

BY DANIEL J. WEITZNER, HAROLD ABELSON, TIM BERNERS-LEE,
JOAN FEIGENBAUM, JAMES HENDLER, AND GERALD JAY SUSSMAN



INFORMATION ACCOUNTABILITY


*With access control and encryption no longer
capable of protecting privacy, laws and systems are needed
that hold people accountable for the misuse of personal
information, whether public or secret.*

Existing legal and technical mechanisms intended to protect our privacy, copyright, and other important values have been overwhelmed by the increasingly open information environment in which we live. These threats follow from the ease of information storage, transportation, aggregation, and analysis. We face the real risk that the technical laws spelled out by Gordon Moore (growth in processing power) and Robert Metcalfe (network effects) will permanently overwhelm our values as enshrined in society's laws.

ILLUSTRATION BY JEAN-FRANÇOIS PODEVIN



John H. Dow Credit Limit based on Annual Income



CALIFORNIA DRIVER LICENSE

FRANCOIS FELIX PODVIN
MARISSON MF:198
EYES:BRN
DOB:19-03-84

OFFICIAL SPECIAL-CLASS ELECTION BALLOT


Ballot form with sections C, D, and E. Section C includes a list of candidates with checkboxes. Section D contains numbered questions with 'YES' and 'NO' options. Section E contains a question about charter amendments with 'YES' and 'NO' options.

Entreprise du pape :

11. Domicile/Residence
5182 - Rue
Whitla

15 OCT 2007

John H. Dow Credit Limit based on Annual Income



WANTED

WANTED
NATIONAL IDENTITY CARD



Victor Szew
CIRCULATION



SAMPLE GENERAL BALLOT CASE

For too long, our approach to information-protection policy has been to seek ways to prevent information from “escaping” beyond appropriate boundaries, then wring our hands when it inevitably does. This hide-it-or-lose-it perspective dominates technical and public-policy approaches to fundamental social questions of online privacy, copyright, and surveillance. Yet it is increasingly inadequate for a connected world where information is easily copied and aggregated and automated correlations and inferences across multiple databases uncover information even when it is not revealed explicitly. As an alternative, accountability must become a primary means through which society addresses appropriate use. Information accountability means the use of information should be transparent so it is possible to determine whether a particular use is appropriate under a given set of rules and that the system enables individuals and institutions to be held accountable for misuse.

Transparency and accountability make bad acts visible to all concerned. However, visibility alone does not guarantee compliance. Then again, the vast majority of legal and social rules that form the fabric of our societies are not enforced perfectly or automatically, yet somehow most of us still manage to follow most of them most of the time. We do so because social systems built up over thousands of years encourage us, often making compliance easier than violation. For those rare cases where rules are broken, we are all aware that we may be held accountable through a process that looks back through the records of our actions and assesses them against the rules.

Personal privacy, copyright protection, and government surveillance are among the more intractable policy challenges in our information society. In each of these policy areas, excessive reliance on secrecy and up-front control over information has yielded policies that fail to meet social needs, as well as technologies that stifle information flow without actually resolving the problems for which they were designed.

Information privacy rights aim to safeguard individual autonomy against the power that institutions or individuals gain over others through the use of personal information.¹ Sensitive, and possibly inaccurate, information may be used against people in financial, political, employment, and health-care settings. In democratic societies, citizens’ behavior is unduly restrained if they fear they are being watched at every turn. They may deliberately avoid reading controver-

sial material or feel inhibited from associating with certain communities and ideas for fear of adverse consequences.

Protecting privacy is more challenging than ever due to the proliferation of personal information on the Web and the increasing analytical power available to large institutions (and to everyone else) through Web search engines and other facilities.² Access control and collection limits over a single instance of personal data are insufficient to guarantee the protection of privacy when either the same information is publicly available elsewhere on the Web or it is possible to infer private details to a high degree of accuracy from other information that itself is public [8, 10]. Worse, many privacy protections (such as lengthy online privacy-policy statements in health care and financial services) are mere fig leaves over the increasing exposure of our social and commercial interactions. In the case of publicly available personal information, people often intentionally make the data available, not always by accident [9]. They may not intend for it to be used for every conceivable purpose but are willing for it to be public nonetheless.

