technology landscape will require sustained institutional capacity. Integrating labeling requirements into broader privacy regulation could reduce fragmentation and improve coherence across policy instruments.

Proposal 4: Pay-for-Privacy (PFP) Model

The Pay-for-Privacy (PFP) model, also known as a consent-or-pay model, offers consumers privacy protections and stratified personal data collection and monetization based on price thresholds. It emerged primarily in response to the consent requirements outlined in the GDPR as a "legally defensible way to obtain the consent required…for behavioural advertising (Team, 2026). In theory, this approach reconciles privacy preferences with existing advertising-based business models.

This report *does not* recommend the adoption of PFP as a mechanism for strengthening consumer data protections in the smartwatch ecosystem. Emerging evidence suggests that it is structurally ill-suited to produce meaningful privacy outcomes, shifting both the financial burden and decision making responsibility onto consumers often in ways that constrain rather than expand genuine choice. Moreover, empirical implementations indicate that PFP models frequently function as mechanisms to preserve data collection practices under a different pricing structure, rather than to reduce them. Given these dynamics, along with concerns related to equity, coercion, and enforceability, the PFP model is better understood as a limited and potentially counterproductive approach within the context of sensitive health and biometric data. The analysis that follows evaluates its design, real-world performance, and key limitations to clarify why it is not recommended as part of a comprehensive privacy policy framework.

Conditions for Viability

The PFP model was developed in digital markets heavily reliant on user data for advertising revenue (Skovgaards, 2025). Depending on the implementation, the model offers smartwatch customers with a spectrum of free, discounted / low-cost, or paid services commensurate with their data collection and monetization preferences. Under ideal conditions, this structure allows firms to maintain revenue while offering privacy-conscious consumers an alternative. To function effectively, several conditions must hold:

- The pricing features appropriately approximate the economic value of user data.
- Consumers understand the implications of data collection and can make informed tradeoffs.
- Core product functionality must remain equivalent across paid and unpaid tiers, such as Meta's "Less Personalized Ads" (LPA), that satisfies end users and business revenue

- The choice between paying and consenting is meaningfully voluntary (as opposed to economically coerced).

In practice, these conditions are difficult to satisfy simultaneously. In addition, there is limited applicability of this model in smartwatch ecosystems. PFP has been tested largely in advertising-driven digital markets, whereas the smartwatch market has complex data flows and payment dependencies whose differentiation makes translation ineffective. Nonetheless, the rest of this section explores each condition in turn and provides concluding remarks on transferability conditions.

Condition 1: Reasonable Pricing and Valuation of Personal Data

Given the intent of the PFP model, a necessary condition to align payment and consent is that the privacy fee must reasonably approximate the economic value of the user's data. In competitive markets such as news publishing, ad-free pricing is typically based on estimated advertising revenue per user, resulting in relatively low monthly costs between €1 and €5 per user per month (Team, 2026). In contrast, large platforms such as Meta derive significantly higher value from user data due to cross-platform integration (e.g., Facebook, Instagram, WhatsApp), advanced behavioral profiling, and high advertiser demand for targeted audiences. This leads to substantially higher privacy fees, often justified using average revenue per user metrics in the range of €9.99 to €12.99 (Meta, 2024).

However, in addition to data value, pricing is also shaped by market power, network effects, and the availability of substitutes. Regulatory analysis by the European Data Protection Board (EDPB) has found that dominant platforms like Meta can set higher privacy fees without losing users, effectively influencing consent decisions through economic pressure (Post, n.d.). To counteract this, the EDPB legally differentiated standard publishers and "Large Online Platforms" (LOPs), ruling that large tech platforms could not use the same revenue pricing justifications as standard publishers given their market concentration (European Data Protection Board, 2024). This dynamic raises concerns about whether PFP models can produce fair or non-coercive outcomes, particularly when users lack viable alternatives.

In the smartwatch context, aligning price with data value is also less applicable because advertising is not the dominant revenue model. Unlike advertising-dependent platforms, smartwatch manufacturers primarily generate revenue through hardware sales, subscription

services, and ecosystem integration. In addition, while data may be monetized through analytics, partnerships, or secondary uses, it is not typically priced per user in a transparent or linear way. This absence of a standardized per-user baseline figure to anchor pricing and the challenge of valuing marginal data value increase with aggregation makes any "privacy fee" in this context likely arbitrary or strategically set.

Condition 2: Consumer Understanding and Ability for Informed Tradeoffs

The PFP model relies on consumers to understand and evaluate the implications of consenting to data collection versus paying for privacy. Researcher Elvy (2017) posits that in complex digital environments, consumers often lack either a sufficient understanding of data practices or a sufficient alternative to consent, effectively limiting their privacy-related decisions-making. In addition, she details the resulting unequal access to privacy, the illusory feeling of control and choice even as ISPs collect user data for other purposes, and the documented instances of predatory and discriminatory behavior against users opting for free data-intensive options. These realities suggest broader concerns of a socioeconomic digital divide among the financially capable and disadvantaged that might arise in the propagation of this model.

These limitations are amplified in the smartwatch ecosystem, where the data flows are integral to device functionality. The process of paying for privacy introduces difficulties that strain how core features and device performance are implemented, creating ambiguity around what consumers are actually purchasing. When layered with previous findings of data transmission to third parties absent clear disclosure among certain manufactures, in addition to the variation in privacy practices across manufacturers and ecosystems, the task of defining a consistent privacy tier grows all the more challenging. Without standardized definitions of privacy protections (e.g., limits on retention, inference, or third-party sharing), PFP models risk offering inconsistent or ill-defined privacy tiers, reducing their reliability as a policy mechanism.

Condition 3: Functionality Differentiation

For a PFP model to operate effectively, the core functionality of a product or service must remain substantively equivalent across privacy tiers. Consumers should not be required to trade essential features, usability, or service quality in exchange for enhanced privacy protections. If

functionality differs meaningfully, the model ceases to represent a true privacy choice and instead becomes a bundled tradeoff between privacy and product performance.

Empirical implementations suggest that this condition is difficult to satisfy in practice. Rather than offering consistent services with varying levels of data protection, PFP models frequently bundle privacy with different underlying products. In digital publishing markets, for example, consumers paying a PFP fee are typically not purchasing complete privacy. Instead, they are buying a narrower restriction on specific practices, such as third-party tracking. First-party data collection often continues for operational or analytics purposes, and users may still receive contextual advertising.

By contrast, platform-based implementations offer a different form of differentiation. When Meta introduced its PFP model, the paid tier removed advertisements entirely. However, this does not eliminate data collection; the platform continues to rely on extensive first-party data to support content personalization, algorithmic ranking, and internal analytics (Mousavi et al., 2026). As a result, the paid product is not simply a "more private" version of the same service, but a different configuration of features and data practices. Across these cases, the nature of what consumers are purchasing varies substantially, limiting comparability and undermining the premise of a standardized privacy choice.

Regulators have attempted to address this gap by encouraging the development of alternatives that better preserve functional equivalence. One such example is the three-tier model introduced by French publisher *Le Monde* in line with France's data protection authority (Team, 2026):

- **Tier 1:** Consent (free). Full access with personalised advertising. Requires compliant consent collection with genuine granularity.
- **Tier 2:** Contextual (free, reduced). Access to a set number of articles per month with contextual advertising only. No tracking consent required.
- **Tier 3**: Pay PFP fee. Full access, no advertising, no third-party tracking. Price set to reflect the actual advertising revenue value foregone.

While this contextual advertising approach represents an improvement over binary consent-or-pay structures (Fishley & Kite, 2024), it also highlights the complexity of maintaining

functional equivalence. The contextual tier often includes reduced access to content or features, indicating that even "improved" models may still rely on differentiated service levels. This suggests that maintaining full functionality while varying only privacy protections remains difficult in practice.

In the context of smartwatches, these challenges are further amplified. Unlike web-based services, limiting data collection on wearable devices that rely on continuous data collection may directly degrade device performance or disable key features altogether. As a result, any attempt to implement a PFP model would likely require users to choose between full device functionality and enhanced privacy protections. This tight coupling between data collection and functionality makes it particularly difficult to design privacy tiers that are meaningfully equivalent.

