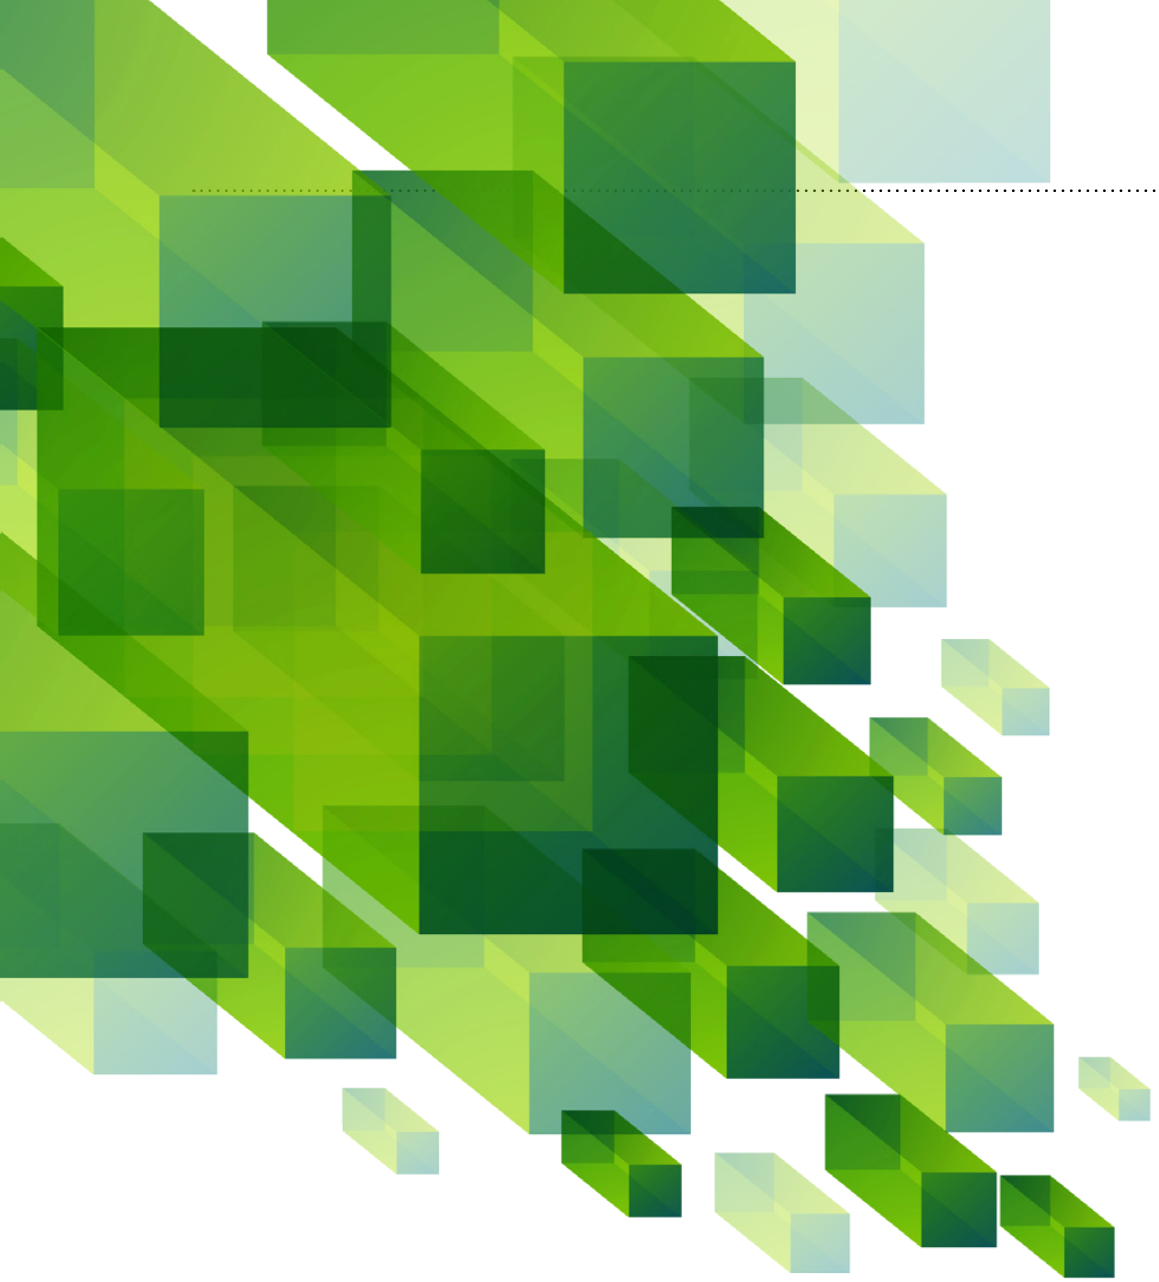

November 2012

Notice and Consent in a World of Big Data

Microsoft Global Privacy Summit Summary Report and Outcomes

Fred H. Cate
Distinguished Professor and C. Ben Dutton Professor of Law,
Maurer School of Law, Indiana University

Viktor Mayer-Schönberger
Professor of Internet Governance and Regulation,
Oxford Internet Institute, University of Oxford

Contents

Preface

One of the great privileges of working at Microsoft is being able to glimpse the future of information technology and envision ways that society can reap the considerable benefits of "Big Data"—the collection, management, and analysis of data on a massive scale. But this privilege also comes with responsibilities, including an obligation to ensure strong information privacy protections. Getting this balance right is crucial not only for Microsoft, but also for policymakers, regulators, industry, educators, and, most importantly, individuals.

