Future of Data Initiative
Massachusetts Institute of Technology

Accountability and Traceability White Paper and Research Roadmap

18 April 2023

https://futureofdata.mit.edu/tr/2023/fod-account-trace-20230418.pdf

Internet Policy Research Initiative

Future of Data - Working Paper Series

2023

MIT CSAIL Computing and Society

https://futureofdata.mit.edu

Accountability and Traceability White Paper & Research Roadmap

April 18, 2023

Status of this document: This document is a draft produced by the MIT Future of Data Initiative (FoD). Technical, public policy and business perspectives reflected here are based on the two roundtable discussions: the Future of Data Research Roundtable on Technical Aspects of Data Governance (April 4, 2022) and The Privacy Tech Marketplace - Designing and deploying modern data governance as a service (11 October 2022). With this public version, we invite comments from interested parties. Comments can be sent to the editor or to ipri-contact@mit.edu. This version of the document has benefitted from contributions by Future of Data Initiative Members, American Family Insurance, Capital One, Fidelity Investments, MassMutual, and Visa, and was edited by Daniel Weitzner (MIT Computer Science and Artificial Intelligence Lab). A full list of contributors is at the end of the document. This document does not represent the formal views of any FoD participating organizations.

1. Overview and Motivation

The MIT Future of Data Initiative is leading a multi-disciplinary research agenda to design and stimulate the deployment of consumer-empowering and accountable systems to provide trusted, traceable uses of personal data on an ecosystem-wide scale. The Initiative has gathered together computer science and Internet policy researchers as well as leading commercial enterprises in financial services, payment technology, cloud platforms, insurance and other sectors to discuss current challenges and opportunities in privacy and data governance. Today's modern privacy laws place appropriately high expectations on organizations processing personal data. At the same time, consumers report declining trust in those who handle their personal data and regulators around the world struggle with the scale of the enforcement challenge. We aim to identify and put into service technical infrastructure for enterprises seeking to handle personal data in a trustworthy and lawful manner with guardrails to enable the traceable, accountable, and scalable use of data.

The Opportunity

Research at MIT has identified new system architectures and software engineering techniques that can help close the gap between the requirements of privacy law and capabilities of data systems. These systems are in early stages and will require close collaboration between computer scientists, legal scholars, and businesses that seek more policy-aware systems. Discussions in Future of Data workshops confirmed that many enterprises, large and small, are eager for new capabilities in their software environments that support privacy compliance, enhance accountability and increase consumer transparency and control. A critical success factor for these new systems will be widespread deployment and cross-enterprise interoperability. Future of Data participants see an opportunity to work in a cross-disciplinary effort with businesses, software vendors, regulators and civil society to propagate these new technical approaches and business practices to increase consumer trust and confidence in modern data services.

Accountable systems provide confidence that personal data will be handled according to established rules and applicable legal frameworks, and that potential violations of those rules can be identified and remediated. (Weitzner et al. 2008; Feigenbaum, Jaggard, and Wright 2020; Kacianka and Pretschner 2021). To build trust in our digital economy, we must deploy systems that prioritize accountability. A recent global consumer study across fifteen markets shows that more than half of consumers believe companies use their data *beyond* the permissioned terms.[1] One cause of this distrust is that consumers have little visibility into how their data is used. Accountability depends, fundamentally, on the traceability of personal data. Without basic traceability to provide transparency into the downstream uses of data, there can be no accountability to demonstrate permissible use or detect misuse. Without accountability, trust in digital systems will continue to be justifiably low.

This white paper describes the goals and broad design requirements of accountable, traceable systems to guide the multi-disciplinary research effort underway at the MIT Future of Data Initiative.

2. Design Principles of Accountable Systems

The goal of deploying accountable systems is to empower consumers and assure consumer protection authorities that data is being handled in a respectful and trusted manner. Respectful use of data means that there are no surprises, and consumers are provided control over their personal data. Trust is earned by clear evidence that data is collected, used, and stored according to applicable law and best practices.