Condition 4: Voluntariness of Choice

An analysis of regulatory experience in the European Union and the United Kingdom, where pay-for-privacy models have been implemented across news publishing and platform markets, highlights persistent challenges in presenting consumers with genuinely voluntary, non-coercive choices between paying and consenting.

Evidence from News Publishing

Evidence from digital publishing markets illustrates how the introduction of a monetized privacy option alters user behavior. In 2023, the United Kingdom required that "organisations using non-essential cookies (such as analytics, performance or marketing cookies) on their website or a mobile app must ask users whether they permit the operator to use such cookies before placing the cookie on the user's device" (Yaros & Hepworth, 2023). Initially, consent rates ranged between 70–80 percent, indicating that a meaningful share of users opted into data collection when no financial tradeoff was introduced (Tobitt, 2025). In reaction to declined revenue (Skovgaards, 2025), publishers introduced PFP models that conditioned privacy on payment. After implementation, acceptance rates increased to near-universal consent, with only a small minority of users choosing to pay for a privacy-preserving alternative (Tobitt, 2025).

This shift suggests that the introduction of a financial barrier materially changes user decision making. Rather than reflecting underlying preferences, observed consent behavior appears to be driven by cost avoidance and convenience, raising concerns that the choice is not meaningfully voluntary. Industry stakeholders have acknowledged that these models are designed to restore advertising revenue, with high consent rates reestablishing the value of personalized data while payments from a small subset of users offset residual losses (Skovgaards, 2025).

While some perspectives suggest that these models may evolve toward greater transparency and more balanced user experiences over time (Elvy, 2017), further evidence reinforces the interpretation that the presence of a paywall systematically nudges users toward consent. For instance in markets with higher privacy awareness such as Germany and France, users were more likely to reject tracking when no monetization layer was present, underscoring the extent to which pricing structures shape outcomes (Tobitt, 2025).

Evidence from Platform Markets

A similar pattern emerged when Meta introduced a PFP model in Europe (Wang, 2025). Despite facing a paid alternative, the overwhelming majority of users continued to consent to data collection. The E.U.'s regulatory body subsequently determined that the model failed to provide a "genuine alternative," particularly given the lack of equivalent free options with reduced data processing. It instead posited that a truly equivalent alternative would be a free version with less data-intensive advertising (e.g. contextual instead of behavioral advertising), allowing for the ad-supported business model while avoiding problematic data processing (Schmalenberger & Heywood, 2025).

The result was Meta's rollout of an alternative "Less Personalized Ads" (LPA) that used 90 percent less data than the personalized ad model, servings ads based only on contextual signals like the specific content being viewed during a session, combined with basic demographic data such as a user's age, gender, and general location. From the business perspective, Meta appealed, citing a reported 800 percent increase in users closing ads because they are repetitive or irrelevant, a 70 percent decrease in advertiser conversion rates, and a 61 percent reduction in offsite conversions. This tension implies that meaningful alternatives may reduce the effectiveness of data-driven business models, creating incentives for firms to structure choices in ways that steer users toward consent. In addition, the legal challenges of enforcing 'equivalent alternatives' in data-intensive ecosystems placed a severe burden on the European

Commission. As such, adoption of this model would likely demand jurisdictional expansion and greater resource diversion for oversight, straining the feasibility of proper adoption

In the smartwatch context, the conditions for meaningful voluntariness are even more constrained. Unlike web-based services, smartwatches operate within tightly integrated hardware and software ecosystems, where switching costs are high and functionality is closely tied to data collection. Users may therefore face compounded pressure to consent, as opting out could involve both financial cost and degraded device utility. Taken together, these dynamics suggest that PFP models are unlikely to produce genuinely voluntary choices in wearable technology environments and may instead reinforce default patterns of consent.

Conclusion

Across all four conditions – pricing, comprehension, functionality, and voluntariness – evidence indicates that PFP models face significant structural limitations. In the context of smartwatches, where data collection is continuous, sensitive, and tightly coupled to functionality, these challenges are further amplified. As such, the PFP model is unlikely to provide a reliable or equitable mechanism for strengthening consumer privacy protections and is not recommended as part of a comprehensive policy framework.

Alignment with Evaluation Criteria

To support a structured assessment, the following table presents a comparative summary of how each proposal performs against the report's core evaluation criteria

Policy Proposal	**Institutional & Political Sustainability**	**Public Comprehension Capacity**	**Consumer Burden Minimization**	**Implementation Feasibility & Enforceability**
Private Right of Action				

Transparency Mandate				
Data Privacy "Nutrition" Label				
Pay-for-Privacy Model				

Lessons Learned

This project identified several cross-cutting insights relevant to the evaluation and governance of privacy risks in the smartwatch data ecosystem. First, the analysis reinforces longstanding limitations of the notice-and-consent model. Even where disclosures are legally sufficient, consumers often lack the capacity to interpret complex data practices or to make informed tradeoffs at the point of decision. Improvements to transparency may enhance comprehension but do not, on their own, resolve underlying structural constraints on consumer choice.

Second, existing legal and regulatory frameworks are not well aligned with the technical realities of modern data systems. While current approaches focus primarily on data collection and sharing, many of the most consequential risks arise from downstream aggregation, inference, and machine learning processes. This creates gaps in oversight where sensitive information can be derived without direct collection or explicit disclosure.

Third, the smartwatch ecosystem highlights a fundamental tension between functionality and privacy. Continuous data collection is integral to device performance, limiting the feasibility of models that rely on restricting data flows without affecting core features. As a result, policy approaches that assume clear separation between service provision and data collection may be difficult to implement in practice.

Fourth, transparency alone does not ensure accountability. While improved disclosures may increase visibility into data practices, they do not create mechanisms for verification or

enforcement, particularly in the context of complex downstream data-sharing relationships. Without structured and accessible records of data practices, both regulators and external stakeholders face challenges in identifying discrepancies between stated and actual behavior.

Finally, the findings suggest that effective privacy protections may require a shift from consumer-centric to system-level approaches. Reliance on individual decision making places significant cognitive and financial burdens on users, while offering limited protection in complex data environments. More durable solutions may involve embedding privacy protections within system design, strengthening institutional oversight, and standardizing data governance practices across the ecosystem.

Recommendations for Future Work

1. Expansion Towards a Standardized Privacy Scoring System

While standardized privacy labels improve the accessibility and comparability of disclosures, they remain primarily descriptive, requiring consumers and intermediaries to interpret multiple attributes simultaneously. A scoring system could further reduce cognitive burden by providing an aggregated assessment of privacy impact. This natural extension would function similarly to the "percent daily value" framework used in nutritional labeling, where consumers have a reference point against which to assess whether a product's data practices fall within, exceed, or fall short of recommended privacy standards.

However, the analogy to nutritional labeling highlights both the promise and the difficulty of this approach. In food policy, raw ingredient lists provide limited value to most consumers because they require domain knowledge to interpret. The introduction of standardized benchmarks, such as the 2,000-calorie daily reference used by the U.S. Food and Drug Administration, enables consumers to contextualize information, have a normative baseline against which individual products can be evaluated, and ultimately make meaningful comparisons (*The Lows and Highs of Percent Daily Value on the Nutrition Facts Label*, 2023).

The primary challenge in developing such a system lies in calibration. Unlike nutritional guidelines, which are grounded in biomedical consensus, privacy lacks universally agreed

quantitative thresholds. Determining how to weight different dimensions of data practice, and how to translate them into composite scores, would require normative judgments that may vary across contexts and stakeholders. As a result, any scoring framework would need to be developed through a transparent, multi-stakeholder process and subject to periodic revision as technologies and social expectations evolve. The following foundational components would be worthwhile considerations:

1. **A Defined Baseline** (**"Recommended Privacy Practices"**)
   A meaningful scoring system would need to be anchored to a clearly defined standard of what constitutes acceptable or optimal data practices. This could draw from existing frameworks, including the Factor Analysis of Information Risk (FAIR) and NIST's Cyber Risk Management Program (CRMP), which provide commentary on data minimization principles, purpose limitation, retention limits, and restrictions on secondary use.
2. **A Transparent and Weighted Scoring Methodology**
   Privacy risk is multi-dimensional. Any composite score would need to account for factors such as:
   - Sensitivity of data collected (e.g., biometric vs. non-sensitive data)
   - Scope and purpose of data use
   - Extent of third party sharing
   - Retention duration
   - Strength of security safeguards
   - Presence of user controls (access, deletion, opt-out)

   These dimensions would need to be weighted according to their relative contribution to consumer harm. For example, unauthorized secondary use of biometric data may warrant significantly greater weight than extended retention of non-sensitive usage data. The weighting process itself would require empirical validation and normative justification.