Microsoft Corporate Vice President Scott Charney put forth ideas for addressing privacy in a Big Data world earlier this year in marking the 10-year milestone of Microsoft's Trustworthy Computing Initiative.[1] Scott suggested that it was time to evolve the privacy frameworks that govern how we manage and use data. He also proposed a framework that would shift the focus from data collection to responsible data management, and he suggested specific ways to hold organizations accountable.

To advance a global conversation on this important topic, and with the aim of generating shared ideas and new thinking in support of alternative approaches to privacy protection, Microsoft invited privacy stakeholders from 19 countries to a series of five regional meetings between May and August of this year—in Brussels, Singapore, São Paulo, Sydney, and Washington, D.C.—and to a final, capstone event in Redmond, Washington.

On behalf of Microsoft, I want to thank all of the participants who provided their valuable input. This report represents a summary of those discussions.

We are clearly at the start of a journey we all need to take, and I am greatly looking forward to participating in the next phases of this important effort.

Peter Cullen
General Manager, Trustworthy Computing
Microsoft Corporation

[1] See the whitepaper titled "Trustworthy Computing Next" by Scott Charney at http://aka.ms/nextwp.

Introduction

Just over four decades ago, the first information privacy statutes were enacted. After intense discussions in North America and Europe, at the end of the 1970s a number of privacy principles emerged under the concept of Fair Information Practices and later became the foundation for the Organisation for Economic Co-operation and Development (OECD) Privacy Guidelines[2] adopted in 1980. Those principles, which seek to balance the "fundamental but competing values" of "privacy and the free flow of information," form the basis of most privacy legislation around the world. At their core, they require that the processing of personal information be lawful, which in practice means that either the processing is explicitly permissible under law or the individual whose personal data is being processed has—after being informed of the reason, context, and purpose of the processing—given consent.

Intuitively, such an approach makes sense. It empowers individuals so that they—rather than a bureaucratic government agency—can exercise their privacy rights as they see fit. Over the years, and especially in the context of the Internet, this system of "notice and consent," originally intended to be only one of multiple ways through which the lawful processing of personal data can take place, has become the dominant mechanism.

Almost everywhere that individuals venture, especially online, they are presented with long and complex privacy notices routinely written by lawyers for lawyers, and then requested to either "consent" or abandon the use of the desired service. That binary choice is not what the privacy architects envisioned four decades ago when they imagined empowered individuals making informed decisions about the processing of their personal data. In practice, it certainly is not the optimal mechanism to ensure that either information privacy or the free flow of information is being protected.

Equally challenging is the fact that in the age of "Big Data," much of the value of personal information is not apparent at the time of collection, when notice and consent are normally given. Because future uses would require going back to individuals for their amended consent, many future uses that have significant individual and societal benefits might be simply too costly to undertake. Moreover, what used to be a relatively simple relationship between individuals and the processors or users of their personal data has often become complicated as data sets are combined and data processors and users change. That also makes it even harder for individuals to fully grasp the complexity of the situation they are asked to assess. Finally, Big Data is not only big, but also collected and processed so often as to make opportunities to consent an unacceptable burden for most individuals.

[2] The full name is Guidelines on the Protection of Privacy and Transborder Flows of Personal Data.

Taken together, these realities challenge the currently dominant privacy mechanism of notice and consent. They can leave individuals' privacy badly exposed, as individuals are forced to make overly complex decisions based on limited information, while data processors can perhaps too easily point to the formality of notice and consent and thereby abrogate much of their responsibility. At the same time, current privacy mechanisms can unduly interfere with the innovation potential of data use. These challenges require a rational reassessment of the privacy landscape, as well as an evaluation of the optimal mix of mechanisms available to protect information privacy in a world that is beginning to realize the latent value of Big Data.

The Regional Privacy Dialogues

Between May and August of 2012, Microsoft sponsored a series of regional privacy dialogues in Washington, D.C., Brussels, Singapore, Sydney, and São Paulo. These events provided an opportunity for small groups of leading regulators, industry executives, public interest advocates, and academic experts to talk frankly about the role of individual control and notice and consent in data protection today, as well as alternative models that might better protect both information privacy and valuable data flows in the emerging world of Big Data and cloud computing. (The participants at each event are listed in Appendix A.)

Each discussion was moderated by a privacy scholar from the region: Professor Fred H. Cate in Washington, D.C., Professor Viktor Mayer-Schönberger in Brussels, former Australian Privacy Commissioner Malcolm Crompton in Singapore and Sydney, and Professor Nelson Remolina Angarita in São Paulo. The discussions followed the Chatham House Rule, under which participants are welcome to use the information learned but agree not to disclose the source of particular comments.

Despite considerable variety in the five regional discussions, there was significant overlap concerning key issues. There was a widely shared sense that notice and consent either have, or are perceived as having, become the dominant means of data protection. Even in countries in which notice and consent are not the primary data protection tools provided by law, they have nevertheless assumed undue importance in policy debates and popular discussions about data protection. As a result, or perhaps as a cause, ensuring individual control over personal data is widely perceived as the goal of data protection and is often highlighted as such by political leaders and commentators.