We can support trustworthy and respectful use of personal data by building systems that meet the following five design principles:

1. *Traceable*: collecting and maintaining data provenance to build trust and provide transparency to consumers, partners, and possibly regulators. Traceability enables consumers to know where their data is, who has it, what it is being used for, and whom it is being shared with. Accurate provenance also enables key privacy law requirements such as the right to erasure and the right to correction. Data traceability is to the benefit of consumers and will help companies earn the trust of their customers and confidence of regulators.
2. *Accountable*: uses of personal data are controlled in a manner to enable monitoring appropriate use as well as detection and consequences for misuse of data (Lampson 2009). Accountable systems will indicate when data uses are tied to a necessary, clear, and legitimate interest such as fraud prevention, legal compliance, and other consumer-

[1] Visa Consumer Empowerment Study (2022), n=19,600, 15 markets include Canada, USA, Germany, France, UK, Japan, Hong Kong, Australia, Indonesia, Vietnam, Saudi Arabia, UAE, South Africa, Nigeria, Guatemala

consented uses, and when, on the other hand, there is misuse. As the complexity of personal data services grows and new privacy laws are enacted, enterprises and regulators face corresponding complexity in navigating their legal obligations regarding processing and transferring personal data. Building on traceability solutions, we can make it possible for organizations to analyze and maintain indicators of data use to monitor compliance with internal policies, consumer consent, contractual obligations and legal requirements.

3. *Granular consent:* respect for consumer rights requires organizations to adhere to the standard of meaningful, freely given, and affirmative consent (GDPR Art. 7, Rec. 32), enabling consumers to exercise granular control to accept or reject specific data use or disclosures (Visa, Inc. 2022). Default consent settings that require individuals to actively opt-out may confuse consumers. Accountable systems must include mechanisms for representing and tracking consent at a granular level, reflecting the often-complex structure of personal data and the many uses to which it may be subject.
4. *Agile*: personal data can be used for analytic and operational purposes while upholding privacy values and retaining trust, transparency, and accountability.
5. *Scalable*: personal data can be processed in an accountable fashion within and across enterprise boundaries as part of global ecosystems.

These design principles cover key aspects of handling personal data in a respectful, trustworthy, and lawful manner, based on the obligations inherent in modern privacy laws. They represent key technical challenges we propose to address in our research, but do not propose to cover every legal privacy requirement.

Public Policy Context for Accountability and Traceability

Information privacy practice and governance has evolved over time to be more sophisticated, largely in response to increased computing power and based on demands from society to protect individuals from intrusive use and misuse of personal data. The 1970s-90s saw mostly articulation of broad privacy principles, notably the Fair Information Practice Principles developed in the US and other democratic societies around the world. Early privacy frameworks in the United States such as the Privacy Act of 1974 (oversight on government information practices), Europe and early OECD guidance (1980 Privacy Guidelines) were similarly articulated at the level of general principle. These tended to focus on guidance for and obligations on institutions, with less direct attention to individual rights.

As Internet commerce developed beginning in the mid-1990s, policymakers in the United States led the way in calling for the adoption of *privacy policies* by all businesses operating online. One of the earliest pieces of legislation that required web operators to have a privacy policy was the 2003 CalOPPA. Those policies became the basis for privacy enforcement by the US Federal Trade Commission. Soon, to increase the likelihood that businesses would actually follow their stated policies, businesses themselves, often pushed by regulators, developed privacy programs, led by

Chief Privacy Officers and growing teams. And in this period of time, the EU enacted the Data Protection Directive of 1995, aimed at harmonizing privacy principles across Europe to promote a single market. These were important advances but were largely legal, administrative, and organizational in nature. The EU Data Protection Directive was the precursor to the General Data Protection Regulation (GDPR) which came into effect May 2018.