This proposed expansion would need to consider measurement validity. Unlike calories, which have a physical basis, privacy risk is probabilistic and context-dependent. The same data practice may pose different risks depending on downstream use, aggregation, or inference capabilities. This makes it difficult to define universally "correct" scores. In addition, scores would need to handle important tradeoffs. For instance, a device with strong security but extensive data sharing could receive a similar score to one with minimal sharing but weak safeguards. Consumers may interpret these as equivalent despite materially different risk

profiles. Lastly, there is heterogeneity in consumer preferences, with some prioritizing functionality or personalization over strict data minimization. A universal score may not reflect these preference differences unless paired with more granular disclosures.

Despite these challenges, a scoring system has the potential to significantly enhance the usability of privacy labels by converting complex disclosures into interpretable signals. If successfully implemented, it could further align market incentives by making privacy performance not only visible, but directly comparable across product

2. Paid Privacy-Enhancing Technologies (PETs) as an Alternative Model

A potential refinement of the Pay-for-Privacy (PFP) concept involves shifting from a "pay to avoid data collection" framework to a model in which consumers pay for enhanced privacy-preserving technologies. Rather than conditioning privacy on the absence of data collection, this approach would allow data processing to continue while incorporating technical safeguards such as end-to-end encryption, federated learning, or homomorphic encryption.

Under this model, smartwatch manufacturers could offer baseline functionality with standard data processing practices, while providing a premium tier that incorporates stronger privacy protections. For example, sensitive biometric data could be processed locally on-device using federated learning, encrypted during transmission and storage, or analyzed using privacy-preserving computation techniques that limit raw data exposure. In theory, this allows firms to maintain data-driven functionality while reducing privacy risks for users willing to pay for enhanced protections.

This approach addresses several limitations associated with traditional PFP models. First, it avoids directly framing privacy as a tradeoff against access or affordability. Users are not required to "pay to avoid surveillance," but instead can opt into higher levels of protection. Second, it may better align with the technical realities of smartwatch ecosystems, where continuous data collection is often necessary for core functionality. Rather than attempting to eliminate data flows, the model seeks to secure them.

However, this approach introduces its own set of challenges. From an equity perspective, it retains elements of a "privacy as a premium feature" model. While less coercive than consent-or-pay structures, it still risks creating differential levels of protection across socioeconomic groups. As noted by Stacy-Ann Elvy (2017), models that tie privacy protections to

financial capacity can contribute to unequal access, particularly when consumers lack the resources or knowledge to evaluate these tradeoffs. Even when baseline protections are present, the availability of enhanced safeguards exclusively through paid tiers may reinforce disparities in exposure to privacy risks.

From a technical and governance standpoint, verifying the effectiveness of privacy-enhancing technologies presents additional complexity. Unlike binary claims such as "no data sharing," mechanisms such as federated learning or homomorphic encryption are difficult for consumers, and in some cases regulators, to independently assess. This raises concerns about transparency, standardization, and the potential for misleading claims regarding the level of protection provided. Finally, the model does not fully resolve issues related to user comprehension. Even if advanced protections are available, consumers must still understand what they are purchasing.

Taken together, while a PET-based pricing model represents a more technically grounded and less coercive alternative to traditional PFP frameworks, it does not eliminate core concerns related to equity, transparency, and consumer understanding. As such, it may be better considered as a complementary innovation within a broader regulatory framework that ensures baseline privacy protections for all users, rather than as a standalone solution.

3. Privacy-Preserving System Design as a Baseline Standard

A third area for future work involves shifting the focus of privacy protections from disclosure and user choice toward the underlying technical architecture of data systems. Emerging privacy-enhancing approaches, including federated learning, differential privacy, homomorphic encryption, and secure computation, demonstrate that it is increasingly possible to extract value from data while limiting direct exposure of raw user information (Thapa & Camtepe, 2021; Zhao et al., 2025). While these techniques vary in implementation and maturity, they share a common design objective: enabling functionality without requiring centralized access to identifiable data.

Rather than recommending specific technologies, future research and policy development could focus on defining a set of baseline design principles for privacy-preserving systems in the wearable technology ecosystem. These principles would emphasize minimizing data visibility, localizing sensitive processing where feasible, and reducing the extent to which raw data must be transmitted, stored, or accessed by manufacturers and third parties. This framing

reconceptualizes privacy as an embedded property of system design, an approach particularly relevant in the context of AI and machine learning systems, where risks increasingly arise not only from data collection, but from the ability to extract sensitive information from model outputs and parameters. As documented in prior sections, techniques such as membership inference, model inversion, and reconstruction attacks demonstrate that even indirect access to trained models can expose underlying training data. At the same time, efforts to improve model transparency through explainability methods may introduce additional privacy risks, creating inherent tradeoffs between interpretability and data protection.

A design-oriented framework would therefore encourage the development of systems that limit both direct and indirect data exposure across the full data lifecycle, including model training and deployment. This could include prioritizing architectures that (1) reduce reliance on centralized data aggregation, (2) incorporate formal privacy guarantees where feasible, and (3) account for downstream risks associated with inference and model leakage.

However, this approach presents several challenges. First, the effectiveness of privacy-preserving system designs can be difficult to evaluate, particularly as attack methods evolve. Second, these approaches may introduce tradeoffs in computational efficiency, model performance, or system complexity. Third, without standardized definitions or verification mechanisms, claims regarding “privacy-preserving” design may be difficult for regulators or consumers to assess.

Despite these limitations, advancing privacy as a system-level design objective represents a meaningful shift from reactive to preventative protection. Future work could explore how such principles might be formalized into technical standards, procurement requirements, or regulatory guidance, particularly for data-intensive consumer technologies such as smartwatches. In this way, privacy protections would be less dependent on user decision making and more consistently embedded within the technologies themselves.

Conclusion

This report examined the privacy and data protection landscape surrounding consumer smartwatches, with a focus on how existing legal, regulatory, and technical frameworks govern the collection, processing, and downstream use of sensitive biometric and behavioral data. Across the data lifecycle, the analysis finds that current protections are fragmented and often misaligned with the realities of continuous, inference-driven data processing in wearable ecosystems. While existing frameworks such as federal consumer protection law, sector-specific health privacy rules, and voluntary risk management standards provide partial safeguards, they do not collectively address the full scope of data generation, sharing, and AI-driven inference enabled by modern smartwatch platforms.

The four policy proposals developed in this report – private right of action, expanded transparency infrastructure, a standardized privacy label, and a critical evaluation of pay-for-privacy models – collectively highlight both the potential and limitations of current regulatory approaches. Taken together, they emphasize that meaningful consumer protection cannot rely on disclosure or consent mechanisms alone, particularly in environments where data collection is continuous, technically complex, and deeply embedded in product functionality.

Across all proposals, a consistent finding emerges: effective governance in the smartwatch ecosystem depends on shifting from static, document-based privacy compliance toward systems that are verifiable, structured, and responsive to real-time data practices. This includes improving the legibility of disclosures, strengthening accountability for downstream data flows, and developing mechanisms that allow oversight entities to independently assess whether stated practices align with actual behavior.

At the same time, the analysis underscores that no single policy intervention is sufficient on its own. The effectiveness of any one mechanism depends on complementary reforms in adjacent areas, including enforcement capacity, technical standardization, and consumer-facing usability. In particular, the emergence of AI-driven inference systems introduces new challenges that extend beyond traditional notions of data collection, requiring regulatory approaches that account for derived data, model behavior, and probabilistic re-identification risks.

Ultimately, the findings of this report suggest that smartwatch data governance is at an inflection point. As wearable devices become increasingly central to health monitoring and

daily decision making, the gap between technological capability and regulatory oversight continues to widen, further complicated by the growing patchwork of state privacy laws. Addressing this gap will require coordinated action across policymakers, regulators, and industry stakeholders to ensure that innovation in wearable technology does not outpace the development of meaningful and enforceable privacy protections.