Even technological tools that help individuals make informed choices about data-collection practices are no longer sufficient to protect privacy in the age of the Web. As a case in point, the growth of e-commerce over the second half of the 1990s sparked concern among Web users worldwide about their personal privacy and led businesses to emphasize Website privacy policies and infrastructure (such as the World Wide Web Consortium’s Platform for Privacy Preferences, or P3P, www.w3.org/P3P/). A fully implemented P3P environment gives Web users the ability to make privacy choices about every single request by business organizations and government agencies to collect information about them. However, the number, frequency, and specificity of these choices would be overwhelming, especially if they were to cover all possible future uses by the data collector and by third parties. Individuals should not have to agree in advance to complex policies with unpredictable outcomes. Moreover, they should be confident that there will be redress if they are harmed by the improper use of the information they provide. Otherwise, individuals cannot be expected to be motivated to attend to privacy choices.

Consider the complexities of protecting privacy in this scenario: Alice is the mother of a three-year-old child with a severe chronic illness. She learns all she can about it, buying books online, searching the Web,

¹There are numerous definitions of privacy. Our chief interest here is understanding privacy rights as they relate to the collection and use of personal information, as opposed to other privacy protections that seek to preserve control over, say, one’s physical integrity.

²See the authors’ technical report; dspace.mit.edu/bitstream/1721.1/37600/2/MIT-CSAIL-TR-2007-034.pdf.

In democratic societies, citizens' behavior is unduly restrained if *they fear being watched at every turn.*

and participating in online parent-support social networks and chat rooms. She then applies for a job and is rejected, suspecting it's because a background check identified her Web activities and flagged her as high risk for expensive family health costs.

Such tales are offered to support the argument for Web privacy. Did, say, the online bookstores assert that the titles of Alice's purchases would be kept confidential? Did AOL promise never to release information about her online searches? Did the chat service guard against lurkers in the chat room, recording the names of every participant? A policy regime based on information hiding would focus on these potential acts of data release, perhaps even taking the position that it is Alice's own personal responsibility to inform herself about the privacy policies of Web sites before using their services. This focus is misplaced. The actual harm was caused not by the disclosure of information by the bookseller, AOL, or chat service, but by the decision to deny Alice the job, that is, by the inappropriate, discriminatory, and possibly illegal use of the information. It is quite conceivable that Alice wants to be publicly identified as someone with an interest in her child's illness. Forcing her to hide it to protect herself against improper information use significantly limits her ability to exercise her right to freedom of association. Rather, Alice (and everyone else) should be able to live in an online environment that provides transparent information use and accountability to rules that limit the harmful use of personal information.

COPYRIGHT

Looking into copyright and government surveillance reveals deficiencies in the reliance on information hiding as a policy tool. In the copyright context, information hiding commonly takes the form of digital rights management (DRM). As with personal privacy, locking up information is extremely difficult, and efforts at up-front control over the information flow results in user frustration and substantially imperfect security. This is a lesson that even the most ambitious online businesses have learned. For example, in early 2007, Apple CEO Steve Jobs wrote that DRM has not worked nor is it

ever likely to work [5]. Soon thereafter, Apple changed the way it sells music online by offering a higher-priced version of its download service unencumbered by DRM. Apple now implements a basic form of information accountability. The newly unlocked tracks include the purchaser's name and other personally identifying information. That way, if he or she shares the purchased music with, say, a hundred million closest friends through the Internet, the purchaser could be held accountable.

The Creative Commons, another approach to online copyright protection, likewise does not rely on up-front enforcement of licenses. Rather, its architecture, based on rights expression, not access restriction, recognizes the value of having information flow freely around the Internet but still seeks to impose certain restrictions on how the information is used.

GOVERNMENT DATA MINING

Recent government use of advanced data mining techniques is another example of the deficiency of access-control and collection-limitation approaches to privacy compliance on the Web. Laws that limit access to information do not protect privacy here because so much of the data is publicly available. To date, neither law nor technology has developed a way to address this privacy loophole [2].

Airline passenger screening by law enforcement and national security agencies illustrates the growing complexity of information handling and transfer. Society may be prepared to accept (and even expect) national security agencies to use aggressive data mining techniques over a range of information in order to identify potential terrorism risks. But citizens find it unacceptable to use the same information with the same powerful analytic tools to investigate domestic criminal activity. Therefore, we need rules in the U.S. (and globally) that address the permissible use of certain classes of information, in addition to simple access and collection limitations.

