STUART S. SHAPIRO

Time to Modernize Privacy Risk Assessment

In 2018, media reports revealed that a company called Cambridge Analytica had harvested data from millions of Facebook users, creating psychometric profiles and models that were then used for political manipulation. For Facebook, it was a high-profile privacy debacle. For the information technology and privacy communities, it was a particularly high-profile wake-up call.

This privacy failure was far too complex to be called simply a breach; it occurred across multiple layers and involved several companies. At the foundation of the scheme was data gleaned from Facebook's "like" button, which Cambridge Analytica used to infer users' personality traits. The company gained access to more people and more information through users' profiles and Facebook friends (i.e., social networks). This unchecked flow of data was enabled by Facebook's privacy policy, the way the platform interacted with third-party apps, and its desire to support social science research. Amazon's Mechanical Turk, a platform for virtual paid piecework, also played a key role, as did various sources of public information, including US Census data. At first glance, the whole mess looks like a textbook example of an emergent property of a complex system: the interactions of multiple actors and systems producing completely unanticipated results.

It's possible that Facebook didn't see the potential for such a disaster brewing in advance because of outdated and inadequate methods for defining and evaluating privacy risks. Despite dizzying socio-technical changes over the past quarter of a century, organizations still rely heavily on assessments of privacy impacts with simplistic forms and functions that are poor matches for the layered complexity of today's technologies. The United States is not alone in this; the data protection impact assessments required by the European Union's General Data Protection Regulation, although an improvement in some respects, are similarly lacking.

As long as this dependence continues, we can expect new, more frequent, and ever-stranger privacy incidents. AI-based decisionmaking tools, for example, which often require large amounts of personal information for training their algorithms, can encode bias in their operation and injure or expose individuals to harm in everything from criminal justice proceedings to benefits eligibility determinations. The Internet of Things raises difficult issues of data aggregation—in which seemingly innocuous data points acquire much greater significance when combined—and of ubiquity, where multiple platforms can create mosaics of individuals' activities. Biometrics, especially facial recognition, create additional potential for persistent surveillance as well as for problematic inference of individual attributes from physiological features. As these technologies develop, public- and private-sector organizations must update their approach to effectively manage privacy risk.

How we got here

In the early 1970s, the public was concerned about the potential implications of modern data processing systems, then standalone mainframes, for civil liberties. The US Department of Health, Education, and Welfare ordered a 1973 report by the Secretary's Advisory Committee

on Automated Personal Data Systems. The committee articulated a set of guidelines called the "Code of Fair Information Practices."

That code formed the basis of the federal Privacy Act of 1974 and prompted the development of numerous, slightly varying, and expanded sets of best practices for protecting privacy. These practices—which typically included consideration of data collection, retention, and use as well as training, security, transparency, consent, access, and redress—were eventually dubbed Fair Information Practice Principles (FIPPs). Around the world, they became the de facto approach to protecting informational privacy. Most privacy statutes and regulations today are built on some version of FIPPs.

Chief among the approaches enabled by FIPPs are Privacy Impact Assessments (PIAs), which bear a name and constitute an approach partly inspired by environmental impact statements and assessments. Echoing these roots, a PIA is both the process of assessing a system's privacy risks and the name of the statement that results. In the evolution of impact assessments, PIAs act as tools for addressing one particular societal value. However, they have been constructed in a way that renders them less about privacy as a human value and more about procedural niceties.

PIAs (and FIPPs) became further embedded in US law and practice when they were required for federal information systems by the E-Government Act of 2002. Today's PIAs largely retain their original form—a set of written questions and answers about each of the FIPPs—and the same function, firmly rooted in identifying potential violations. Because FIPPs provide the principal structure of PIAs, they have become so intertwined with these processes and artifacts that they have together taken on a perception of inseparability. And as the interactions between society and computational technology become more complex, that perceived indivisibility increasingly poses problems.

Problems with the status quo

By using FIPPs to define privacy practices without requiring more expansive analysis, PIAs today maintain a relatively narrow and inelastic view of privacy. This static conception offers a very circumscribed model for imagining and understanding the risks that technological systems could pose to privacy. An ideal risk model describes possible threats, identifies vulnerabilities that might be exploited by them, and lays out what would happen if each exploit were realized, including its likelihood and severity. However, because FIPPs are the risk model utilized by PIAs, today's consideration of privacy risks is largely restricted to violations of FIPPs. Furthermore, the close integration of PIAs and FIPPs, together with FIPPs-based compliance obligations, effectively discourages the use of other privacy risk models and assessment methods.

PIAs also suffer from two problems that have been significantly exacerbated by the evolution of technologies. First, PIAs tend to emphasize description over analysis, which prejudices them toward addressing privacy in a checklist fashion. Second, even when PIAs do explicitly invite discussion of possible privacy risks and potential mitigation strategies, risks are typically construed narrowly. They tend to be first-order problems, issues that might arise as the immediate result of system operation. Potential knock-on effects are seldom considered, nor are potential problems involving indirect cause and effect.

These problems are compounded by the largely procedural nature of FIPPs. Consequences for individuals' privacy are often framed only as possible FIPP violations—as violations of privacy-related procedure—rather than as violations of privacy per se. Many real results from privacy violations, such as embarrassment, lost opportunities, discrimination, physical danger (e.g., stalking), and more, are overlooked.