Definitions

To ensure understanding and consistency within this report's documents, this paper has constructed definitions for common words used. The words defined in this report shall have the following meaning when found, to wit:

"Privacy" is the functional ability to either maintain control over, or possess timely, intelligible, and meaningfully enforceable knowledge of, the use, collection, dissemination, accumulation, storage, and disclosure of personally identifiable information, linked information, or inferred information concerning an individual (Solove, 2006, *GSA Rules of Behavior for Handling Personally Identifiable Information (PII)*, 2019; Moore, 2003) .

"Smartwatch" is a wearable Class B digital device intended to be worn on the wrist, is marketed as a general wellness product, and incorporates programmable computing capability, a user interface, wireless connectivity, and the capacity to record biometric information (*47 CFR Part 15 -- Radio Frequency Devices*, n.d.; *General Wellness: Policy for Low Risk Devices - Guidance for Industry and Food and Drug Administration Staff*, 2016; *Policy Statement of the Federal Trade Commission on Biometric Information and Section 5 of the Federal Trade Commission Act 1*, 2023) .

"Artificial Intelligence" is any computer system(s) that performs tasks that typically require human intelligence, such as reasoning, learning, pattern recognition, and decision making.

"Personal data" is understood as defined in the GDPR Chapter (1) Article (4) (1)

"Third party" is understood as defined in the GDPR Chapter (1) Article (4) (10)

"Biometric data" is understood as defined in the GDPR Chapter (1) Article (4) (14)

References

47 CFR Part 15 -- Radio Frequency Devices. (n.d.). Code of Federal Regulations. https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-15

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, *347*(6221), 509–514. https://doi.org/10.1126/science.aaa1465

All About Lawyer. (2026, February 26). *Apple 95 Million Lawsuit, Apple Siri Privacy Settlement Explained Payments & Deadlines*. All about Lawyer. https://allaboutlawyer.com/apple-95-million-lawsuit-apple-siri-privacy-settlement-explained-payments-deadlines/

ALM Corp. (2026, March 9). *Meta Ray-Ban AI Smart Glasses Face UK Investigation and US Lawsuit: What 7 Million Users Need to Know About the Privacy Scandal*. ALM Corp. https://almcorp.com/blog/meta-ray-ban-ai-smart-glasses-privacy-investigation-uk-us/

Apple. (2021). *Measuring Walking Quality Through iPhone Mobility Metrics*. https://www.apple.com/lae/healthcare/docs/site/Measuring_Walking_Quality_Through_iPhone_Mobility_Metrics.pdf

Bélanger, F., & Crossler, R. E. (2011). Privacy in the Digital Age: a Review of Information Privacy Research in Information Systems. *MIS Quarterly*, *35*(4), 1017–1041. https://doi.org/10.2307/41409971

Bernard, S. (2024). *Artificial Intelligence and Health Data: a promising combination with potential privacy pitfalls*. PRINEOS S.r.l. https://prineos.com/en/blog/artificial-intelligence-and-health-data/

Bleich, S. N., & Wolfson, J. A. (2015). Differences in consumer use of food labels by weight loss strategies and demographic characteristics. *BMC Public Health*, *15*(1). https://doi.org/10.1186/s12889-015-2651-z

Bracy, J. (2024). *Meta fined GDPR-record 1.2 billion euros in data transfer case*. Iapp.org. https://iapp.org/news/a/meta-fined-gdpr-record-1-2-billion-euros-in-data-transfer-case

Brennan Center. (2026, February 18). *Congress Must Close Data Broker Loophole by third-partyProhibiting Government Purchases of Americans' Sensitive Data*. Brennan Center for Justice. https://www.brennancenter.org/our-work/research-reports/congress-must-close-data-broker-loophole-prohibiting-government-0

Broms, A. J., & Lundin, G. (2025). *Number of personal data breaches in Europe increased by 22 per cent in 2025*. DLA Piper. https://sweden.dlapiper.com/en/news/number-personal-data-breaches-europe-increased-22-cent-2025

Buschhaus, C., Butting, A., Michael, J., Nitsch, V., Pütz, S., Rumpe, B., Stellmacher, C., & Theis, S. (2025). Overcoming the hurdle of legal expertise: A reusable model for smartwatch privacy policies. *Data & Knowledge Engineering*, *159*, 102443. https://doi.org/10.1016/j.datak.2025.102443

California Privacy Protection Agency. (2018a). *About DROP and the Delete Act*. Privacy.ca.gov. https://privacy.ca.gov/drop/about-drop-and-the-delete-act/

California Privacy Protection Agency. (2018b). *California Consumer Privacy Act of 2018*. https://cppa.ca.gov/regulations/pdf/ccpa_statute.pdf

California Privacy Protection Agency (CPPA). (n.d.). Cppa.ca.gov. https://cppa.ca.gov/

Cannovale, M. (2025, January 23). *GDPR fines across Europe total €1.2 billion in 2024, according to DLA Piper*. Legalcommunitygermany.com. https://legalcommunitygermany.com/gdpr-fines-across-europe-total-e1-2-billion-in-2024-according-to-dla-piper/

Castro, D., Dascoli, L., Diebold, G., & January. (2022). *The Looming Cost of a Patchwork of State Privacy Laws*. https://www.congress.gov/118/meeting/house/115376/documents/HHRG-118-IF17-20230301-SD021.pdf

Chan, H. Y., & Toh, H. J. (2025). A document analysis of international data transfer terms in smartwatch privacy policies: inadequacies and recommendations for improvements. *International Review of Law, Computers & Technology*, 1–22. https://doi.org/10.1080/13600869.2025.2510440

Chan, K. (2023, January 4). *European Union fines Facebook parent Meta 390M euros for privacy violations*. PBS NewsHour. https://www.pbs.org/newshour/world/european-union-fines-facebook-parent-meta-390m-euros-for-privacy-violations

Chase, S. E., Liddell, R. G., McGonagle, C. L., & Ives, S. J. (2024). Accuracy of the Apple Watch Series 9 for Measures of Energy Expenditure and Heart Rate at Rest and During Exercise: Impact of Skin Pigmentation. *Journal of Functional Morphology and Kinesiology*, *9*(4), 275–275. https://doi.org/10.3390/jfmk9040275

Chen, C. C., Shu, D., Ravishankar, H., Li, X., Agarwal, Y., & Cranor, L. F. (2024). *Is a Trustmark and QR Code Enough? The Effect of IoT Security and Privacy Label*

Information Complexity on Consumer Comprehension and Behavior. https://doi.org/10.1145/3613904.3642011

Cheong, B. C. (2024). Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making. *Frontiers in Human Dynamics*, *6*. https://doi.org/10.3389/fhumd.2024.1421273

Chico, M. J. (2025, March 13). *Interpretation of BIPA Amendment Splits the Courts: Implications for Businesses*. The Firewall. https://www.thefirewall-blog.com/2025/03/interpretation-of-bipa-amendment-splits-the-courts-implications-for-businesses/

Choy, C., Young, E., Li, M., Cranor, L. F., & Peha, J. M. (2024). Consumer-driven design and evaluation of broadband labels. *Telecommunications Policy*, *48*(5), 102717. https://doi.org/10.1016/j.telpol.2024.102717

Clements, R. (2018). *Evaluating the Costs and Benefits of a Smart Contract Blockchain Framework for Credit Default Swaps*. William & Mary Law School Scholarship Repository. https://scholarship.law.wm.edu/wmblr/vol10/iss2/3/

Code of Federal Regulations. (2024, May 30). *Part 318—Health Breach Notification Rule*. Code of Federal Regulations; National Archives. https://www.ecfr.gov/current/title-16/chapter-I/subchapter-C/part-318

Comparison of US State Privacy Laws Defining Covered and Sensitive Data. (2025). In *Information Policy Centre*. https://www.informationpolicycentre.com/wp-content/uploads/2025/11/cipl_comparison_us_state_privacy_laws_nov25-5.pdf

Congress. (2021, March 1). *S.500 - SMARTWATCH Data Act*. Www.congress.gov; Congress. https://www.congress.gov/bill/117th-congress/senate-bill/500

