



August 15, 2014

Federal Trade Commission
Office of the Secretary
Room H-113 (Annex X)
600 Pennsylvania Avenue NW
Washington, DC 20580

Re: Big Data: A Tool for Inclusion or Exclusion? Workshop, Project No. P145406

Dear Sir/Madam:

The Future of Privacy Forum (FPF) is a think tank seeking to advance responsible data practices and includes leaders in business, academia, and consumer advocacy. In response to the Federal Trade Commission's (FTC) public workshop exploring the use of big data and its impact on American consumers,¹ including low income and underserved consumers, FPF submits these Comments to suggest how organizations can explore the benefits of big data to themselves, consumers, and society at large.

As FPF explained in our submission to the NTIA on how big data impacts the White House Consumer Privacy Bill of Rights, the challenges and opportunities presented by big data will require organizations to develop new accountability mechanisms to ensure that they are managing personal information in responsible – and ethical – ways.² (Our full comments are attached as Appendix C.)

I. Assessing the Benefits of Big Data

Any successful approach to big data must recognize the different ways that data can now be used. While many uses of data present no new difficulties or privacy concerns, the more challenging uses of big data should be guided by a risk-benefit analysis that takes into account exactly how the benefits of big data will be distributed. So far, our procedural frameworks are largely focused on traditional privacy risks and assessing what measures can be taken to mitigate those risks. In 2010, for example, the Department of Commerce's Internet Policy Task Force endorsed the use of privacy impact assessments (PIAs) both to help organizations decide whether it is appropriate to engage in innovative data uses and to identify alternative approaches that

¹ Fed. Trade Comm'n, Big Data: A Tool for Inclusion or Exclusion?, <http://www.ftc.gov/news-events/events-calendar/2014/09/big-data-tool-inclusion-or-exclusion> (last visited Aug. 15, 2014).

² NTIA, Big Data and Consumer Privacy in the Internet Economy, Request for Public Comment, 79 Fed. Reg. 32,714 (June 5, 2014).

could reduce relevant privacy risks.³ However, the benefits that come with data use should also inform any analysis. Human research review boards, for example, have long taken into account anticipated benefits and even the importance of any knowledge that could result from human testing and basic research.⁴

Organizations and privacy professionals have become experienced at evaluating risk, but they should also engage in a rigorous data benefit analysis in conjunction with traditional privacy risks assessments. FPF suggests that organizations could develop procedures to assess the “raw value” of a data project, which would require organizations to identify the nature of a project, its potential beneficiaries, and the degree to which those beneficiaries would benefit from the project. This value would be discounted by the probability that a big data benefit can actually be accomplished. While there are no definitive rules of what probability of benefit is needed to overcome a presumption against exposing a beneficiary to privacy risk, it is clear that the mere assertion that a product or service can be improved is inadequate; yet demanding organizations provide proof beyond any doubt would be an impossible standard. Further, any analysis would need to take into account different community standards and expectations, as well as an inherent degree of subjectivity in assessing the relative merits of various projects.

FPF’s initial guidance on this type of analysis is attached as Appendix A. Subjective challenges and limitations should not be used as an excuse not to further explore how benefits can be weighed. The federal government has been a leader in this respect. The White House Office of Management and Budget, for example, requires federal agencies to engage in comprehensive cost-benefit analysis. This analysis recognizes that even where values are not computable, “efforts to measure it can produce useful insights even when the monetary values of some benefits or costs cannot be determined.”⁵ Although a more nimble process is needed for fast moving private sector activity, the concept of weighing benefits and risks is well established.

II. Using Big Data to Fight Discrimination and Empower Groups

As the White House Big Data Report recognizes, there is sometimes potential tension between socially beneficial and privacy impactful uses of information in everything from educational technology to consumer generated health data. But the Report also highlights the many positive uses of big data, such as identifying health risks at an early stage, creating efficiencies in energy distribution, and uncovering fraud through predictive analytics.⁶

FPF in partnership with the Anti-Defamation League has prepared a discussion draft paper that looks at how businesses, governments, and civil society organizations are leveraging data to protect and empower vulnerable groups. The draft FPF-ADL report is attached as Appendix B.

³ U.S. DEP’T OF COMMERCE, INTERNET POL’Y TASK FORCE, COMMERCIAL DATA PRIVACY AND INNOVATION IN THE INTERNET ECONOMY: A DYNAMIC POLICY FRAMEWORK 34-35 (2010), <http://www.commerce.gov/sites/default/files/documents/2010/december/iptf-privacy-green-paper.pdf>.

⁴ 45 CFR § 46.111.

⁵ OFFICE OF MANAGEMENT & BUDGET, CIRCULAR NO. A-94 REVISED, MEMORANDUM FOR HEADS OF EXECUTIVE DEPARTMENTS AND ESTABLISHMENTS, GUIDELINES AND DISCOUNT RATES FOR BENEFIT-COST ANALYSIS OF FEDERAL PROGRAMS (Oct. 29, 1992), http://www.whitehouse.gov/omb/circulars_a094.

⁶ WHITE HOUSE EXEC. OFFICE OF THE PRESIDENT, BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES (May 2014), http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.

Big data can be used to provide access to job markets, to uncover discriminatory practices, and to develop new tools to improve education and provide assistance to those most in need. The report's case studies look at different ways data can be used to see beyond bias to provide new opportunities. For example, the state of New York recognized that clinicians and researchers are faced with incomplete and inadequate information about the health status of lesbian, gay, bisexual, and transgender (LGBT) groups. In response, the state is developing a coordinated, multi-agency effort to strengthen data collection about the state's LGBT community in order to develop more tailored approaches to reduce health disparities among LGBT individuals.⁷

Big data also provides a needed tool to shine a light on discriminatory practices and increase transparency. For, example, the implementation of the electronic information portals by the Equal Employment Opportunity Commission streamlines the collection and analysis of government-agency data, improving the capacity to spot trends and uncover discrimination across the federal government. The FPF-ADL report provides additional examples.