LEGAL FRAMEWORK

The information-accountability framework more closely mirrors the relationship between the law and human behavior than do the various efforts to

Information accountability means that *information usage should be transparent* so it is possible to determine whether a use is appropriate under a given set of rules.

enforce policy compliance through access control over information. As an early illustration of information accountability at work today, consider credit bureaus and their vast collections of personal information. When these databases came on the scene in the consumer financial markets of the 1960s, policymakers recognized the public imperative to protect individual privacy and assure data accuracy, all while maintaining enough flexibility to allow analysis of consumer credit data based on the maximum amount of useful information possible. Under the Fair Credit Reporting Act (enacted 1970) [3], privacy is protected not by limiting the collection of data, but by placing strict rules on how the data may be used. Analysis for the purpose of developing a credit score is essentially unconstrained, but the resulting information can be used only for credit or employment purposes. It cannot be used for marketing and other profiling. Strict penalties are imposed by the FCRA for the breach of these use limitations. Data quality is protected by giving all consumers the right to see the data held about them (transparency). If a user of the data makes a decision adverse to the consumer (such as denial of a loan or rejection of an employment application) the decision must be justified with reference to the specific data in the credit report on which the decision was based (accountability). If the consumer discovers that the data is inaccurate, he or she may demand that it be corrected. Stiff financial penalties are imposed by the FCRA against the credit bureau if it fails to make the appropriate corrections.

The typical consumer appreciates the paradox associated with protecting privacy and other information policy values through increased transparency. As the FCRA illustrates, we achieve greater information accountability only by making better use of the information that is collected and by retaining the data that is necessary to hold data users responsible for policy compliance. The success of this accountability regime for the past 40 years over a very large set of data—credit reports on nearly every adult in the U.S.—makes it a worthy model for considering policy compliance in other large systems.

TECHNICAL ARCHITECTURES

What technical architecture should be required to support information accountability? Our goal in promoting accountability systems is to build into our information infrastructures the technology necessary to make acts of information usage more transparent in order to hold the individuals and institutions who misuse it accountable for their acts. Systems supporting information accountability require three basic architectural features:

Policy-aware transaction logs. In a decentralized system each endpoint must assume the responsibility of recording information-use events that may be relevant to the assessment of accountability to some set of policies.

Policy-language framework. Assessing policy compliance over a set of transactions logged at a heterogeneous set of endpoints by diverse human actors requires a common framework for describing policy rules. Drawing on semantic Web techniques, larger and larger overlapping communities on the Web can develop shared policy vocabularies in a bottom-up fashion. A lack of perfect global interoperability of these policies is not a fatal flaw. Just as human societies learn to cope with overlapping and sometimes contradictory rules, so too are policy-aware systems likely to develop at least partial interoperability [1].

Policy-reasoning tools. Accountable systems must be able to assist users in answering such questions as: Is this data allowed to be used for a given purpose? and Can a given string of inferences be used in a given context, in light of the provenance of the data and the applicable rules? One possible approach to designing accountable systems is to place a series of accountable appliances throughout the system that communicate through Web-based protocols [7]. Accountability appliances would serve as proxies to data sources, mediating access to the data, and maintain provenance information and logs of data transfers. They could also present accountability reasoning in human-readable ways and allow annotation, editing, and publishing of the data and reasoning being presented [6]. This aspect of the accountability and


transparency perspective is closely related to the issue of maintaining provenance for scientific data [4, 11].

CONCLUSION

Alan Westin published his landmark study *Privacy and Freedom* in 1967 [12]. Still in the age of mainframe computers, it set the stage for thinking about privacy over the next three decades. Westin presented what has become a classic definition of privacy, emphasizing the individual's right to control how personal information "is communicated to others." An information-accountability perspective on privacy would reframe this definition, shifting toward the use of any information. Following Westin, we would say that privacy is the claim of individuals, groups, and institutions to determine for themselves when, how, and to what extent information about them is used lawfully and appropriately by others.

Westin's work is essential today for identifying the role of privacy in a free society. However, advances in communications and information technology and the ease of data searching and aggregation have rendered his definition incomplete as a framework for information policy and information architectures that are intended to be policy aware.