Congress. (2025, November 4). *Text - S.3097 - 119th Congress (2025-2026): Health Information Privacy Reform Act*. Congress. https://www.congress.gov/bill/119th-congress/senate-bill/3097/text

Consumer Privacy Act. (n.d.). *Colorado Privacy Act (CPA)*. Consumer Privacy Act. https://www.consumerprivacyact.com/colorado-privacy-act-cpa/

Coordes, L. (2021). *Articles Unmasking the Consumer Privacy Ombudsman*. https://scholarworks.umt.edu/cgi/viewcontent.cgi?article=2473&context=mlr

Cornell Law. (2019). *11 U.S. Code § 363 - Use, sale, or lease of property*. Legal Information Institute. https://www.law.cornell.edu/uscode/text/11/363

Cranor, L. F. (2025). *Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice*. Colorado Law Scholarly Commons. https://scholar.law.colorado.edu/ctlj/vol10/iss2/7/

Digital Medicine Society. (2026). *The DiMe Seal*. Dimeseal.org. https://dimeseal.org/

Doran, B., Valdetero, J., & Pestine, Z. (2024). *BIPA update: Illinois limits liability and clarifies electronic consent for biometric data collection*. Gtlaw.com. https://www.gtlaw.com/en/insights/2024/8/bipa-update-illinois-limits-liability-and-clarifies-electronic-consent-for-biometric-data-collection

Drichoutis, A. C., Nayga, Jr., R. M., & Lazaridis, P. (2009). Can Nutritional Label Use Influence Body Weight Outcomes? *Kyklos*, *62*(4), 500–525. https://doi.org/10.1111/j.1467-6435.2009.00448.x

Drye, K. (2011, May 13). *Disney's Playdom Charged with Violating Children's Online Privacy*. Kelley Drye & Warren LLP. https://www.kelleydrye.com/viewpoints/blogs/ad-law-access/disneys-playdom-charged-with-violating-childrens-online-privacy-enters-3-million-ftc-settlement

Duha Ibdah, M., Raparthi, A., & Bacha. (2021). *"Why Should I Read the Privacy Policy, I Just Need the Service": A Study on Attitudes and Perceptions Towards Privacy Policies*. https://anysbacha.github.io/publications/ibdah_access21.pdf

Dusheck, J. (2018). *Wearable sensors can tell when you are getting sick*. Medical Center Development. https://med.stanford.edu/medicalgiving/news/wearable-sensors-can-tell-when-you-are-getting-sick.html

Ebrahim, Z. (2026, March 20). *Meta Ray-Ban Smart Glasses Are Sending Footage of People Having Sex, Undressing, and Using the Bathroom to Employees*. The Hearty Soul. https://theheartysoul.com/meta-ray-ban-glasses-recording-private-situations/

Electronic Code of Federal Regulations. (2025, April 22). *Part 312— Children's Online Privacy Protection Rule (COPPA Rule)*. Code of Federal Regulations; National Archives. https://www.ecfr.gov/current/title-16/chapter-I/subchapter-C/part-312

Electronic Frontier Foundation. (2026). *The Law and Medical Privacy*. Electronic Frontier Foundation. https://www.eff.org/issues/law-and-medical-privacy#main-content

Electronic Privacy Information Center. (2025). *What is Real Time Bidding?* Electronic Privacy Information Center. https://epic.org/what-is-real-time-bidding/

Elvy, S.-A. (2017). Paying for privacy and the personal data economy. *Colombia Law Review*, *117*(6), 1369–1460. HeinOnline.

https://heinonline.org/HOL/Page?handle=hein.journals/clr117&id=1437&collection=journals&index=

Emami-Naeini, P., Agarwal, Y., & Cranor, L. (2020). *Specification for CMU IoT Security and Privacy Label (CISPL 1.0)*. https://www.iotsecurityprivacy.org/downloads/Privacy_and_Security_Specifications.pdf

Emami-Naeini, P., Dheenadhayalan, J., Agarwal, Y., & Cranor, L. F. (2021). An Informative Security and Privacy "Nutrition" Label for Internet of Things Devices. *IEEE Security & Privacy*, *20*(2), 31–39. https://doi.org/10.1109/msec.2021.3132398

Energy Star. (2024). Energy Star Annual Overview Report. In *Energy Star*. United States Environmental Protection Agency. https://www.energystar.gov/sites/default/files/2024-05/Annual%20Overview%20Report_2023.pdf

European Data Protection Board. (2024). *Opinion 08/2024 on Valid Consent in the Context of Consent or Pay Models Implemented by Large Online Platforms*. European Data Protection Board. https://www.edpb.europa.eu/system/files/2024-04/edpb_opinion_202408_consentorpay_en.pdf

Evans, G. J. (2015). How state laws can shore up the FTC'S authority to regulate data breaches, privacy, and more. *Administrative Law Review*, *67*(1), 187–219. JSTOR. https://doi.org/10.2307/26424827

Fabian, B., Ermakova, T., & Lentz, T. (2017). Large-scale readability analysis of privacy policies. *Proceedings of the International Conference on Web Intelligence*. https://doi.org/10.1145/3106426.3106427

Fair, L. (2021, January 13). *Health app broke its privacy promises by disclosing intimate details about users*. Federal Trade Commission. https://www.ftc.gov/business-guidance/blog/2021/01/health-app-broke-its-privacy-promises-disclosing-intimate-details-about-users

Fair, L. (2022, December 14). *Shedding Light on Dark Patterns: Protecting Consumers from Digital Deception*. National Association of Attorneys General. https://www.naag.org/attorney-general-journal/shedding-light-on-dark-patterns-protecting-consumers-from-digital-deception/

Faircloth, F., McNicholas, E., Fasone, B., & Varon, R. (2026, January 15). *California's CCPA Cybersecurity Audit Rule Takes Effect: What Businesses Need to Know*. Ropesgray.com. https://www.ropesgray.com/en/insights/alerts/2026/01/californias-ccpa-cybersecurity-audit-rule-takes-effect-what-businesses-need-to-know

Federal Communications Commission. (2022, February 23). *Broadband Consumer Labels*. https://www.fcc.gov/broadbandlabels

Federal Communications Commission. (2024). *U.S. Cyber Trust Mark*. Fcc.gov; Federal Communications Commission. https://www.fcc.gov/CyberTrustMark

Federal Reserve. (n.d.). Federal Trade Commission Act Section 5: Unfair or Deceptive Acts or Practices Background. In *Federal Reserve*. Retrieved April 5, 2026, from https://www.federalreserve.gov/boarddocs/supmanual/cch/200806/ftca.pdf

Federal Trade Commission. (2000, July 21). *FTC Announces Settlement With Bankrupt Website, Toysmart.com, Regarding Alleged Privacy Policy Violations*. Federal Trade Commission. https://www.ftc.gov/news-events/news/press-releases/2000/07/ftc-announces-settlement-bankrupt-website-toysmartcom-regarding-alleged-privacy-policy-violations

Federal Trade Commission. (2013, July 19). *Federal Trade Commission Act*. Federal Trade Commission. https://www.ftc.gov/legal-library/browse/statutes/federal-trade-commission-act

Federal Trade Commission. (2014). *Data Brokers: A Call for Transparency and Accountability*. https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf

Federal Trade Commission. (2022, September 15). *FTC Report Shows Rise in Sophisticated Dark Patterns Designed to Trick and Trap Consumers*. Federal Trade Commission. https://www.ftc.gov/news-events/news/press-releases/2022/09/ftc-report-shows-rise-sophisticated-dark-patterns-designed-trick-trap-consumers

Federal Trade Commission. (2023a, March 2). *FTC to Ban BetterHelp from Revealing Consumers' Data, Including Sensitive Mental Health Information, to Facebook and Others for Targeted Advertising*. Federal Trade Commission. https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-ban-betterhelp-revealing-consumers-data-including-sensitive-mental-health-information-facebook

Federal Trade Commission. (2023b, May 31). *FTC and DOJ Charge Amazon with Violating Children's Privacy Law by Keeping Kids' Alexa Voice Recordings Forever and Undermining Parents' Deletion Requests*. Federal Trade Commission. https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever

Federal Trade Commission. (2024, May 6). *BetterHelp Customers Will Begin Receiving Notices About Refunds Related to a 2023 Privacy Settlement with FTC*. Federal Trade Commission.