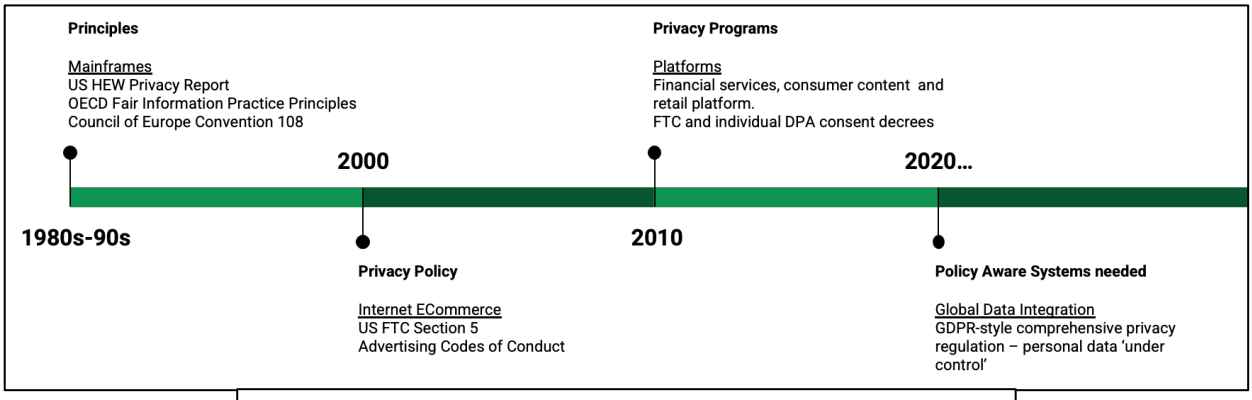

Fig 1. Developments in Modern Privacy Law, Institutions and Technology

The Challenge: The gap between privacy law and computer systems

Modern data privacy law, as exemplified by the GDPR, has caught up with many of the privacy risks posed by advanced data analytics, but now technology has fallen behind. Hence, there is a growing gap between privacy policy and the systems which process personal data. The challenge presented is to understand the gap between privacy law and enterprise systems and develop a strategy to bridge that gap. Our research has shown that today's database systems and other foundations of modern systems development environments require more technical and organizational controls to operationalize privacy concepts such as purpose limitation, data minimization, and the right to be forgotten (Kraska et al. 2019; Wang et al. 2022). These challenges are especially acute when processing data at scale or across enterprise boundaries. The gap between law and technology ultimately makes it hard for well-meaning organizations to do the right thing with personal data and makes it hard for regulators to understand and address emergent risks to data or to detect violations in certain cases. Most importantly, organizational and technical obstacles such as lack of governance tools may hinder organizations from assessing whether systems are handling data according to policy requirements and thus respecting the privacy commitments that were made to consumers.

With the advent of modern privacy law such as the European Union General Data Protection Directive (GDPR) and the California Consumer Privacy Act (CCPA), there is a growing gap between practices stated in organizational policies and the ability to demonstrate that actual systems handle personal data according to those rules. Given these challenges, the next step in the timeline (Fig.1) must be to introduce accountable, policy-aware systems.

Consumer Research on Traceability

Recent studies of consumer attitudes about privacy show a substantial trust gap. In a global survey across 15 markets, nearly 50% of consumers surveyed think that companies use data that goes beyond terms of consent. Similarly, only about half of consumers say they have "at least some visibility" into their data use post-sharing. In this same survey population, two-thirds of consumers surveyed would find value and comfort in a solution to trace their data post-sharing. 61% of consumers surveyed believe that added visibility would boost trust and 55% of consumer surveyed believe that they would be more likely to share their data if a "trace" solution were available.[2]

Limits on Scope of Research

There are some areas that are out of scope for the Future of Data Initiative: We do not seek to change privacy law, nor to remedy common and important complaints with online advertising, and we do not believe that there are quick, cryptographic fixes to today's privacy challenges. Privacy depends on the security of personal data, a vital topic, but one which is also beyond the scope of our research. Developing policy-aware privacy architecture is a cross-disciplinary problem, which spans business, technical, and legal research.

3. Representative Use Cases

Research challenges in building accountable systems abound. To give focus to our early rounds of research, we will concentrate work on two use cases in which comprehensive accountability and traceability are needed: open banking systems and cloud-based data integration with specific personal data governance requirements. We describe each use case here. In the following sections we will suggest broad technical challenges to address these use cases (Section 4) and describe a research sandbox we are building to give researchers the ability to test their hypotheses and solution builds.

Open Banking Systems

Consumer data-sharing ecosystems being deployed around the world will facilitate innovation in consumer banking services, but also raise novel questions regarding consumer trust and the need for personal data governance across organizational boundaries. The growing open banking environment will depend on accountability and traceability features to assure respectful use of personal data while enabling more open flow and analysis of personal financial information. Both consumers and regulators are demanding that personal data governance capabilities be deployed alongside open banking APIs, but there is much to learn about how to design and deploy such services at scale.

---

[2] Visa Consumer Empowerment Study (2022), n=19,600, 15 markets include Canada, USA, Germany, France, UK, Japan, Hong Kong, Australia, Indonesia, Vietnam, Saudi Arabia, UAE, South Africa, Nigeria, Guatemala

One common consumer data sharing system is "Open Banking,' a system through which consumers or businesses authorize third-parties to access their financial information such as bank and investment account data (e.g., transaction or payment history) or services (e.g., making a payment or requesting a loan). When consumers or businesses choose to share their financial data with third parties, the third parties can, in turn, provide several products and services, including budgeting, credit checks or help initiating payments. Research shows that 87 percent of U.S. consumers are using open banking to link their financial accounts to third parties, however only 43 percent of U.S. consumers are aware that they are using open banking.[3]

The financial services industry in the United States has collaborated to develop a technical framework for exchange of personal financial information, the Financial Data Exchange (FDX). While FDX has done considerable work to enable actual exchange of data, much work needs to be done, both on the technical and policy fronts, to enable well governed exchange of that personal data. The research sandbox described in Section 5 will provide sample data sets, APIs, and relevant policies for researchers to experiment with in building research projects to address accountability and traceability needs in open banking use cases.