III. The Value of Big Data Review Processes

In some cases, analyzing the benefits of big data and maximizing the benefits that can come from innovative data projects will require new internal review processes.⁸ Moreover, industry increasingly faces ethical considerations over how to minimize data risks while maximizing benefits to all parties.⁹

A documented review process could infuse ethical considerations into how big data is used without requiring radical changes to business practices or placing undue limitations on innovation.¹⁰ Institutional Review Boards (IRBs), which are the chief regulatory response to decades of questionable ethical decisions in the field of human subject testing, provide a useful precedent for focusing on good process controls as a way to address potential privacy concerns. While IRBs play a key role for research subject to the Common Rule and thus, would be inappropriate for wholesale use in the wide range of big data decision-making,¹¹ they could provide inspiration for how projects can be evaluated based on prevailing community standards and assessments of risks and benefits, particularly in cases involving greater privacy risks.

Traditional IRBs provide only one example of what this might look like, however. Privacy officers and other internal processes, such as privacy impact assessments and review boards, can provide an important outlet for evaluating certain categories of data uses. Moving forward, broader big data ethics panels could provide a commonsense response to public concerns about data misuse. While these institutions could provide a further expansion of the role of privacy professionals within organizations, they might also provide a forum for a diversity of viewpoints inside and out of organizations. Ethics reviews could include members with different

⁷ Press Release, NY State Office of Children and Family Services, Governor Cuomo Announces Multi-Agency State Effort to Address LGBT Disparities (July 23, 2014), http://www.ocfs.state.ny.us/main/view_article.asp?ID=833.

⁸ See Ryan Calo, *Consumer Subject Review Boards: A Thought Experiment*, 66 STAN. L. REV. ONLINE 97 (2013).

⁹ danah boyd, *What Does the Facebook Experiment Teach Us?*, MEDIUM (July 1, 2014), <https://medium.com/message/what-does-the-facebook-experiment-teach-us-c858c08e287f>.

¹⁰ Matthew Salganik, *After the Facebook Emotional Contagion Experiment: A Proposal for a Positive Path Forward*, FREEDOM TO TINKER (July 7, 2014), <https://freedom-to-tinker.com/blog/mjs3/after-the-facebook-emotional-contagion-experiment-a-proposal-for-a-positive-path-forward/>.

¹¹ See Calo, *supra* note 8.