Further, there was broad general agreement that privacy frameworks that rely heavily on individual notice and consent are neither sustainable in the face of dramatic increases in the volume and velocity of information flows nor desirable because of the burden they place on individuals to understand the issues, make choices, and then engage in oversight and enforcement. In short, ensuring individual control over personal data is not only an increasingly unattainable objective of data protection, but in many settings it is an undesirable one as well.

The discussions also addressed the advent of Big Data, the increasingly pervasive nature of data collection and use, and the technological developments that expand our capacity to interconnect, analyze, identify, and extract new and unanticipated value from even old or seemingly worthless data as factors that require new approaches to data protection. A key sentiment expressed in all of the discussions is that those new approaches must shift responsibility away from data subjects toward data users and toward a focus on accountability for responsible data stewardship—rather than mere compliance while ensuring that expectations and protection of privacy are preserved.

As to additional mechanisms to ensure privacy, one of the most widely discussed alternatives was focusing more attention on the "use" of personal information rather than on its "collection," given the increasingly pervasive nature of data collection and surveillance, inexpensive data storage and sharing, and the development of valuable new uses for personal data. Many participants were careful to note that focusing on the use of personal data does not mean that there should not be responsibilities or regulation relating to data collection, nor should a focus on data collection in specific or sensitive circumstances be abandoned. Rather, in most situations, a more practical, as well as sensitive, balancing of valuable data flows and more effective privacy protection is likely to be achieved by focusing more attention on appropriate, accountable use.

Many of the dialogues devoted considerable time to what constitutes a "use" of personal data, what uses should be permitted or prohibited (or should require some greater form of authorization—for example, specific affirmative consent), and by what standards these determinations should be made. As many participants noted, the failure to build consensus around the standards that data protection laws should implement currently impedes effective regulation and efforts at international harmonization.

There seemed to be broad agreement that "uses" should include disclosure, but there was uncertainty about whether "uses" should include analysis of data within an institution if the data is not used to make a decision or create new information. Similarly, there was nearly universal agreement that the "harms" or "impacts" that data protection laws should be designed to avoid must include not only physical and financial injury but also broader concepts consistent with protecting privacy as a human right—such as reputational or social harm and the chilling effect of surveillance. But there was little consensus as to precisely which other impacts should be included or how they might be determined.

At the same time, while recognizing privacy as a human right and the need to more clearly define impacts, there was recognition of the need to resolve conflicts with other fundamental rights. For example, privacy can be in conflict with the right to engage in free speech or to live in a society free from the threat of terrorism. The quest needs to be for more effective and efficient protection of privacy, not a weakening of the protections that existing frameworks are intended to provide, even if they do not always do so successfully.

In addition to considering how the standards that guide data protection should be determined, the participants also devoted considerable energy to the question of how those standards should be implemented as a matter of both law and individual entity policy. For example, one approach that emerged emphasized the following: 1) legal standards that identify the rights of data subjects and the responsibilities of data users to the data subjects, 2) adherence by data users to those standards so data subjects and regulators alike can have a high degree of confidence that the data user both intends to comply with and is capable of complying with them, and 3) transparency to data subjects and regulators, an opportunity for individuals to exercise their rights, tools for data users to demonstrate responsible stewardship of personal data, and oversight by regulators. Irrespective of the specific elements, there was broad agreement that implementation should be practical, flexible, and focused on data users ensuring and demonstrating accountability for their responsible use of personal data.

Despite the limits of notice and consent, many participants noted that this mechanism might continue to play a role in the future, even if in a modified form from today. For example, notice may be a key tool for transparency, although this may suggest that disclosure to a regulator or a central, accessible repository might be more efficient than individual notice. Similarly, consent may be necessary for the use of certain types of data or for certain uses of data. Some participants expressed the hope that if notice and consent are reserved for more appropriate uses, individuals might pay more attention when this mechanism is used. Recognizing that notice and consent will have continuing value in certain settings also reflects an "evolutionary" rather than "revolutionary" approach to updating data protection principles, which many of the participants found desirable. As a result, while there was broad agreement in all of the dialogues that merely fine-tuning notice and consent will not provide the sort of new approaches to data protection that are widely thought necessary, this does not suggest that notice and consent should not be improved where possible so that when used they are more effective.

One key element of responsible data stewardship that emerged at all five events was the need for better security to protect personal data against unintended access, loss, alteration, or disclosure. Any new model of data protection must ensure a high degree of confidence that personal data will be appropriately protected. While standards for data security are increasingly being implemented around the globe, there was discussion about the need to ensure that security standards remain flexible, given the constantly evolving nature of security challenges, and that they be more substantive rather than focused on providing notice to people whose data may have been compromised. Again, a key goal of many participants is to shift the responsibility for data protection away from the data subject toward the data user.

A number of discussions touched on enforcement of data protection laws and policies. Participants agreed that enforcement is a critical element of data protection, but some placed special emphasis on enforcement as a way to transform "self-regulation" into "co-regulation," by giving the force of

law to institutional or sectoral privacy undertakings that meet minimum requirements. There was also discussion of the extent to which relying on multinational enforcement mechanisms (such as designated lead enforcement agencies, an international enforcement body, or binding arbitration) might help build cross-border accountability and trust while reducing the costs of enforcement and avoiding duplicative enforcement.