Multi-Enterprise Cloud-based Data Integration with Data Governance

Multi-party sharing can include models such as collaboration platforms, data exchanges, data trusts, and cloud-based federated ecosystems. The models and mechanisms are varied, but the aim of each is to derive greater value from data through some form of sharing, whether that involves the transfer of data sets, granting access to data sets, or some form of data collaboration. Each participant contributes new data sources or services, thereby adding additional value to the broader ecosystem. Early Future of Data research demonstrated the use of Secure Multiparty Computation systems for assessing cybersecurity risk across a large number of enterprises (Castro et al. 2020).

Within multi-party data sharing opportunities, standards of data security, privacy, control, transparency, monitoring, permissioning, and auditability must apply. Data exchange must be facilitated between participants in a trusted, legally compliant environment. Accountability and governance mechanisms that measure, trace, track, and report how data moves through and outside of such data integration platform need to be designed to build trust and safeguard against data misuse and unauthorized access.

Future of Data Initiative partners have a particular interest in enabling private computation and privacy-preserving machine learning in order to test properties of various models for bias and discrimination. In certain cases, regulated industries are required to demonstrate the impact of underwriting, credit scoring, or other decision-making models on protected classes of customers

[3] Visa Open Banking Consumer Survey, n=1,500, representative sample of U.S. population based on Census Bureau. Administered digitally, April 2022.

(generally race, gender, etc.). At the same time, laws may prevent those same firms from collecting such sensitive identifying information about their customers, precisely to reduce the risk that sensitive data will be used in the decision-making process. Due to overlapping regulatory frameworks, firms using predictive models may be reluctant to share the details of those systems as they reveal proprietary information. Private computation approaches that allow measurement of model behavior without sharing protected class information directly with companies, and without requiring companies to share proprietary information would help achieve important regulatory goals and earn trust in machine learning-driven decision making. For example, a recent UK Financial Conduct Authority Hackathon investigated the value of these techniques for Anti-Money Laundering approaches.[4]

Subsequent releases of the research sandbox described in Section 5 will provide sample data sets, APIs and relevant policies for researchers to experiment with in building research projects to address accountability and traceability needs in these multi-enterprise cloud environments.

4. Research Challenges

These are the key technical research challenges to be explored in developing accountable, traceable systems. The MIT Future of Data Initiative will encourage and support research in 4 principal research areas:

A. Multi-enterprise, open Internet and cross-enterprise audit and traceability protocols
B. Privacy-preserving data analytics
C. Database architectures for accountability, traceability and personal data governance
D. User Experience design for accountability and traceability

These are the descriptions of research topics in each of the four areas.

Multi-enterprise, Open Internet and Cross-enterprise Audit and Traceability Protocols

Cross-enterprise flows of personal data, along with complex intra-enterprise uses, require new provenance tracking and usage logging to enable traceability and accountability. How should those facilities be designed and how will they scale? What degree of information assurance is required for the logging and audit features? Do we need public ledger techniques for provable log integrity? If so, how should those be designed? What are the threat models against which accountability and traceability systems should be designed? What kinds of protocols can be designed to manage personal data flow and enable accountability and traceability in use cases such as Open Banking as described in Section 3?

The Internet is an open data platform and as such consumer expectations of the interconnected digital experience is driving more data to be shared between companies. At the same time, the

[4] https://www.fca.org.uk/firms/innovation/techsprints ; https://www.future-fis.com/the-pet-project.html

consumer holds their primary institution accountable as their trust center even after the data has left their walls. This creates significant challenges for companies maintaining trust of usage while balancing customer experience. Technological innovation can help increase trust with better accountability and traceability, while lowering compliance costs and increasing data agility.

Organizations that enjoy broad trust from consumers could act as trust centers to log and audit data tracing requests for consented data as it propagates through multi-party ecosystems. This is not dissimilar to the role played by a Root Certificate Authority in SSL/TLS models. Existing PKI users may want to interact directly with the third-party trust center with regards to audits of their data. An internet scale, secure, easy to implement system would need to be designed.