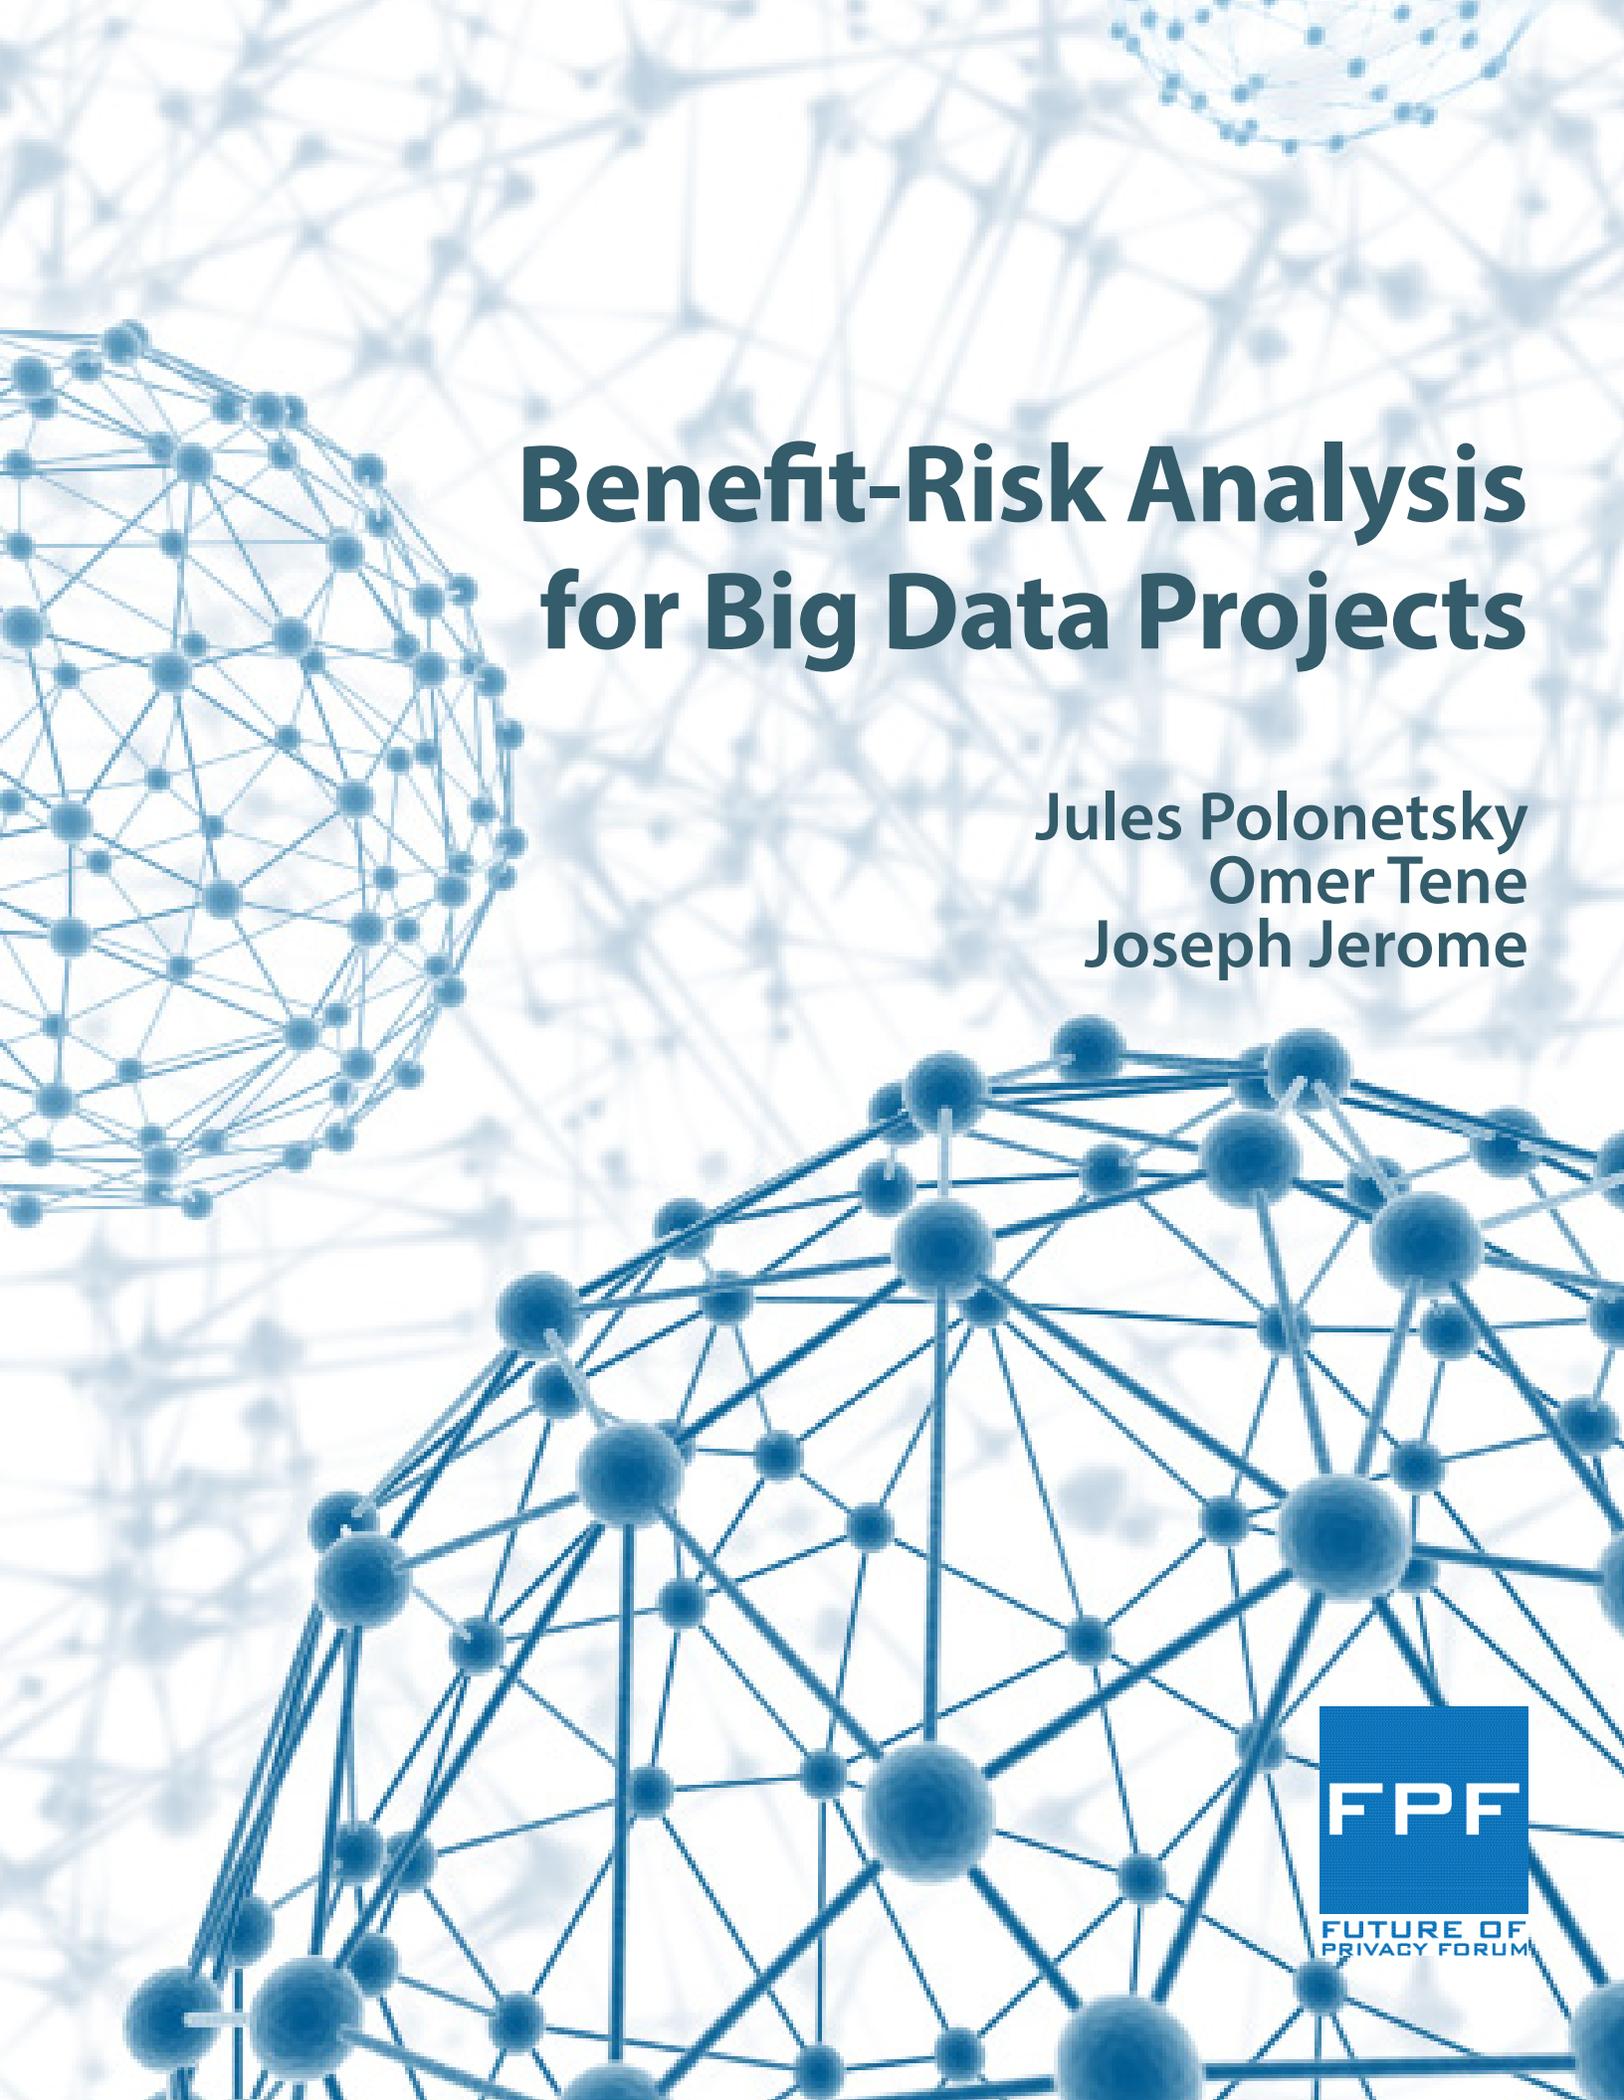
backgrounds, training, and experience, and could seek input from outside actors including consumer groups and regulators. It is likely that these review boards will vary from organization to organization, and standardization would present a unique set of challenges, but organizations should be encouraged to develop structures and personnel to grapple with big data and provide an important check on potential data misuse.