Discussions at all five events addressed the need for greater harmonization and interoperability in data protection across national borders in a way that does not lead to a "race-to-the-bottom, lowest-common-denominator" result. Harmonization and interoperability have assumed even greater importance with the growth of cloud computing and e-commerce, which often involve instant flows of data across geographic boundaries. One advantage of developing new approaches to data protection based on widely shared 21st-century standards of appropriate data stewardship is that it might well lead to more consistent privacy protection across borders and greater harmonization of privacy practices and obligations. Moreover, greater harmonization and interoperability are potentially effective tools for maximizing the scarce resources available for data protection and enforcement, and for ensuring that individuals enjoy commensurate levels of privacy protection—and that their rights can be vindicated affordably and easily—no matter where they travel, browse, or shop, or where their data is stored.

It is always risky to try to summarize rich and varied discussions among talented privacy professionals, but two themes seemed to dominate most of the discussions at all five locations: that society should, to the greatest extent possible, shift responsibility for protecting privacy from the individual data subject to the data user, and that the tools used to do so should be as flexible, efficient, practical, interoperable, and sensitive to competing values and realities as possible to achieve responsible data stewardship.

Microsoft Global Privacy Summit

Following the five regional events, Microsoft invited more than 70 privacy and data protection experts from government, industry, nonprofit organizations, and academia to a global privacy summit in Redmond, Washington. Drawn from 19 countries on five continents, the participants came together to consider the future of data sources and uses and practical steps to enhance privacy protection. Many had participated in the regional discussions, and all received in advance of the summit a summary of the key points from those discussions. (The summit participants are listed in Appendix B.)

Following presentations on innovative new uses of personal data by Craig Mundie, chief research and strategy officer at Microsoft; Leroy Hood, president of the Institute for Systems Biology; Kush Parikh, senior vice president of business development at Inrix; and Kenn Cukier, data editor at the *Economist*, the bulk of the summit was spent in interactive discussions reflecting on the themes

identified in the regional workshops and considering how best to address them in practice.

The participants were able to respond to speakers, pose questions, and interact with one another not only face to face, but also using an interactive tool that allowed participation by everyone and permitted every idea to be captured (anonymously and with consent). To facilitate maximum engagement, discussions took place as a single large group as well as in seven smaller groups, in all cases with professional facilitators and rapporteurs. Every effort was made to ensure not only that all voices were heard and all interjections included, but also that the discussion progressed toward a practical and useful outcome.

After the summit, participants had another two weeks to review the presentations and documents online and add additional comments to the record. In addition, participants had the opportunity to review and comment on this report in draft form. The remainder of this report captures the major themes that emerged during the summit. Given the breadth and depth of the discussions, this report is necessarily selective, and while it is based on the discussions in Redmond, it does not purport to reflect a consensus view of the participants.

Significant Challenges

Participants identified a number of privacy challenges in the near future, but five broad themes emerged:

- There was considerable concern about the need for greater public awareness of privacy issues, increased transparency about the uses of personal data, and more effective education about privacy and the valuable uses of personal data. There was broad agreement that even if data protection systems come to rely less on notice and consent, practical and ethical considerations require that the public and other key stakeholders (including policymakers, regulators, the press, and business) be better informed about data processing activities and the benefits and risks of those activities.

- The pressing need for increased standardization, consistency, and interoperability across data protection laws and practices was emphasized by many participants. While participants also recognized that there are distinctive national and cultural aspects of privacy and stressed the continuing role of national data protection laws, there was a widely shared belief that individuals, societies, and data users can benefit from greater consistency and interoperability across national systems.

- A number of participants noted the values in tension with privacy and therefore stressed the importance of "balance" when protecting privacy and the need to restore the balance between privacy and the free flow of information reflected in the OECD Privacy Guidelines. For exam-

ple, one common refrain was the importance of not suppressing innovation with overly restrictive privacy laws. Also, closely related to the importance of balance was the perceived need for flexibility in data protection regimes, especially in light of rapidly changing technologies and applications, evolving expectations of privacy, and recognition that different types and uses of personal data will inevitably need to be treated differently.

- One theme that became increasingly prevalent as the discussion focused on practical steps for moving forward was the need for clear, specific terms and definitions. Participants noted that as privacy protection increasingly focuses on permitting or preventing certain uses of information or guarding against certain risks or side effects of data use, the more important it is to define uses, risks, and other terms clearly and concretely. This is obviously a challenge and may well be in tension with other themes, but it is nevertheless important to heed.

- The final prevalent theme was the need for an updated or enhanced framework for protecting personal data. The OECD Privacy Guidelines, on which most modern data protection laws are based, was crafted more than 30 years ago, before the advent of the World Wide Web, cloud computing, smart phones, or Big Data. As discussed in greater detail below, while the guidelines seem to have continuing relevance, they are no longer adequate as a guide for 21st-century data protection or as the basis for greater interoperability among national data protection regimes.

Privacy Principles for the 21st Century

To help focus thinking and facilitate practical outcomes, participants at the summit worked in small groups to consider how the OECD Privacy Guidelines might be updated in light of the themes that emerged in the regional discussions. To aid in their task, they were given a draft version of revised principles reflecting the regional discussions. The original 1980 OECD version and the revised version appear below, along with a brief overview of the discussions comparing them. (Appendix C presents the two versions of each principle side by side, in a table.)