Accountability protocols might be implemented at a number of different points of the data workflow. If it is enforced at the transaction level, it can be within the API call or at the level of the driver. Wherever it is implemented, there will need to be some standards that need to be adhered to by the processing system. For instance, database processing engines like AWS RDS: MySql, Aurora, etc. These standards will generally ensure the valid processing of data (consented or otherwise), full compliance with the protocol will ensure that technical standards for data protection are met, regardless of external data sharing or not. Some of the requirements associated with such protocols include:

- Data Protection at Rest: How is data protected at storage time. How does that data behave when cut off from a trust authority for an extended period of time?
- Data Protection at Query Time: How can non-consented data be filtered out or masked out at query time?
- Consent: How can we trace data sharing conditional on permissioned consent, with attention paid to purpose and key terms of data access?
- Data Retention Policies: What can be done when the consumer revokes access to the data? Does the data reach a state of entropy after a certain period of time regardless of consent grants?
- Audit: How can we design clear, predictable means to ensure that all data protection standards associated with any given data set are adhered to?

Research here should consider the scaling requirements associated with a large number of parties participating in data sharing systems, and the need to design user experiences that provide both consumers and regulators with meaningful indicia of trust and control over personal data flowing through these complex environments.

Privacy-preserving Data Analytics

Data sharing and aggregation across institutions can enable business analytics from a multi-institutional or even ecosystem-wide vantage point. And even within institutional boundaries, privacy-preserving computation can enable enterprises to realize aggregate insights from personal

data without revealing data from individual sources.[5] While aggregated analytics can be mutually beneficial for participants, collecting them may require participants to reveal sensitive business information to competitors. Recent research on a wide variety of *privacy preserving* data analytics protocols (e.g., multi-party computation (MPC), federated learning, private set intersection) enables mutually untrusting parties to jointly compute public outputs over private inputs without relying on a trusted third-party. The design of a privacy-preserving data analytics protocols for the outlined use cases will depend on various parameters: What kinds of aggregate statistics are useful? What kinds of trust assumptions are permissible? What kind of privacy guarantees (e.g., cryptographic, differential privacy, or other new definitions) are required under different legal and contractual frameworks? How can we establish protocols for transparency in analytics conducted on the data? What threat models should be considered when employing MPC approaches and what methods can be used to test privacy? Research proposals on this topic should address techniques to address use cases in open banking and regulatory assessments of fairness, bias and disparate impact. As this research develops, FoD partner companies will work to provide synthetic data to serve as testbeds to evaluate proposed designs as they develop.

Database Architectures for Accountability, Traceability and Personal Data Governance

What new data management architectures will provide enterprises with data control, provenance, lifecycle oversight, and accountability tools for managing personal data according to legal frameworks and institutional commitments? What kinds of protocols can be designed to manage personal data flow and enable accountability and traceability in use cases such as Open Banking as described in Section 3? The EU GDPR, along with other current and proposed privacy laws, call on institutions to rethink their current data governance structures to meet evolving requirements. Beyond requirements in statute, enterprises have many internal commitments and contractual requirements they seek to respect in their data analytic processes. What kinds of system architectures can support those data use controls? How can they be built to operate efficiently at scale? While clean-slate solutions may hold some appeal, we are particularly interested in what strategies are available to address the heterogeneous legacy data systems that characterize almost all enterprise data environments today? Given the complexities of most enterprise data systems, it is highly unlikely that any single database architecture will solve a meaningful problem.

User Experience Design for Accountability and Traceability

User experience research will play a vital role in understanding how do design accountability and traceability features to build consumer trust. We invite research that answers the following questions about the user experience in the use cases described here:

[5] World Economic Forum, "The Next Generation of Data-Sharing in Financial Services: Using Privacy Enhancing Techniques to Unlock New Value", Sep 2019

- What information affects a consumer's stated level of trust and how will that sense of trust or distrust be reflected in sharing behaviors? What indications can we look to for informed, freely given sharing decisions?
- How will consumers give, manage, and revoke consent between various stakeholders (e.g., data providers, recipients, and aggregators)?
- How will consumers request deletion of their data?
- What user experience considerations should researchers explore to enhance consumer trust and inform the ability of consumer to trace their data from consent, collection, use and sharing to deletion?
- What can we learn about the impact of the volume of decisions consumers are asked to make with respect to controlling their data?
- What do consumers want to know regarding third party use of their personal data usage? How do consumers want to be notified when their personal data is used, and how do consumers want to be informed on their personal data usage post-revocation of consent to ensure that there is no misuse of their data?
- How can enterprises in complex multi-party ecosystems manage and distribute responsibility for operationalizing these consumer expectations?