Big data is increasingly serving as an important catalyst for conversations about privacy and ethics. It will require organizations to create new processes and innovative tools to engender trust and address privacy, and FPF looks forward to further engagement and collaboration with the FTC on these tasks.

Sincerely,

Jules Polonetsky
Director and Co-Chair
Future of Privacy Forum

Appendix A



Benefit-Risk Analysis for Big Data Projects

Jules Polonetsky
Omer Tene
Joseph Jerome

FPF

FUTURE OF
PRIVACY FORUM

Benefit-Risk Analysis for Big Data Projects

Jules Polonetsky
Omer Tene
Joseph Jerome

August 2014



About the Future of Privacy Forum

Future of Privacy Forum (FPF) is a Washington, DC based think tank that seeks to advance responsible data practices. The forum is led by Internet privacy experts Jules Polonetsky and Christopher Wolf and includes an advisory board comprised of leading figures from industry, academia, law and advocacy groups.

To learn more about FPF, please visit www.futureofprivacy.org

Introduction: This analysis provides guidance for organizations in their weighing of the benefits of new or expanded data processing against attendant privacy risks.

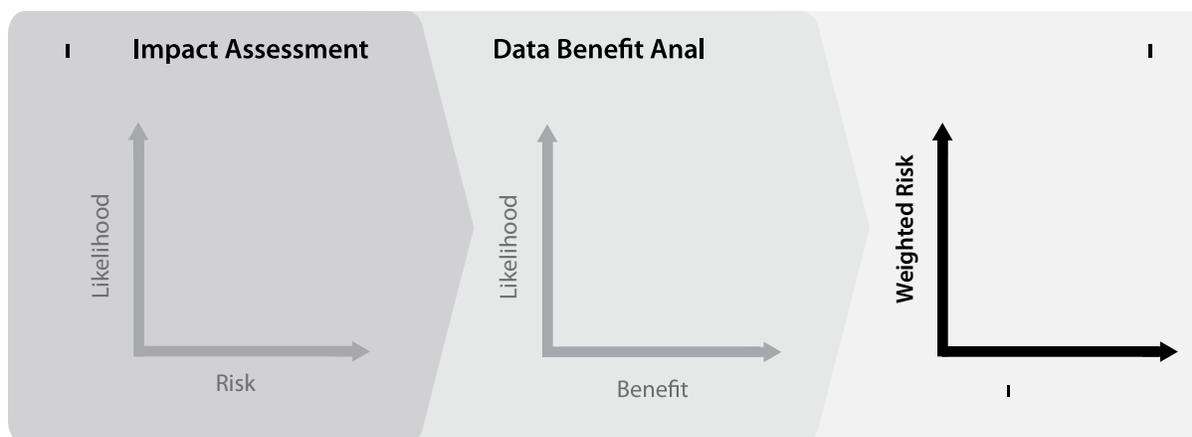
When responsible organizations identify new ways to process data, for example, when launching a new program, product, system or service, they utilize **Privacy Impact Assessments (PIA)** to conduct a systematic analysis to identify and address privacy issues. Current PIA practice includes detailed frameworks to help privacy professionals understand and quantify privacy risks.¹ Yet accounting for *risks* is only part of a balanced value equation. Decision-makers must also assess, prioritize, and to the extent possible, quantify a project's *benefits* in order to understand whether assuming the risk is ethical, fair, legitimate and cost-effective.

The phenomenon of “Big Data” exacerbates the tension between potential benefits and privacy risks by upping the ante on both sides of the equation. On the one hand, big data unleashes tremendous benefits not only to individuals but also to communities and society at large, including breakthroughs in health research, sustainable development, energy conservation and personalized marketing.² On the other hand, big data introduces new privacy and civil liberties concerns including high-tech profiling, automated decision-making, discrimination, and algorithmic inaccuracies or opacities that strain traditional legal protections.³

Decision-makers need to engage in a **Data Benefit Analysis (DBA)**.

This document offers decision-makers a framework for a reasoned analysis to balance big data benefits against privacy risks. This process of identifying both benefits and risks is grounded in existing law. The Federal Trade Commission weighs benefits to consumers when evaluating the *unfairness* of business practices under Section 5 of the Federal Trade Commission Act. Similarly, the European Article 29 Data Protection Working Party applied a balancing test in its opinion interpreting the *legitimate interest* clause of the European Data Protection Directive.⁴ The White House Office of Science and Technology Policy, which has recently studied the social and technical ramifications of big data, recognized the need to strike an appropriate balance between new opportunities and individual values.⁵

Structures and processes for sound benefit analysis are already well established. For example, in 1992, the White House Office of Management and Budget (OMB) issued guidelines for cost-benefit analysis of federal government programs and projects.⁶ The OMB stressed that the criterion for deciding whether a government program can be justified is net present value, which is “computed by assigning monetary values to benefits and costs, discounting future benefits and costs using an appropriate discount rate, and subtracting the sum total of discounted costs from the sum total of discounted benefits.” The OMB’s guidance recognizes that some benefits may not be computable, but efforts to measure value can nevertheless produce useful insights. The same holds true with big data projects.



Privacy Impact Assessment (PIA)

What is a Privacy Impact Assessment (PIA)?