Another limitation of FIPPs is that they ignore the social context of systems, preventing analysts from considering potential harms originating in the external environment. Finally, FIPPs are so dependent on a system's purpose, without carefully evaluating whether that purpose is fundamentally objectionable, that an unethical purpose can sometimes serve as the basis for satisfied FIPPs. If a system had the purpose of maintaining individual political dossiers on members of the general public, for example, a purely FIPPs-based analysis would take this as its unquestioned starting point and assess each principle relative to that disturbing purpose.

Other risk models and methods

Over the past two decades, other, more capable privacy risk models and assessment methods have been developed that could address the inadequacies of FIPPs and PIAs. Law professor Ryan Calo's dichotomous privacy harms, for example, categorizes all privacy injuries as either subjective or objective, with the former forcing explicit consideration of potential impacts on individuals' mental states—something often ignored by FIPPs-based models. Another model, a taxonomy of privacy developed by law professor Daniel J. Solove, proposes 16 different kinds of privacy problems divided into four groups relating to information collection, information processing, information dissemination, and invasions. This granular categorization enables more precise identification of privacy harms, again forcing more nuanced consideration of potential adverse privacy consequences.

Other risk models address vulnerabilities or threats. The contextual integrity heuristic, developed by information science professor Helen Nissenbaum, aims to identify violations of informational norms, which

can be construed as privacy vulnerabilities. The model is noteworthy for explicitly recognizing the existence of social standards of privacy in various spheres of life, something FIPPs avoid by design. In contrast, frameworks such as LINDDUN, a privacy threat model and methodology, focus on modeling threats at the level of system architecture, considering factors such as potential attempts to link together system elements pertinent to individuals (data, processes, flows, etc.). Although situated at notably different levels, all these models attempt to discover issues that might ultimately affect privacy as experienced by individuals, rather than primarily looking for procedural problems.

Just as there are models beyond FIPPs, there are also new privacy risk assessment methodologies that could replace or complement PIAs. For example, the National Institute of Standards and Technology has developed a Privacy Risk Assessment Methodology that addresses systemic privacy vulnerabilities, defined as "problematic data actions," and consequences, defined as "problems for individuals." This methodology features numeric scores that explicitly estimate the likelihood and severity of privacy consequences. There are also more advanced quantitative options using statistical analysis, such as privacy expert R. Jason Cronk's adaptation of the Factor Analysis for Information Risk framework, as well as a wholly qualitative but rigorous methodology called System-Theoretic Process Analysis for Privacy (STPA-Priv), which uses an approach originally developed to address the safety properties of system control structures. These models and methods have distinct emphases and orientations and could be mixed and matched for best effect.

Cambridge Analytica revisited

Could using other risk models and assessment methods have helped Facebook avert the Cambridge Analytica scandal? It's not clear what, if any, privacy risk analysis was performed by the company, but it's likely that more innovative approaches could have anticipated some of the factors that eventually led to the attempted political manipulation of Facebook users.

One fundamental reason Facebook users were vulnerable is that they were part of a social network. The privacy of any specific Facebook user depends in part on others in their network. This is, of course, part and parcel of being on Facebook, but that connectedness tends to be viewed exclusively as a feature—not a vulnerability. Cambridge Analytica exploited this weakness, termed "passthrough" by information science professors Solon Barocas and Karen Levy, in which the connections of one user enable access to the information of other users.

More recent risk models could have illuminated the threat of manipulation amid the tempestuous political climate of the time. Solove's taxonomy, which considers decisional interference a significant privacy problem, might have suggested the potential consequence of inappropriately influencing voters. And if Facebook had performed an analysis using STPA-Priv, looking at the combination of technology and social context through the lens of hierarchical control and feedback loops, it might even have found the specific control failure scenarios that actually led to abuse.

Using one or several of the models available, Facebook almost certainly could have identified and addressed at least some of the relevant control weaknesses, which might have prevented the debacle. These weaknesses included inadequate monitoring of researcher data use, insufficient restrictions on app access to user and friend profiles, and targeting of users across platforms. That apparently none of these weaknesses provoked concern at the time highlights the importance of adopting more capable approaches to privacy risk.

The complex role of technology in society demands that public and private entities expand their privacy risk assessment toolbox beyond FIPPs and PIAs. In the United States, the Federal Trade Commission should issue guidance for the private sector encouraging the adoption of a broader range of privacy risk models and assessment processes. The National Institute of Standards and Technology, through its privacy engineering program, should develop guidance and tools to assist organizations in comparing and selecting appropriate privacy risk models and assessment methods.

The White House Office of Management and Budget should update and supplement its existing PIA guidance for federal agencies, directing them to actively consider and deploy privacy risk models and assessment methods in addition to FIPPs and PIAs. Finally, the National Science Foundation should encourage and support research explicitly focused on enhancing privacy risk models and assessment methods, consistent with the 2016 National Privacy Research Strategy.

FIPPs and PIAs were innovative in their early days, but the world has changed dramatically. Modern technologies and systems require complementary and flexible approaches to privacy risk that are more likely to discover serious and unexpected issues. FIPPs and PIAs by themselves are no longer enough. Moving forward, organizations need to employ privacy risk assessments that ultimately serve the public interest.

Stuart S. Shapiro *is a principal cyber security and privacy engineer at the MITRE Corporation, a nonprofit company operating federally funded research and development centers.*
Approved for MITRE public release. Distribution unlimited 21-01156-5.