How can design decisions can be examined through user studies (Nouwens et al. 2020)? Economics, psychology and sociology all have long research traditions of how trust is gained, lost and regained. We are especially interested in bringing together UX designers, social scientists and computer scientists on this topic.

5. Research Sandbox

A research sandbox soon to be released will provide resources (e.g., specifications, code, synthetic data sets, relevant policies, computing infrastructure) to support the research problem statements outlined in Section 4. Our goal in introducing the research sandbox is to provide researchers with sample data, APIs and policies relevant to the specific accountability and traceability scenarios to which we hope research results can apply.

The initial release of the sandbox focuses on the open banking use case and will include:

- A synthetic data set of customer transactions
- Open banking API specifications (e.g., FDX[6]) and code implementations
- Computing infrastructure (e.g., Amazon S3 for storage and EC2 for compute) for running a simulated open banking environment.

[6]https://financialdataexchange.org/FDX/News/Press-Releases/Financial%20Data%20Exchange%20Releases%20FDX%20API%205.1.aspx

Subsequent releases will expand the scope of the sandbox to include the multi-party data sharing use case as well.

6. Next steps

Research on the topics described in this White Paper is already underway. As we have results from that work, the Future of Data Initiative will publish our findings. We will also continue to engage in discussions on the design and deployment of consumer-empowering, accountable, and traceable systems with industry partners, civil society organizations and regulators around the world. Our aim is both to advance the technical state of the art in our theoretical understanding of accountable systems and to encourage and support the deployment of accountability and traceability techniques in publicly deployed infrastructure.

Contributors

Nirmal Baid, Vice President, Data & AI Initiatives, Visa
Bayan Bruss, Sr. Director, Applied ML Research, Capital One
Kevin Fitzpatrick, Head of Privacy, Data & AI Governance, MassMutual
Adam Fox, Head of Distribution Technology and Data Science, MassMutual
Robert Hedges, Chief Data Officer, Visa
Ilaria Liccardi, Research Scientist, MIT
Awah Teh, Vice President, Data Governance and Privacy Engineering, Capital One
Mona Vernon, Head of Fidelity Labs, Fidelity Investments,
Daniel Weitzner, 3Com Founders Senior Research Scientist, MIT (editor)

References

Castro, Leo de, Andrew W. Lo, Taylor Reynolds, Fransisca Susan, Vinod Vaikuntanathan, Daniel J. Weitzner, and Nicolas Zhang. 2020. "SCRAM: A Platform for Securely Measuring Cyber Risk." *Harvard Data Science Review*, July. https://doi.org/10.1162/99608f92.b4bb506a.

Feigenbaum, Joan, Aaron D Jaggard, and Rebecca N Wright. 2020. "Accountability in Computing: Concepts and Mechanisms." *Foundations and Trends®in Privacy and Security* 2 (4): 247–399.

Kacianka, Severin, and Alexander Pretschner. 2021. "Designing Accountable Systems." In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 424–37.

Kraska, Tim, Michael Stonebraker, Michael Brodie, Sacha Servan-Schreiber, and Daniel Weitzner. 2019. "SchengenDB: A Data Protection Database Proposal." In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*, 24–38. Springer.

Lampson, Butler. 2009. "Privacy and Security: Usable Security: How to Get It." *Communications of the ACM* 52 (11): 25–27. https://doi.org/10.1145/1592761.1592773.

Nouwens, Midas, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. "Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence." In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13.

Visa, Inc. 2022. "Visa Consent Management Specification." https://images.globalclient.visa.com/Web/InovantElqVisaCheckout/%7B3c16d577-f2c7-43a8-8d57-543ff7d19c82%7D_Consent_Management_Specification_v2_Public_Final.pdf.

Wang, Lun, Usmann Khan, Joseph Near, Qi Pang, Jithendaraa Subramanian, Neel Somani, Peng Gao, Andrew Low, and Dawn Song. 2022. "PrivGuard: Privacy Regulation Compliance Made Easier." In . https://www.usenix.org/conference/usenixsecurity22/presentation/wang-lun.

Weitzner, Daniel J., Harold Abelson, Tim Berners-Lee, Joan Feigenbaum, James Hendler, and G.J. Sussman. 2008. "Information Accountability." *Communications of the ACM* 51 (6): 82–87. https://doi.org/10.1145/1349026.1349043.