A PIA is a decision-making tool used to identify and mitigate privacy risks at the beginning and throughout the development life cycle of a program, product, system or service.⁷ While a formalized review process is not necessary for every use of data, particularly if the data is neither sensitive nor identifiable, a PIA process helps organizations understand what personal information they are collecting, how it will be used, stored, accessed and shared, and how privacy risks can be mitigated.⁸



PIA Goals

1) IDENTIFY privacy risks arising from the collection, storage, or dissemination of information in a potentially identifiable form.

2) EVALUATE compliance obligations and possible ways to mitigate privacy risks.

Many organizations have gained experience incorporating PIA into project management. A PIA is a necessary and proactive feature of managing risk in a responsible organization.

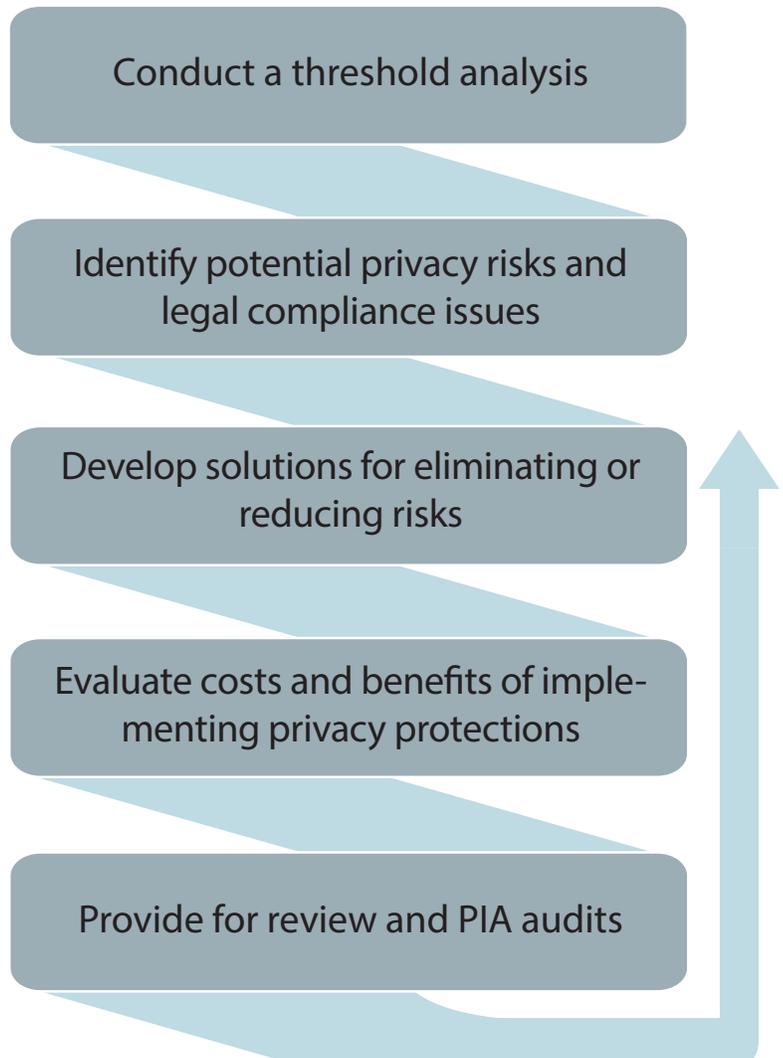
At its core, a PIA requires a value judgment to be made concerning the estimated **level of privacy risk** and the **likelihood** that such risk would materialize. In addition, a PIA involves determining how an organization can best employ risk mitigation tools to comply with privacy principles, generally captured in the **Fair Information Practice Principles (FIPPs)**, as well as with individuals' expectations of privacy.

Mitigation measures can include policies around notice and choice, data minimization, or limited retention of data. A PIA allows an organization to adjust its project to avoid using certain types of data or to further aggregate information being used in a project.



PIA Steps

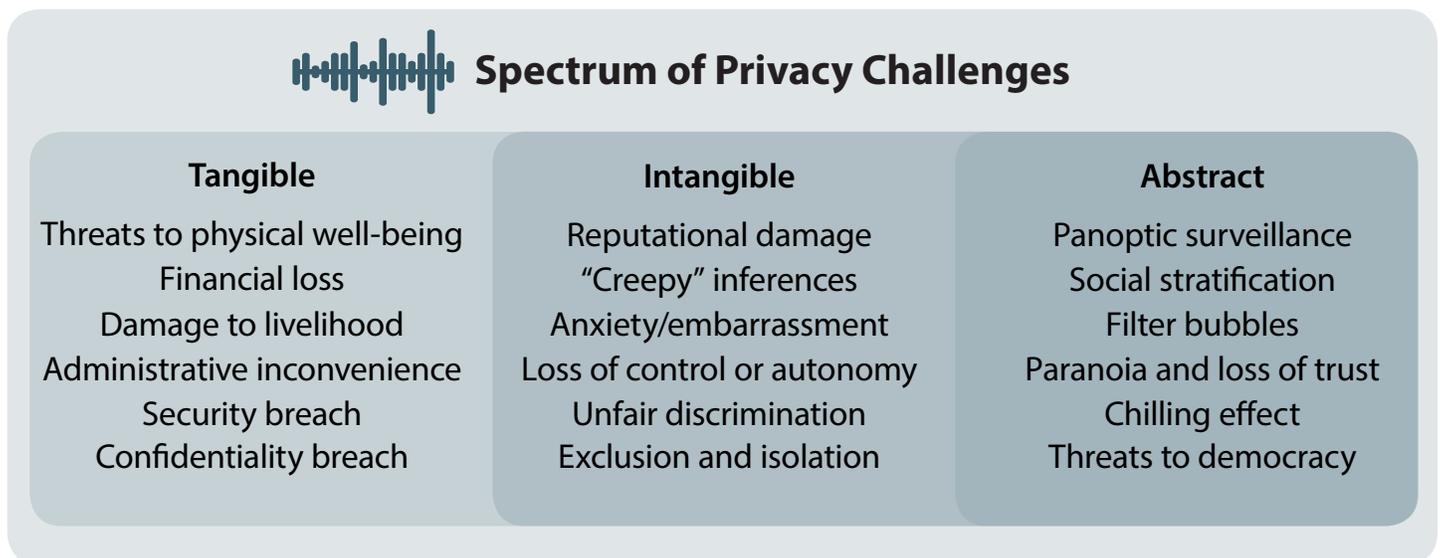
A considerable body of knowledge has been created to help guide organizations through the implementation of a PIA. While specific PIA policies and procedures vary depending upon the nature of an organization and scope of data use, a typical PIA process consists of the following stages:⁹



Privacy Risks

Organizations have come to realize that privacy risk, sometimes conceptualized as *privacy harm*,¹⁰ comes in different flavors. Various frameworks have developed to help organizations categorize privacy risk. Daniel Solove’s taxonomy classifies privacy risks into four categories – information collection, information processing, information dissemination, and invasion – which, in turn, are broken out into 16 sub-categories.¹¹ Richard Thomas distinguishes between material privacy harms, moral privacy harms, and broader democratic and societal harms.