https://www.ftc.gov/news-events/news/press-releases/2024/05/betterhelp-customers-will-begin-receiving-notices-about-refunds-related-2023-privacy-settlement-ftc

Ferguson, M. (2025, May). *Meta's Controversial Data Policy on Ray-Ban Smart Glasses Sparks Privacy Debate*. Meta's Controversial Data Policy on Ray-Ban Smart Glasses Sparks Privacy Debate; OpenTools. https://opentools.ai/news/metas-controversial-data-policy-on-ray-ban-smart-glasses-sparks-privacy-debate

Financial Accounting Standards Board. (2021, March). *Accounting Alternative for Evaluating Triggering Events*. Fasb.org; Financial Accounting Standard Board. https://www.fasb.org/page/ShowPdf?path=ASU_2021-03.pdf&title=ACCOUNTING%20STANDARDS%20UPDATE%202021-03%E2%80%94INTANGIBLES%E2%80%94GOODWILL%20AND%20OTHER%20%28TOPIC%20350%29:%20ACCOUNTING%20ALTERNATI

Fishley, B., & Kite, C. (2024). *Consent or pay: One rule for some (Large Online Platforms), another rule for everyone else?* https://www.weil.com/-/media/mailings/2024/q3/consent-or-pay_one-rule-for-some_large-online-platforms_another-rule-for-everyone-else---august-2024.pdf

Fiveable. (2025). *Targeted advertising and manipulation*. Fiveable. https://fiveable.me/digital-ethics-and-privacy-in-business/unit-5/targeted-advertising-manipulation/study-guide/erL4S1ZLPaZRasYN

Food and Drug Administration. (2023, October 4). *The Lows and Highs of Percent Daily Value on the Nutrition Facts Label*. U.S. Food and Drug Administration.

https://www.fda.gov/food/nutrition-facts-label/lows-and-highs-percent-daily-value-nutrition-facts-label

GDPR. (2018). *General Data Protection Regulation (GDPR)*. GDPR; European Union. https://gdpr-info.eu/

General Wellness: Policy for Low Risk Devices - Guidance for Industry and Food and Drug Administration Staff. (2016). U.S. Food and Drug Administration. https://www.fda.gov/media/90652/download

Gilbert, S., Baca-Motes, K., Giorgio Quer, Wiedermann, M., & Brockmann, D. (2024). Citizen data sovereignty is key to wearables and wellness data reuse for the common good. *Npj Digital Medicine*, *7*(1). https://doi.org/10.1038/s41746-024-01004-z

Google. (2024). *Fitbit Privacy Policy - Help*. Google.com. https://support.google.com/product-documentation/answer/14815921

Government Accountability Office. (2019a). *Consumer Privacy: Changes to Legal Framework Needed to Address Gaps*. Www.gao.gov; Government Accountability Office. https://www.gao.gov/products/gao-19-621t

Government Accountability Office. (2019b). *Internet Privacy: Additional Federal Authority Could Enhance Consumer Protection and Provide Flexibility*. Government Accountability Office. https://www.gao.gov/assets/gao-19-52.pdf. GAO-19-52.

Government Accountability Office. (2022, January 5). *Credit Scoring Alternatives for Those Without Credit*. Www.gao.gov. https://www.gao.gov/blog/credit-scoring-alternatives-those-without-credit

GSA rules of behavior for handling Personally Identifiable Information (PII). (2019, October 8). U.S. General Services Administration.

https://www.gsa.gov/directives-library/gsa-rules-of-behavior-for-handling-personally-identifiable-information-pii-2

Gupta, R., Iyengar, R., Sharma, M., Cannuscio, C. C., Merchant, R. M., Asch, D. A., Mitra, N., & Grande, D. (2023). Consumer Views on Privacy Protections and Sharing of Personal Digital Health Information. *JAMA Network Open*, *6*(3), e231305. https://doi.org/10.1001/jamanetworkopen.2023.1305

Harvard Law Review. (2025, March 10). *Data Privacy in Bankruptcy: The Consumer Privacy Ombudsman*. Harvard Law Review. https://harvardlawreview.org/print/vol-138/data-privacy-in-bankruptcy-the-consumer-privacy-ombudsman/

IAPP. (2025, July 8). *The digital body: Rethinking privacy and security in wearable health trackers*. Iapp.org. https://iapp.org/news/a/the-digital-body-rethinking-privacy-and-security-in-wearable-health-trackers

IAPP. (2026, April 6). *US State Privacy Legislation Tracker*. Iapp.org. https://iapp.org/resources/article/us-state-privacy-legislation-tracker

Identity Theft Resource Center. (2025, January 27). *ITRC Annual Data Breach Report*. ITRC. https://www.idtheftcenter.org/publication/2024-data-breach-report/

Kelley, P. G., Bresee, J., Cranor, L. F., & Reeder, R. W. (2009). A “nutrition label” for privacy. *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS ’09*, 1–12. https://doi.org/10.1145/1572532.1572538

Kelly, J. (2025). *From Copyright Case to AI Data Crisis: How The New York Times v. OpenAI Reshapes Companies’ Data Governance and eDiscovery Strategy*. Nelson Mullins Riley

& Scarborough LLP. https://www.nelsonmullins.com/insights/blogs/corporate-governance-insights/all/from-copyright-case-to-ai-data-crisis-how-the-new-york-times-v-openai-reshapes-companies-data-governance-and-ediscovery-strategy

Kim, S., Nayga, R. M., & Capps, O. (2000). The Effect of Food Label Use on Nutrient Intakes: An Endogenous Switching Regression Analysis. *Journal of Agricultural and Resource Economics*, *25*(1), 215–231. JSTOR. https://doi.org/10.2307/40987057

Konnoth, C. (2024). AI and data protection law in health. *Edward Elgar Publishing EBooks*, 111–129. https://doi.org/10.4337/9781802205657.ch07

LegalClarity Team. (2025, March 29). *15 U.S.C. 1681t: How Federal Preemption Affects State Laws*. LegalClarity. https://legalclarity.org/15-u-s-c-1681t-how-federal-preemption-affects-state-laws/

Lomas, E. (2025, February 3). *Oura's Ovulation Detection Algorithm Outperforms Calendar Tracking Method in New Validation Study*. The Pulse Blog. https://ouraring.com/blog/oura-ovulation-detection-algorithm-validation-study/?srsltid=AfmBOopXEniKiFc-PEISRaWfXt-KMl6fRjj59u0rqtXg5UFB_C00x6Z4

McDonald, A. M., Reeder, R. W., Kelley, P. G., & Cranor, L. F. (2009). A Comparative Study of Online Privacy Policies and Formats. In I. Goldberg & M. Atallah (Eds.), *Privacy Enhancing Technologies* (Vol. 5672, pp. 37–55). Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-03168-7_3

McDonald, A., & Cranor, L. (2008). The cost of reading privacy policies . *2008 Privacy Year in Review*, *4*(3), 543–569. A Journal of Law and Policy for the Information Society. https://lorrie.cranor.org/pubs/readingPolicyCost-authorDraft.pdf

Meta. (2022, July 26). *Meta Privacy Policy - How Meta collects and uses user data*. Www.facebook.com; Meta. https://www.facebook.com/privacy/policy/

Meta. (2024, November 12). *Facebook and Instagram to Offer Subscription for No Ads in Europe*. Meta. https://about.fb.com/news/2024/11/facebook-and-instagram-to-offer-subscription-for-no-ads-in-europe/

Miratech Staff. (2019, October 31). *The Marriott/Starwood Data Breach: Why Third-Party Risk Management is Critical During M&A*. Mitratech. https://mitratech.com/resource-hub/blog/the-marriott-starwood-data-breach-why-third-party-risk-management-is-critical-during-m-a/

Monteith, J., Shapcott, O., Talas, A., & Pranav Dahiya. (2023). Who Is Benefiting from Your Fitness Data? A Privacy Analysis of Smartwatches. In F. Stajano, V. Matyáš, B. Christianson, & J. Anderson (Eds.), *Security Protocols XXVIII* (Vol. 14186, pp. 97–112). Springer Cham. https://doi.org/10.1007/978-3-031-43033-6_11