1. **Original OECD Version:** *Collection Limitation Principle—There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.*

   Discussion Version: *Collection Limitation Principle—Data should be obtained by lawful and fair means and in a transparent manner. Data should not be collected in a manner likely to cause unjustified harm to the individual unless required by law. "Harm" may include more than physical injury or financial loss.*

Three broad issues arose in connection with the first principle. First, with respect to what is meant by "personal data," there is a growing awareness that many forms of previously unidentifiable data might become personally identifiable in a world of Big Data and advanced analytics. Applying the Collection Limitation Principle to all data seems unworkably broad, but to limit it to data already recognized as "personal" seems too narrow. This theme continued in the discussion of many of the other principles.

Second, when linking collection with fairness, much depends on what "fairness" means. Without a clear definition, using the term seems unnecessarily vague and creates the risk of circular reasoning (i.e., data processing in accordance with these guidelines is fair, but data must be collected by fair means in order to be processed in accordance with these guidelines).

Third, the requirement in the original OECD principle that data be collected, "when appropriate," with the "knowledge or consent of the data subject," seems to ignore the reality of the extraordinary volume of data that is generated today through routine activities and transactions and near-ubiquitous sensors (such as surveillance cameras, location monitoring by smart phones, and embedded computers in cars and other devices). Often, knowledge or consent of data collection in these situations is either nonexistent or likely to be so vague as to be meaningless. No one suggested that knowledge is not important, or that consent may not be appropriate in some settings, but there seems a real risk that the "where appropriate" exception could swallow the entire principle, given today's technology landscape.

2. **Original OECD Version:** *Data Quality Principle—Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.*

   Discussion Version: *Data Quality Principle—Data should be relevant to the purposes for which they are used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.*

   This principle seemed to strike most participants as useful and relevant, with one possible exception. The language "to be used" in the original OECD version could be interpreted as suggesting that the determination as to relevance might need to be made only at the time of collection, with an eye toward intended use. This, of course, is inconsistent with the world of Big Data, in which new uses for data are discovered over time. The principle might be more consistent with 21st-century reality and offer better, continuing protection for personal privacy if the words "to be" were removed, as in the discussion version, and if relevance were evaluated for each use at the time of the use.

3. **Original OECD Version:** *Purpose Specification Principle—The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfillment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.*

   Discussion Version: *Purpose Specification Principle—Organizations should collect and store data only as necessary to serve lawful purposes and for appropriate uses. Information about both purposes and uses should be readily available.*

   Many participants found the Purpose Specification Principle to be largely inconsistent with the ways in which data is used today and will be used in the future, and it is the one principle that many participants suggested might be omitted entirely or at the very least dramatically reshaped. There seemed to be a broadly shared sentiment that there of course should be limits on uses of data, but that those limits need not necessarily be linked to the purposes for which the data was originally collected. Use limits are discussed further under the next principle. Some participants suggested that the Purpose Specification Principle could survive if it is limited to the purpose of the primary (or first) use of personal data and not applied to secondary uses.

4. **Original OECD Version:** *Use Limitation Principle—Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with [the Purpose Specification Principle] except with the consent of the data subject or by the authority of law.*

   Discussion Version: *Use Limitation Principle—Data may not be used if the use is fraudulent, unlawful, deceptive or discriminatory; society has deemed the use inappropriate through a standard of unfairness; the use is likely to cause unjustified harm to the individual; or the use is over the well-founded objection of the individual, unless necessary to serve an over-riding public interest, or unless required by law.*

   Many participants considered the Use Limitation Principle, originally adopted in 1980, unworkably narrow in the 21st century because it restricts uses of data to those required by law or to which the individual has consented. The Use Limitation Principle threatens medical research, fraud prevention, identity verification, credit-worthiness assessment, and many other valuable uses of data, not to mention valuable benefits from Big Data.

   The harder question, as many summit participants noted, is what should replace this outdated principle. There was considerable discussion about restructuring the principle to permit all uses of data other than those meeting certain criteria. But determining those criteria proved quite difficult. Seemingly everyone agreed that uses causing financial or physical injury to the data subject or that involve discrimination or some other illegal act would clearly be prohibit-

ed, but efforts to move beyond the obvious prohibitions rapidly ran into the twin objections of being vague and of potentially ignoring distinct cultural sensibilities. As was the case with the regional dialogues, there seemed a broad willingness to go further than just restricting uses likely to cause "harms," but far less agreement emerged on whether prohibited uses should extend to those causing reputational injury (even more controversial if the data is true) or those causing apprehension or discomfort on the part of the data subject but not otherwise violating the law.

In addition, the discussion version of principle 8, the Accountability Principle, includes a requirement that organizations develop and implement risk assessments. This requirement seems to relate to the Use Limitation Principle and raises the question of what risks a responsible data steward should assess against.

This is clearly an area for future discussion, and one of the most challenging areas in which to be specific but also to find common ground for harmonization of divergent data protection laws.

5. **Original OECD Version:** *Security Safeguards Principle—Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.*

   Discussion Version: *Security Safeguards Principle—Data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.*

   There was widespread agreement that security is a key component of privacy, and while both versions of the Security Safeguards Principle reflect that, many participants suggested that security should be the subject of greater attention and enforcement as a practical matter.