Big data presents new challenges impacting the entire risk spectrum. It accentuates not only the traditional tangible privacy harms but also the more abstract, ethical challenges requiring businesses and governments to make weighty value choices. Existing risk assessment frameworks are geared to identify and address tangible harms, such as financial loss or security vulnerabilities.



Yet big data situations require a broader view of risk as well as additional analysis of a project’s ethical implications. As the risk taxonomy above suggests, some of the new risks commonly associated with big data are not easily mapped to any traditional recognizable harms. Other concerns are that data analysis could permit new forms of unfair discrimination, stigmatization and narrowcasting. All of these new concerns must also be incorporated into a PIA, which may therefore require careful consideration of a project’s abstract or unintended consequences.

Many new uses of data conducted by organizations continue to be routine, involving uses of data that do not create new risks, or uses that are subject to well-defined measures for risk elimination. The analysis documented here is required when a minimum threshold is surpassed. Organizations must have personnel and processes in place to spot new issues and concerns and select those issues that should be subjected to further **PIA** and **Data Benefit Analysis**.

Risk Mitigation Strategies

Traditionally, organizations mitigated privacy risks by operationalizing the Fair Information Practice Principles (FIPPS), including enhancing notice and choice or limiting data retention. Alas, the onset of big data practices has introduced formidable challenges to some of these fundamental principles, including the scope of the framework (often addressed by defining the term “personally identifiable information” (PII)), the concepts of data minimization, purpose limitation and consent, and the right of individual access.¹² This has required policymakers and professionals to develop new privacy solutions to augment traditional tools. These enhanced solutions address the broader categories of privacy risks that are created by big data.¹³



Enhanced Transparency:

Like any interpretative process, big data analysis is prone to errors, inaccuracies and bias.¹⁴ Consequently, organizations should provide more transparency into their automated processing operations and decision-making processes, including eligibility factors and marketing profiles.



Featurization:

Organizations should increase the ability of individuals to access to their data in intelligible, usable form in ways that allow them to analyze and utilize their own information. Featurization will allow individuals to declare their own policies, preferences and terms of engagement and “share the big data wealth” with organizations.



Privacy by Design:

Organizations should integrate privacy considerations early into lifecycle of new products and services. The assessment of privacy challenges at the design stage helps stem privacy risks at their outset. Moreover, privacy by design processes encourage organizations to revisit privacy issues throughout a project’s life.



De-Identification:

While there are many different understandings of what constitutes effective anonymization, organizations should implement practical de-identification processes that make use of legal and administrative safeguards in addition to reasonable technical measures. Both the sensitivity of the data and the utility of the data must be considered.

Risk mitigation strategies are essential for protecting privacy, yet at the same time they can constrain beneficial uses of data, thereby minimizing data utility. In addition, mitigation strategies alone do not help organizations decide **when is it worthwhile to proceed with a big data project despite residual privacy risks**. For example, if big data analysis can generate a health benefit that will improve the lives of millions of people, it may be ethical to allow a project to proceed even if privacy risks cannot be completely eliminated. Conversely, if the likelihood of accomplishing a benefit is extremely remote or if the contemplated benefit is minor, large privacy risks would not be justified.

Introducing Data Benefit Analysis (DBA)

By focusing exclusively on privacy risk, existing PIA practice does not account for the tremendous variance in anticipated big data benefits. This drives policymakers and corporate decision-makers into rote discussions of an almost ideological nature, with each side claiming the moral high ground and fully discounting arguments made by the other side. What is needed is a more thorough vocabulary of big data benefits. The following analysis proposes a methodology to better structure the discussion of big data benefits, assessing such variables as the *nature* of the benefit, the *identity* of the beneficiary and the *likelihood* of success. The results of this process, in turn, will feed into existing PIA practice to form a balanced, comprehensive view of big data risks and rewards.

Big data promises extraordinary benefits ranging from breakthroughs in medical research to enhancement of product offerings.

A Global Human Trafficking Hotline Network.¹⁵ Non-profit organizations collaborate to establish an international information-sharing database to collect sensitive information about human trafficking. Together with law enforcement authorities, these organizations can use this information to help combat organized crime.

Internet Searches Reveal Harmful Drug Interactions. Medical researchers use massive datasets of de-identified Internet search results to discover harmful drug interactions, by comparing individuals' search queries against "fingerprints" of adverse side effects.¹⁶

Gathering Voice Data to Improve Speech Recognition.¹⁷ Directory assistance services collect millions of voice samples in order to help create effective digital assistants that are embedded into mobile devices.

Ensuring High School Students Graduate.¹⁸ Using data collected across school districts, analytic tools predict which students are at risk of dropping out of school, providing schools and educators with a mechanism for early intervention.

Improve Newspapers' Ability to Compete Online.¹⁹ Traditional news publishers like the New York Times are turning to data in order to better serve subscribers with targeted reporting and interest-based content.

Location Data Create Intelligent Highways.²⁰ Geolocation information automatically generated by commuters' mobile devices is used to visualize traffic patterns in real time, helping urban planners manage traffic to deliver cost savings, environmental benefits, and a higher quality of life for commuters.

Using TV Viewing Habits to Create Better Entertainment.²¹ By analyzing the viewing habits of millions of users, entertainment streaming services can not only recommend better programming to viewers but also create new shows and programs that are tailored to their viewers' tastes.