Moore, A. D. (2003). Privacy: Its Meaning and Value. *American Philosophical Quarterly*, *40*(3), 215–227. JSTOR. https://doi.org/10.2307/20010117

Mousavi, S., Dash, A., Zannettou, S., & Gummadi , K. P. (2026). *Does Ad-Free Mean Less Data Collection? An Empirical Study of Platform Data Practices and User Expectations*. Arxiv.org. https://arxiv.org/html/2602.01231v1

Mozilla Foundation. (2022, August 17). *In Post Roe v. Wade Era, Mozilla Labels 18 of 25 Popular Period and Pregnancy Tracking Tech With *Privacy Not Included Warning*. Mozilla Foundation.

https://www.mozillafoundation.org/en/blog/in-post-roe-v-wade-era-mozilla-labels-18-of-25-popular-period-and-pregnancy-tracking-tech-with-privacy-not-included-warning/

Mozilla Foundation. (2024, November 19). *Every Breath You Take, Every Move You Make: The Hidden Costs of a $500B Market Centered on Our Bodies*. Mozilla Foundation. https://www.mozillafoundation.org/en/blog/every-breath-you-take-every-move-you-make-the-hidden-costs-of-a-500b-market-centered-on-our-bodies/

Na, L., Yang, C., Lo, C.-C., Zhao, F., Fukuoka, Y., & Aswani, A. (2018). Feasibility of Reidentifying Individuals in Large National Physical Activity Data Sets From Which Protected Health Information Has Been Removed With Use of Machine Learning. *JAMA Network Open*, *1*(8), e186040–e186040. https://doi.org/10.1001/jamanetworkopen.2018.6040

National Institute of Standards and Technology. (2023, July 12). *AI Risk Management Framework*. NIST. https://www.nist.gov/itl/ai-risk-management-framework

National Institute of Standards and Technology. (2025). *NIST Privacy Framework 1.1*. https://doi.org/10.6028/nist.cswp.40.ipd

Navlakha, M. (2024, April 2). *Google agrees to delete billions of Incognito mode data records*. Mashable. https://mashable.com/article/google-incognito-mode-data-lawsuit

Nguyen, T. T., Huynh, T. T., Ren, Z., Nguyen, T. T., Nguyen, P. L., Yin, H., & Nguyen, Q. V. H. (2024, March 31). *A Survey of Privacy-Preserving Model Explanations: Privacy Risks, Attacks, and Countermeasures*. ArXiv.org. https://doi.org/10.48550/arXiv.2404.00673

Office of the Assistant Secretary for Planning and Evaluation. (1996, August 20). *Health Insurance Portability and Accountability Act of 1996*. ASPE; U.S. Department of Health

and Human Services. https://aspe.hhs.gov/reports/health-insurance-portability-accountability-act-1996

ORCHA. (2022, July 27). *84% of Women's Health Period Apps Share Data with Third Parties*. ORCHA. https://us.orchahealth.com/84-of-25-womens-health-period-tracker-apps-share-data-with-third-parties/

Oregon Consumer Justice. (2026). *Consumer-Led Enforcement*. OCJ. https://ocj.org/glossary/consumer-led-enforcement

Oura Team. (2021, August 23). *UCSD Study Shows Clear Oura Temperature Increase After Conception*. The Pulse Blog. https://ouraring.com/blog/temperature-increase-pregnancy/

Peremore, K. (2023, July 18). *HIPAA compliance in wearable devices*. Www.paubox.com. https://www.paubox.com/blog/hipaa-compliance-in-wearable-devices

Perez, S. (2026, March 5). *Meta sued over AI smart glasses' privacy concerns, after workers reviewed nudity, sex, and other footage*. TechCrunch. https://techcrunch.com/2026/03/05/meta-sued-over-ai-smartglasses-privacy-concerns-after-workers-reviewed-nudity-sex-and-other-footage/

Persoskie, A., Hennessy, E., & Nelson, W. L. (2017). US Consumers' Understanding of Nutrition Labels in 2013: The Importance of Health Literacy. *Preventing Chronic Disease*, *14*. https://doi.org/10.5888/pcd14.170066

Policy Statement of the Federal Trade Commission on Biometric Information and Section 5 of the Federal Trade Commission Act 1. (2023). U.S. Federal Trade Commission. https://www.ftc.gov/system/files/ftc_gov/pdf/p225402biometricpolicystatement.pdf

Post, S. (n.d.). *The Meta Smokescreen*. Www.beuc.eu. https://www.beuc.eu/enforcement/meta-smokescreen

Proctor, R., Ali, A., & Vu, K.-P. (2008). *Examining the Usability of Web Privacy Policies*. https://www.researchgate.net/profile/Athar-Ali-7/publication/220302320_Examining_Usability_of_Web_Privacy_Policies/links/58468d5c08ae2d21756c9680/Examining-Usability-of-Web-Privacy-Policies.pdf

Protecting Consumer Privacy and Security. (n.d.). Federal Trade Commission. https://www.ftc.gov/news-events/topics/protecting-consumer-privacy-security

Radanliev, P. (2025). Privacy, ethics, transparency, and accountability in AI systems for wearable devices. *Frontiers in Digital Health*, *7*. https://doi.org/10.3389/fdgth.2025.1431246

Rahkovsky, I., Martinez, S., & Kuchler, F. (2012). New Food Choices Free of Trans Fats Better Align U.S. Diets With Health Recommendations. In *Economic Research Service*. United States Department of Agriculture. https://ers.usda.gov/sites/default/files/_laserfiche/publications/44672/18236_eib95.pdf?v=84899

Reidenberg, J. R., Breaux, T., Cranor, L. F., French, B. M., Grannis, A., Graves, J. T., Liu, F., McDonald, A., Norton, T. B., Rohan Ramanath, Russell, N. C., Sadeh, N., & Schaub, F. (2016). *Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding*. FLASH: The Fordham Law Archive of Scholarship and History. https://ir.lawnet.fordham.edu/faculty_scholarship/619/

Reidenberg, J. R., Breaux, T., Cranor, L. F., French, B., Grannis, A., Graves, J., Liu, F., McDonald, A., Norton, T., Ramanath, R., Russell, N. C., Sadeh, N., & Schaub, F. (2014).

Disagreeable Privacy Policies: Mismatches between Meaning and Userss Understanding. *SSRN Electronic Journal*, *30*(1). https://doi.org/10.2139/ssrn.2418297

Reviglio, U. (2022, August 4). *The untamed and discreet role of data brokers in surveillance capitalism: a transnational and interdisciplinary overview*. Policyreview.info. https://policyreview.info/articles/analysis/untamed-and-discreet-role-data-brokers-surveillance-capitalism-transnational-and

Ridgway, W. E., Grismer, M. J., Holmstrand, D. E., & Zivkovic, L. V. (2025, April 16). *District Court Rulings Could Signal Expansion of California Consumer Privacy Right of Action*. Skadden.com; Skadden, Arps, Slate, Meagher & Flom LLP. https://www.skadden.com/insights/publications/2025/04/district-court-rulings-could-signal-expansion

Rochford, J. (2023). *Cothron v. White Castle System, Inc.* Justia Law. https://law.justia.com/cases/illinois/supreme-court/2023/128004.html

Ruby, D. (2022, December 20). *Smartwatch Statistics 2023: How Many People Use Smartwatches?* Demand Sage. https://www.demandsage.com/smartwatch-statistics/

Santos, C., Morozovaite, V., & De Conca, S. (2025). No harm no foul: how harms caused by dark patterns are conceptualised and tackled under EU data protection, consumer and competition laws. *Information & Communications Technology Law*, 1–47. https://doi.org/10.1080/13600834.2025.2461958

Schaub, F. (2015, July 24). *A Design Space for Effective Privacy Notices*. Www.usenix.org. https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub

Schmalenberger, A., & Heywood, D. (2025, October 13). *Meta's "pay or OK" dilemma: the clash with EU digital regulation*. Taylorwessing.com; Taylor Wessing.

https://www.taylorwessing.com/en/global-data-hub/2025/eu-digital-laws-and-gdpr/gdh---metas-pay-or-ok-dilemma

SEC. (2023). *Document*. Sec.gov. https://www.sec.gov/Archives/edgar/data/1326801/000132680124000010/meta-12312023xexhibit991.htm