6. **Original OECD Version:** *Openness Principle—There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.*

   Discussion Version: *Openness Principle—There should be a general policy of openness about developments, practices and policies with respect to data. Means should be readily available of establishing the existence and nature of data, and the main purposes of their use, how it will be protected, as well as the identity and usual residence of the data controller.*

   Similarly, there was broad agreement that openness or transparency is a critical element of any data protection system. A number of participants expressed the view that greater trans-

parency should be required if there are fewer opportunities for consent or if personal data can be lawfully collected without consent. This was also seen by some as a key opportunity to help educate data subjects, perhaps by requiring as part of the Openness Principle that information about data subjects' legal rights, and ways to exercise them, be made available along with information about data processing activities. With respect to the "identity and usual residence of the data controller," some participants suggested that this might be expanded to require that data controllers provide information on how they can be contacted over the Internet.

7. **Original OECD Version:** *Individual Participation Principle—An individual should have the right:*

   a) *to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him;*

   b) *to have communicated to him, data relating to him within a reasonable time; at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him;*

   c) *to be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and*

   d) *to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.*

   Discussion Version: *Fairness Principle—Individuals should have the right to be provided reasonable preferences with respect to the use of data relating to them; have a means to obtain redress with respect to data uses relating to them that are deemed to be unfair; and have the right with regard to data relating to them concerning, or used in a manner affecting, employment, health care, financial matters, or legally protected rights:*

   a) *to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him;*

   b) *to have communicated to him, data relating to him within a reasonable time; at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him;*

   c) *to be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.*

   The OECD Individual Participation Principle sparked considerable discussion—not because there seemed to be any opposition to the concept, but rather because of the potential risk that the principle could generate significant burdens for both data processors and data subjects in a world of Big Data. Three broad strands of concern were evident in the discussion.

First, some participants noted that with the exponential growth of data collected about individuals, responding to requests for access to such data could be prohibitively expensive and seemingly of little value if the data is not being used for any significant purpose. Inaccurate data, for example, might be very significant if it is being used to determine eligibility for a job or a government benefit, but it might be less relevant as a small part of a vast data set being used to determine normal spending patterns for fraud detection purposes. One response to this concern would be to focus some or all of the legal obligations created under this principle on data that is used, or is likely to be used, for some significant activity or in a manner affecting a legally protected right, as suggested in the discussion version.

Second, some participants worried about the burden this principle would place both on individuals, who might find it difficult to even know all of the parties with access to their data, and on organizations, which in many cases, due to the distributed nature of data, would have difficulty verifying the identity of individuals and determining which data pertains to them. Some also expressed concern about issues that would arise if data that is being used in a de-identified format must be re-identified purely to respond to a general access request. It is therefore important that this principle not be seen as an excuse for weaker oversight and enforcement or for decreased accountability obligations on the part of the data processor.

Third, participants expressed concern that the principle, at least in its original formulation, is too narrow. By focusing on "individual participation" rather than on fundamental fairness in data processing, the principle as written in 1980 might exacerbate the concern that in a world of Big Data individuals are being asked to do more, when it should be data processors that bear the burden of responsible data stewardship. However, discussions about ways to improve or expand this principle raised concerns about vagueness and the difficulty of defining key terms such as "fair." It was nevertheless clear that while consent should be deemphasized as a primary means of data protection, this should not be confused with deemphasizing individual choices or preferences.

8. **Original OECD Version:** *Accountability Principle—A data controller should be accountable for complying with measures which give effect to the principles stated above.*

Discussion Version: *Accountability Principle—A data controller should be accountable for complying with measures which give effect to the principles stated above, and should be liable for the reasonably foreseeable harm caused by his failure to do so. To be "accountable" means that the organization has developed and implemented risk assessments, policies, processes, and procedures reasonably designed to enforce data usage rules consistent with these Principles and to ensure that data is effectively secure. Accountable organizations must be able to, upon request of a regulator, demonstrate how they are being accountable. Enforcement of data protection laws should achieve*

effective compliance with these principles and applicable law, while minimizing the burden on individuals or lawful information flows.

Participants broadly supported this principle, although some felt that the original version does not go far enough. A number of participants noted that the original appears to focus on compliance when, in fact, it should focus more on responsible data stewardship and the broad mechanisms that data processors can use to ensure compliance and demonstrate it to regulators and the public, as the discussion version does. Accountability is already the subject of considerable attention, and the discussion in Redmond suggested that the attention is justified.

These concerns and thoughts reflect the major streams of a rich and far-ranging debate. It appeared obvious to the participants that the unique challenges presented by Big Data, as well as the complexity and multitude of social interactions online, have created an urgent need to adjust information privacy regulations, and the principles that underlie them, to meet the needs of a new era. The goal is to make privacy protections more effective and efficient as well as to revisit the balance between privacy and information flows in a world of not only vastly more data, but also more and rapidly changing valuable uses of that data.

In important ways, participants were in agreement on the general direction of such adjustments. Much more, of course, needs to be done—and swiftly, given the pace of technological change—to ensure that individuals remain protected and data processors embrace their responsibilities while innovation is not artificially constrained. Looking at how "use" can be defined, and acceptable uses delineated from harmful ones, is an evident next step. This is but the beginning of a crucial debate about how we want our digital future to be.

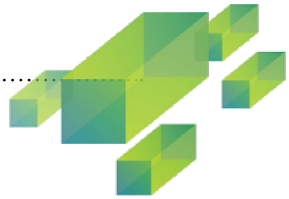

Appendix A
Regional Privacy Dialogues: List of Participants

Note: Participants were invited as individuals; institutional affiliations are listed for identification purposes only.