Tracking lost baggage and improving customer service.²² Airlines are increasingly using data to offer travelers the ability to track their bags from curb to baggage claim, and big data analytics is directly improving customer experiences by allowing airlines to understand what travelers want.

Introducing Data Benefit Analysis (DBA)

So far, there has been no framework in place to assess big data benefits in a way commensurate with existing PIA risk frameworks. Yet accounting for costs is only part of a balanced and ethical value equation. In order to complete a cost-benefit analysis, organizations need tools to help them assess, prioritize, and to the extent possible, quantify a project's *rewards*. Not all benefits are or should be treated as equal: a potentially big benefit with a high likelihood of success must be treated differently than a smaller benefit with a similarly high likelihood of success – or a big benefit that is unlikely to ever be accomplished.

The scope and dimensions of big data benefits have not been accounted for under current PIA practice. Yet existing legal frameworks already recognize the need to balance privacy risks against data rewards. For example, Section 5 of the Federal Trade Commission Act defines as “unfair” a practice that “causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers them-

The Federal Trade Commission engages in a balancing test when determining when a practice is “unfair” by assessing whether any potential injury to consumers is **not outweighed by countervailing benefits** to consumers.

selves *and not outweighed by countervailing benefits to consumers or to competition.*”²³ Similarly, the European Data Protection Directive²⁴ and new draft Regulation²⁵ authorize the processing of personal data based upon a “legitimate interest” of an organization, requiring organizations to perform a balancing test between individual risks and organizational rewards.

The Article 29 Working Party recognizes that in determining the “legitimate interests” of organizations, those interests must be weighed against **the potential effect on individual rights**.

The European Article 29 Data Protection Working Party has recently presented a balancing test to help organizations determine whether their legitimate interest in processing data outweighs the rights or interests of individual data subjects. The test recognizes that benefits may range “from insignificant through somewhat important to compelling.”²⁶

In similar vein, institutional review boards, which evaluate the ethics of human subject research proposals, are instructed to evaluate risks in relation to anticipated benefits, taking into account both prevailing community standards and subjective risk and benefit determinations.²⁷

Introducing Data Benefit Analysis (DBA)

Maximizing the potential of big data requires a new framework for evaluating not only the risks but also the benefits created by novel data uses. To account for unique big data benefits and risks, organizations should engage in the following data benefit analysis.

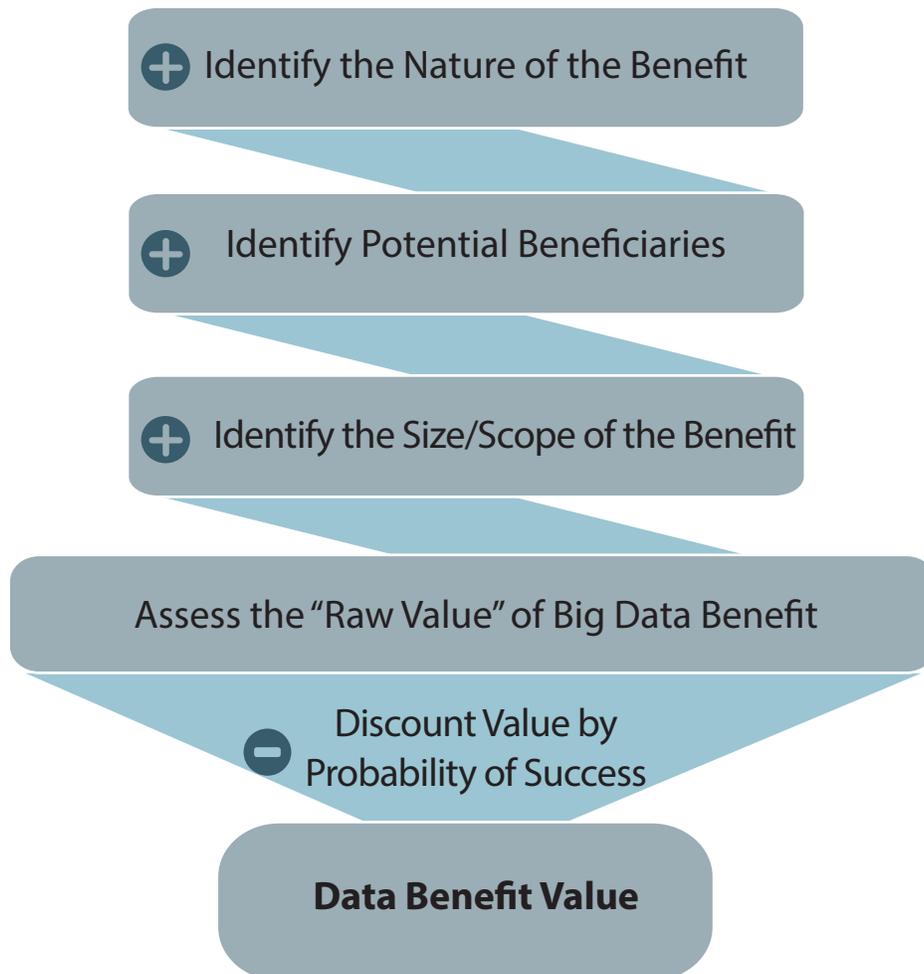
Data Benefit Analysis comprises two elements. **First**, organizations should assess the “raw value” of a benefit, which consists of (1) the nature of the benefit, (2) the potential beneficiaries, and (3) the degree (or size and scope) of the benefit. **Second**, organizations should discount the raw value score by the probability that the benefit can be achieved to obtain a discounted value score.²⁸

This process intertwines with and complements any risk assessments. The end goal is to achieve an optimal balance between organizational needs and individual privacy.

There are no definitive rules on what degree or probability of benefit is needed to overcome presumptions against creating privacy risk. It is clear that the mere assertion that a product or service can be improved is not enough; yet proof beyond any doubt is an unreasonable standard.

Any analysis must take into account culture-specific differences in evaluating the relative weight of each parameter. For example, the relative value of a health or national security benefit may differ from society to society. Some societies may place a high value on individual benefits, while others give greater weight to community values.

Depending upon how each of these factors compute, an organization will compile a raw value that reflects the potential benefit of a project – before taking into account uncertainty and weighing the benefits against privacy risks.