Sivakumar, V., Mone, V., & Rakhmanov Abdumukhtor. (2024). Addressing privacy concerns with wearable health monitoring technology. *Wiley Interdisciplinary Reviews. Data Mining and Knowledge Discovery/Wiley Interdisciplinary Reviews. Data Mining and Knowledge Discovery*, *14*(3). https://doi.org/10.1002/widm.1535

Skovgaards, K. (2025, November 4). *Consent or Pay: What Publishers Need to Know*. Consent Management Platform (CMP) Usercentrics. https://usercentrics.com/knowledge-hub/consent-or-pay-what-publishers-should-expect/

Smalley, S. (2025). *Major location data broker reports hack to Norwegian authorities*. Therecord.media. https://therecord.media/location-data-broker-gravy-breach

Smigielski, M., & Kantrow, J. M. (2022). *BIPA Class Actions Explode as Small Businesses Struggle with Lawsuits*. Lewis Brisbois. https://lewisbrisbois.com/insights/podcasts/bipa-class-actions-explode-as-small-businesses-struggle-with-lawsuits

Solove, D. J. (2006). A Taxonomy of Privacy. *University of Pennsylvania Law Review*, *154*(3), 477–564. https://doi.org/10.2307/40041279

Solove, D. J., & Hartzog, W. (2013). The FTC and the New Common Law of Privacy. *Colombia Law Review*, *114:583*, 583–676. https://doi.org/10.2139/ssrn.2312913

State of California. (2026). *About DROP and the Delete Act*. Privacy.ca.gov. https://privacy.ca.gov/drop/about-drop-and-the-delete-act/

State of California Department of Justice. (2024, March 13). *California Consumer Privacy Act (CCPA)*. State of California - Department of Justice - Office of the Attorney General. https://oag.ca.gov/privacy/ccpa

Team, C. (2025, May 30). *Consent or Pay: What You Need to Know*. CookieYes. https://www.cookieyes.com/blog/consent-or-pay/

Team, U. (2026, March 9). *Consent-or-Pay for Publishers: Compliance Guide for the UK and EU (2026)*. UniSignIn. https://www.unisignin.com/blog/consent-or-pay-publishers-compliance-guide

Text of proposed laws: Proposition 24. (n.d.). https://vig.cdn.sos.ca.gov/2020/general/pdf/topl-prop24.pdf

Thapa, C., & Camtepe, S. (2021). Precision health data: Requirements, challenges and existing techniques for data security and privacy. *Computers in Biology and Medicine*, *129*(1), 104130. https://doi.org/10.1016/j.compbiomed.2020.104130

The Use of Alternative Data in Credit Risk Assessment: Opportunities, Risks, and Challenges 2024. (2024). In *World Bank Group*. https://documents1.worldbank.org/curated/en/099031325132018527/pdf/P179614-3e01b947-cbae-41e4-85dd-2905b6187932.pdf

Tim Lamb. (2025, July 3). *Why the Commission's Decision Undermines the Goals of the DMA*. Meta Newsroom. https://about.fb.com/news/2025/07/why-the-commissions-decision-undermines-the-goals-of-the-dma/

Tobitt, C. (2025, August 21). *More UK news publishers are adopting "consent or pay" advertising model*. Press Gazette. https://pressgazette.co.uk/marketing/consent-or-pay-uk-publishers-advertising-2025/

Todd, J., & Variyam, J. (2008). *The Decline in Consumer Use of Food Nutrition Labels*. United States Department of Agriculture. https://ers.usda.gov/sites/default/files/_laserfiche/publications/46061/12042_err63_1_.pdf?v=51392

Torchinsky, R. (2022, June 24). How period tracking apps and data privacy fit into a post-Roe v. Wade climate. *NPR*. https://www.npr.org/2022/05/10/1097482967/roe-v-wade-supreme-court-abortion-period-apps

Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2011). The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research*, *22*(2), 254–268. https://doi.org/10.1287/isre.1090.0260

Union of Concerned Scientists. (2016). *Transparency in Food Labeling: Food Labels Inform Consumer Choices—and Industry Pushes Back*. Union of Concerned Scientists. http://www.jstor.org/stable/resrep17295

UnitedHealthCare. (2025, October 6). *Members can earn up to $300 for reaching program daily health and activity goals*. Www.uhc.com. https://www.uhc.com/news-articles/benefits-and-coverage/wellness-rewards

Vandenberghe, M. (2026, January 23). *Health Data Risk in AI: The Hidden Threats You Should Know*. MyData-TRUST. https://www.mydata-trust.com/2026/01/23/health-data-risk-ai/

Vermont Attorney General. (2018). *Guidance on Vermont's Act 171 of 2018 Data Broker Regulation*. https://ago.vermont.gov/wp-content/uploads/2018/12/2018-12-11-VT-Data-Broker-Regulation-Guidance.pdf

Waldman, A. E. (2020). Cognitive biases, dark patterns, and the "privacy paradox." *Current Opinion in Psychology*, *31*(1), 105–109. https://doi.org/10.1016/j.copsyc.2019.08.025

Wang, Z. (2025, December 9). *Meta will offer EU users option of less personalized ads, caving to EU pressure | The Current*. The Current. https://www.thecurrent.com/data-privacy-meta-users-option-less-personalized-ads-eu

West, D. M. (2025, April 15). *How AI can enable public surveillance*. Brookings. https://www.brookings.edu/articles/how-ai-can-enable-public-surveillance/

What good and effective data privacy accountability looks like: Mapping organisations' practices to the CIPL Accountability Framework. (2020). Centre for Information Policy Institute (CIPL). https://www.informationpolicycentre.com/wp-content/uploads/2025/04/cipl_accountability_mapping_report__27_may_2020__v2.0-2.pdf

WHOOP. (2025). *Full Privacy Policy*. WHOOP. https://www.whoop.com/us/en/full-privacy-policy/?srsltid=AfmBOoqMn3OwTvqZ4uuaqYupWyLd4u41_8MpnVkqd9tSKqdHTILYOwpz

Winder, D. (2025, May 15). *Healthcare Hackers Compromise 276 Million Patient Records*. Forbes. https://www.forbes.com/sites/daveywinder/2025/05/15/276-million-patient-records-compromised---what-you-need-to-know/

Wolford, B. (2024, February 8). *What's your data really worth?* Proton. https://proton.me/blog/what-is-your-data-worth

Yaros, O., & Hepworth, E. (2023, July 17). *Cookie Banners Missing "Reject All" Buttons Face Investigation by the UK Information Commissioner's Office*. Www.mayerbrown.com. https://www.mayerbrown.com/en/insights/publications/2023/07/cookie-banners-missing-reject-all-buttons-face-investigation-by-the-uk-information-commissioners-office

Zabierek, L., Bolton, T., Pugh, B., & Lesmes, S. (2022, June 14). *Limiting a Private Right of Action in Federal Data Security and Privacy Legislation*. The Belfer Center for Science and International Affairs. https://www.belfercenter.org/publication/limiting-private-right-action-federal-data-security-and-privacy-legislation

Zhang, S., Feng, Y., Yao, Y., Cranor, L. F., & Sadeh, N. (2022). How Usable Are iOS App Privacy Labels? *Proceedings on Privacy Enhancing Technologies*, *2022*(4), 204–228. https://doi.org/10.56553/popets-2022-0106

Zhao, J. C., Elkordy, A. R., Sharma, A., Ezzeldin, Yahya H, Avestimehr, S., & Bagchi, S. (2023). *The Resource Problem of Using Linear Layer Leakage Attack in Federated Learning*. ArXiv.org. https://arxiv.org/abs/2303.14868

Zhao, J., Bagchi, S., Avestimehr, S., Chan, K., Chaterji, S., Dimitriadis, D., Li, J., Li, N., Nourian, A., & Roth, H. (2025). The Federation Strikes Back: A Survey of Federated Learning Privacy Attacks, Defenses, Applications, and Policy Landscape. *ACM Computing Surveys*, *57*(9), 1–37. https://doi.org/10.1145/3724113

Zhu, L., Liu, Z., & Han, S. (2019). Deep Leakage from Gradients. *ArXiv:1906.08935 [Cs, Stat]*. https://arxiv.org/abs/1906.08935