Washington, D.C., United States, May 9, 2012

Joseph Alhadeff, Oracle

Howard Beales, George Washington University

Julie Brill, Federal Trade Commission

Justin Brookman, Center for Democracy & Technology

Fred H. Cate, Indiana University

Peter Cullen, Microsoft

Pam Dixon, World Privacy Forum

Marc Groman, Network Advertising Initiative

Maneesha Mithal, Federal Trade Commission

Stuart Pratt, Consumer Data Industry Association

Peter Swire, Ohio State University

Hilary Wandall, Merck

Danny Weitzner, White House Office of Science and Technology Policy (now at Massachusetts Institute of Technology)

Joel Winston, Hudson Cook, LLP

Brussels, Belgium, July 12, 2012

Marty Abrams, Centre for Policy Information Leadership

Joseph Alhadeff, Oracle

Rosa Barcelo, European Commission

Marie-Helène Boulanger, European Commission

Dan Cooper, Covington & Burling

Peter Cullen, Microsoft

Nicolas Dubois, European Commission

Peter Hustinx, European Data Protection Supervisor

Christopher Kuner, International Chamber of Commerce

Sophie Kwasny, Council of Europe

Viktor Mayer-Schönberger, University of Oxford

Kostas Rossoglou, BEUC, The European Consumer Organisation

Scott Taylor, Hewlett-Packard Development Company

Cecilia Verkleij, European Commission

Republic of Singapore, August 6, 2012

Edgardo Angara, Senate of the Philippines

Yuanyuan Chen, National University of Singapore

Seng Choon Seah, Consumer Association of Singapore

Malcolm Crompton, Information Integrity Solutions

Benedict Hernandez, Business Processing Association of the Philippines

Tamai Katsuya, University of Tokyo

Lim Kian Kim, Singapore Cloud Forum

Li-Na Koh, Infocomm Development Authority of Singapore

Brendon Lynch, Microsoft

Nguyen Manh Quyen, Vietnam e-Commerce & IT Agency

Abu Bakar Munir, University of Malaya

Teow Hin Ngair, Singapore IT Federation

Vu Quoc Thanh, Vietnam Information Security Association

Koh See Khiang, Baker & McKenzie

Nakorn Serirak, Formerly from the Prime Minister's Office of the Official Information Commission, Thailand

Sudaryatmo, Indonesian Consumer Protection Foundation

Lisa Watson, Direct Marketing Association of Singapore

John Yong, Infocomm Development Authority of Singapore

Sydney, Australia, August 8, 2012

David Batch, Commonwealth Bank of Australia

Andrea Beatty, HWL Ebsworth Lawyers

Gary Blair, Commonwealth Bank of Australia

Malcolm Crompton, Information Integrity Solutions

Ray Delany, New Zealand Institute of IT Professionals

Richard Glenn, Australian Government Attorney-General's Department

Patrick Gunning, King & Wood Mallesons

Vikram Kumar, InternetNZ

Brendon Lynch, Microsoft

Matthew Minogue, Australian Government Attorney-General's Department

Timothy Pilgrim, Office of the Australian Information Commissioner

Paul Roth, University of Otago

Jodie Sangster, Association for Data-Driven Marketing & Advertising

Peter Timmins, Open and Shut blog

Nigel Waters, Australian Privacy Foundation; Privacy International

Anthony Wong, AGW Consulting Pty Ltd

São Paulo, Brazil, August 20, 2012

H.D. Gonzalo Arenas, House of Representatives, Chile

Ana Brian, University of the Republic, Uruguay

Marcelo Vieira de Campos, Former Secretary of State at the Ministry of Justice and Senate Special Adviser, Brazil

Arlene Gonzalez, Data Protection Agency, Costa Rica

Bruno Magrani, Getulio Vargas Law School, Brazil

Silvia Melchior, Melchior & Micheletti Law Firm, Brazil

Lina Ornelas, Mexico's Federal Institute for Access to Information and Data Protection (IFAI)

Pablo Palazzi, Universidad de San Andrés, Argentina

Oscar Puccinelli, Universidad Nacional de Rosario, Argentina

Carolina Purdas, Baker & McKenzie

Enrique Rajevic, Consejo para la Transparencia, Chile

Nelson Remolina Angarita, La Universidad de los Andes, Colombia

Rodrigo Rojas, Abdala & Co. Attorneys, Chile

Maria José Viega, Agencia para el Desarrollo del Gobierno de Gestión Electrónica y la Sociedad de la Información y del Conocimiento (AGESIC)

Diego Zegarra, Benites, Forno & Ugaz Law Firm, Peru

Appendix B
Global Privacy Summit: List of Participants

Note: Participants were invited as individuals; institutional affiliations are listed for identification purposes only.

Redmond, Washington, September 12, 2012

Marty Abrams, Centre for Information Policy Leadership

Joseph Alhadeff, Oracle

Maria Alvarez, Microsoft

Gonzalo Arenas, House of Representatives, Chile

Elizabeth Arguello, Secretary of Economy, Mexico

Andrea Beatty, HWL Ebsworth Lawyers

Ana Brian, University of the Republic, Uruguay

Marcelo Vieira de Campos, Former Secretary of State at Ministry of Justice and Senate Special Adviser, Brazil

Fred H. Cate, Indiana University

Scott Charney, Microsoft

Yuanyuan Chen, National University of Singapore

Malcolm Crompton, Information Integrity Solutions

Dan Cooper, Covington & Burling

Stanley Crosley, Indiana University Center for Law, Ethics, and Applied Research in Health Information