Will the new tools and laws we've described here put an end to all privacy invasion, unfair misuse of personal information, copyright infringement, and identity theft? Of course not. Perfect compliance is not the proper standard by which to judge laws or systems that help enforce them. Rather we should ask how to build systems that encourage compliance and maximize the possibility of accountability for violations. We should see clearly that our information-policy goals cannot be achieved by restricting the flow of information alone. While the accountability approach is a departure from contemporary computer and network policy techniques, it is far more consistent with the way legal rules traditionally work in democratic societies.

Contemporary information systems depart from the norm of social systems in the way they seek to enforce rules up front by precluding the possibility of violation, generally through the application of strong cryptographic techniques. In contrast, we follow rules because we are aware of what they are and because we know there will be consequences, after the fact, if we violate them. Technology will better support freedom by relying on these social compacts than by seeking to supplant them. 

REFERENCES

1. Barth, A., Mitchell, J., and Rosenstein, J. Conflict and combination in privacy policy languages. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society* (Washington, D.C., Oct. 28). ACM, New York, 2004, 45–46.

2. Dempsey, J. and Flint, L. Commercial data and national security. *The George Washington Law Review* 72, 6 (Aug. 2004).
3. Fair Credit Reporting Act, 15 U.S.C.1681; www.law.cornell.edu/uscode/15/usc_sup_01_15_10_41_20_III.html.
4. Golbeck, G. and Hendler, J. A semantic Web approach to the provenance challenge. *Concurrency and Computation: Practice and Experience* (2000); www.mindswap.org/~golbeck/downloads/pc.pdf.
5. Jobs, S. *Thoughts on Music* (Feb. 6, 2007); www.apple.com/hotnews/thoughtsonmusic/.
6. Kagal, L., Hanson, C., and Weitzner, D. Integrated policy explanations via dependency tracking. In *Proceedings of the IEEE Workshop on Policies for Distributed Systems and Networks* (June 2–4, 2008).
7. Lunt, T. *Protecting Privacy in Terrorist-Tracking Applications*. Presentation to the Department of Defense Technology and Privacy Advisory Committee (Washington, D.C., Sept. 29, 2003).
8. Samarati, P. Protecting respondent's privacy in microdata release. *IEEE Transactions on Knowledge and Data Engineering* 13, 6 (Nov./Dec. 2001), 1010–1027.
9. Solove, D. *The Digital Person*. New York University Press, New York, 2004.
10. Sweeney, L. K-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness, and Knowledge-based Systems* 10, 5 (2002), 557–570.
11. Szomszor, M. and Moreau, L. Recording and reasoning over data provenance in Web and grid services. In *Proceedings of the International Conference on Ontologies, Databases, and Applications of Semantics 2888* (Catania, Sicily, Italy, 2003), 603–620.
12. Westin, A. *Privacy and Freedom*. Atheneum Press, New York, 1967.

DANIEL J. WEITZNER (djweitzner@csail.mit.edu) is Director of the Massachusetts Institute of Technology Decentralized Information Group and principal research scientist in the MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, and Technology and Society Policy Director of the World Wide Web Consortium.

HAROLD ABELSON (hal@mit.edu) is the Class of 1922 Professor of Computer Science and Engineering at the Massachusetts Institute of Technology, Cambridge, MA.

TIM BERNERS-LEE (timbl@csail.mit.edu) is Director of the World Wide Web Consortium and holds the 3Com Founders chair and is a senior research scientist in the Laboratory for Computer Science and Artificial Intelligence at the Massachusetts Institute of Technology, Cambridge, MA.

JOAN FEIGENBAUM (joan.feigenbaum@yale.edu) is the Grace Murray Hopper Professor of Computer Science at Yale University, New Haven, CT.

JAMES HENDLER (hendler@cs.rpi.edu) is the Tetherless World Professor of Computer and Cognitive Science at Rensselaer Polytechnic Institute, Troy, NY.

GERALD JAY SUSSMAN (gjs@mit.edu) is the Panasonic Professor of Electrical Engineering at the Massachusetts Institute of Technology, Cambridge, MA.

The authors would like to thank Randy Davis and Butler Lampson for their insightful comments on accountability, copyright, and privacy.

The work reported here was conducted at MIT, RPI, and Yale with support from the National Science Foundation Cybertrust Grant (award #04281) and IARPA (award #FA8750-07-2-0031).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2008 ACM 0001-0782/08/0600 \$5.00

DOI: 10.1145/1349026.1349043