Peter Cullen, Microsoft

Ray Delaney, Institute of IT Professionals (formerly the New Zealand Computer Society)

Marc Davis, Microsoft

Elizabeth Denham, Office of the Information & Privacy Commissioner, British Columbia

Pam Dixon, World Privacy Forum

Makarim Edmon, Faculty of Law, University of Indonesia

Jeffrey Friedberg, Microsoft

Jennifer Glasgow, Acxiom

Arlene González, Data Protection Agency, Costa Rica

Marc Groman, Network Advertising Initiative

Leslie Harris, Center for Democracy & Technology

Peter Haynes, Microsoft

Mike Hintze, Microsoft

Tamai Katsuya, Research Center for Advanced Science and Technology, University of Tokyo

See Khiang Koh, Baker & McKenzie

David Ku, Microsoft

Yu-Chuang Kuek, Yahoo!

Sophie Kwasny, Council of Europe

Brendon Lynch, Microsoft

Bruno Magrani, Getulio Vargas Law School, Brazil

Nguyen Manh Quyen, Vietnamese Ministry of Industry and Trade

Viktor Mayer-Schönberger, University of Oxford

Silvia Melchior, Melchior & Micheletti Law Firm, Brazil

Christopher Millard, Queen Mary, University of London

Teow Hin Ngair, Singapore IT Federation

Carolyn Nguyen, Microsoft

Alfonso Oñate, Mexico's Federal Institute for Access to Information and Data Protection (IFAI)

Lina Ornelas, Mexico's Federal Institute for Access to Information and Data Protection (IFAI)

Pablo Palazzi, Universidad de San Andrés, Argentina

Oscar Puccinelli, Universidad Nacional de Rosario, Argentina

Nelson Remolina Angarita, La Universidad de los Andes, Colombia

Rodrigo Rojas, Abdala & Co. Attorneys, Chile

Paul Roth, University of Otago

Nakorn Serirak, Freedom of Information expert, Thailand

David Simon, Association for Data-Driven Marketing & Advertising

Sudaryatmo, Indonesian Consumer Protection Agency

Scott Taylor, Hewlett-Packard Development Company

Julie César Vega, Asociación Mexicana de Internet (AMIPCI)

Maria José Viega, Agencia para el Desarrollo del Gobierno de Gestión Electrónica y la Sociedad de la Información y del Conocimiento (AGESIC)

Hilary Wandall, Merck

Lisa Watson, Direct Marketing Association of Singapore

Danny Weitzner, Massachusetts Institute of Technology

Joel Winston, Hudson Cook, LLP

Anthony Wong, AGW Consulting Pty Ltd

Mark Young, Covington & Burling

Diego Zegarra, Benites, Forno & Ugaz Law Firm, Peru

Appendix C
OECD Privacy Principles and Discussion Versions

OECD Version	Discussion Version
1. Collection Limitation Principle There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.	**1. Collection Limitation Principle** Data should be obtained by lawful and fair means and in a transparent manner. Data should not be collected in a manner likely to cause unjustified harm to the individual unless required by law. "Harm" may include more than physical injury or financial loss.
2. Data Quality Principle Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.	**2. Data Quality Principle** Data should be relevant to the purposes for which they are used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.
3. Purpose Specification Principle The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfillment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.	**3. Purpose Specification Principle** Organizations should collect and store data only as necessary to serve lawful purposes and for appropriate uses. Information about both purposes and uses should be readily available.
4. Use Limitation Principle Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with Paragraph 9 [3] except: a) with the consent of the data subject; or b) by the authority of law.	**4. Use Limitation Principle** Data may not be used if the use is fraudulent, unlawful, deceptive or discriminatory; society has deemed the use inappropriate through a standard of unfairness; the use is likely to cause unjustified harm to the individual; or the use is over the well-founded objection of the individual, unless necessary to serve an over-riding public interest, or unless required by law.
5. Security Safeguards Principle Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.	**5. Security Safeguards Principle** Data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.

OECD Version	Discussion Version
6. Openness Principle There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.	**6. Openness Principle** There should be a general policy of openness about developments, practices and policies with respect to data. Means should be readily available of establishing the existence and nature of data, and the main purposes of their use, how it will be protected, as well as the identity and usual residence of the data controller.
7. Individual Participation Principle An individual should have the right: a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him; b) to have communicated to him, data relating to him within a reasonable time; at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him; c) to be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and d) to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.	**7. Fairness Principle** Individuals should have the right to be provided reasonable preferences with respect to the use of data relating to them; have a means to obtain redress with respect to data uses relating to them that are deemed to be unfair; and have the right with regard to data relating to them concerning, or used in a manner affecting, employment, health care, financial matters, or legally protected rights: a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him; b) to have communicated to him, data relating to him within a reasonable time; at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him; c) To be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.
8. Accountability Principle A data controller should be accountable for complying with measures which give effect to the principles stated above.	**8. Accountability Principle** A data controller should be accountable for complying with measures which give effect to the principles stated above, and should be liable for the reasonably foreseeable harm caused by his failure to do so. To be "accountable" means that the organization has developed and implemented risk assessments, policies, processes, and procedures reasonably designed to enforce data usage rules consistent with these Principles and to ensure that data is effectively secure. Accountable organizations must be able to, upon request of a regulator, demonstrate how they are being accountable. Enforcement of data protection laws should achieve effective compliance with these principles and applicable law, while minimizing the burden on individuals or lawful information flows.