Hypothetical Case Study

Acme Corporation develops Road Runner, a fitness app that collects and analyzes information about users' diet, health, exercise and sleep. The app's data analysis provides users with helpful insights about their lifestyle, enabling them to optimize calorie consumption, reduce blood sugar and cholesterol levels, create a balanced exercise schedule, comply with doctors' prescriptions, and more.

A free app, Road Runner quickly gains traction, achieving a strong following with millions of users across the world. Acme collects and stores granular information about Road Runner users' habits, compiling statistics and creating graphs and indices that are accessible by users through an easy to use dashboard. In addition, Road Runner gives users real time notifications to inform them of any developing health conditions, such as lack of sleep, hypertension, dehydration, or failure to take medication.

Acme incentivizes employers to pre-package Road Runner into their mobile application management platforms by promising potential savings on their health care benefit plans. In turn, some employers are offering bonuses to employees who lose weight, optimize their body fat percentage, or exercise more. In addition, healthcare providers are urged to recommend to their patients usage of the app to enhance adherence to prescription medicine regimens.

Acme retains user data indefinitely, but keeps it in de-identified form by assigning random identifiers to individual users. Acme's CEO argues that in the future, the data retained by the company could be used to prevent epidemics, cure lethal disease, and increase life expectancy by up to 40 years.

To help the research community, Acme provides health researchers in accredited schools with access to its information. According to a recent article in the American Journal of Medication, researchers have been able to utilize Road Runner data to find a concealed harmful interaction between two best-selling drugs.

In the U.S., Acme provides periodic reports on longitudinal studies about users' health and behavior to the U.S. Department of Health and Human Services. The reports, which are in aggregated form, help the federal government make decisions concerning public health and research funding.

While Road Runner users' de-identified information could conceivably be linked to PII with varying degrees of certainty, this would have to be done through highly complex (and expensive) processes of data matching and analysis, which neither Acme nor its researchers and business partners have a clear interest to partake.

How to approach:

1 Conduct a full assessment of the **benefits** of the proposed data project (see pages 9-10)

2 Recognize and account for traditional **privacy risks** as well as risks that are unique to big data, and explore strategies for **risk mitigation** (see pages 2-4)

3 **Weigh** unmitigated privacy risks against big data rewards to determine how to proceed with a project (see page 11)

1) Identify the nature of the benefit:



Big data projects increasingly promise wide-ranging benefits to scientific research, public health, national security, law enforcement, energy conservation and economic efficiency. Organizations should recognize that the nature of the benefit must be accounted for by an analysis that measures social and cultural priorities and sensitivities.



The Road Runner app promises better information about personal health, cost savings to companies and communities, and additional knowledge to help inform public health policy and funding decisions.

2) Identify the potential beneficiaries:

Data projects can affect a wide-variety of stakeholders. These include not only the individual whose data is processed and the business that is processing the information, but increasingly also the government, a community, or society at large. As the OMB explains, "Analyses should include comprehensive estimates of the expected benefits and costs to society based on established definitions and practices for program and policy evaluation."²⁹

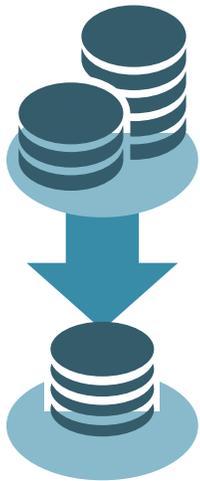


The Road Runner app primarily benefits the individual whose data it collects; but Acme also promises benefits to organizations through insurance cost savings, as well as to government, the research community, and potentially, society at large..

3) Assess the size or scope of the benefit and assign a Raw Value Score:

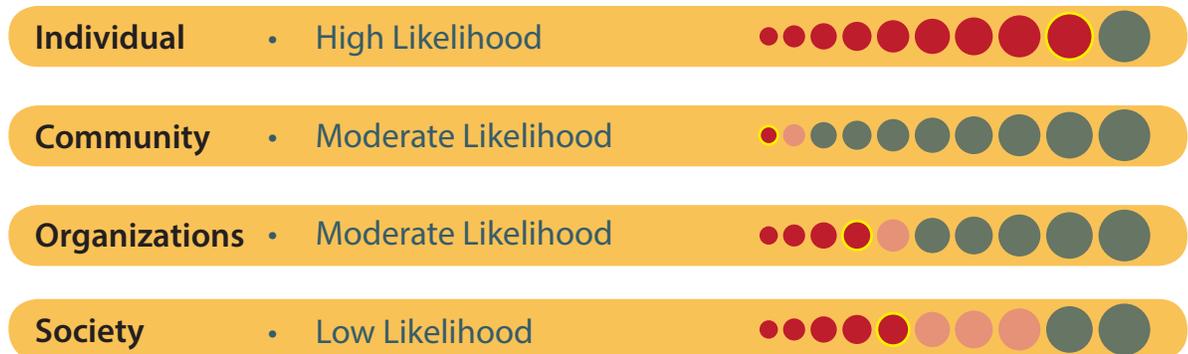


A raw value score combines the assessment of the beneficiary the nature, size and scope of the benefit. The raw value score represents the absolute value of a project prior to its discounting by probability and accounting for risk.



4) Discount by the probability of success:

After computing a raw value score, an organization must assess the likelihood that the benefits of a project will in fact come to pass. Uncertainty constitutes a **discount factor** that reduces the initial raw value score. **The certainty of obtaining the desired benefit is an essential element in determining the desirability of assuming related privacy risks.**



The Data Benefit Analysis should not be viewed as a static framing exercise.

The ultimate value of a data project's benefits as well as the magnitude of its privacy asks are linked to the risk mitigation strategies that have been implemented.

In many cases, mitigation techniques may impact data utility by reducing the potential benefit. This means that the **Data Benefit Analysis** is a dynamic process, through which mitigation techniques are carefully calibrated to optimize the risk-benefit equation in order to reach the apex point. The OMB calls this exercise "sensitivity analysis," noting that "[m]ajor assumptions should be varied and net present value and other outcomes recomputed to determine how sensitive outcomes are to changes in the assumptions."³⁰ Of course, in many cases, a baseline level of protection against risk will mandatory under regulation in order to support the legitimacy of the data processing.

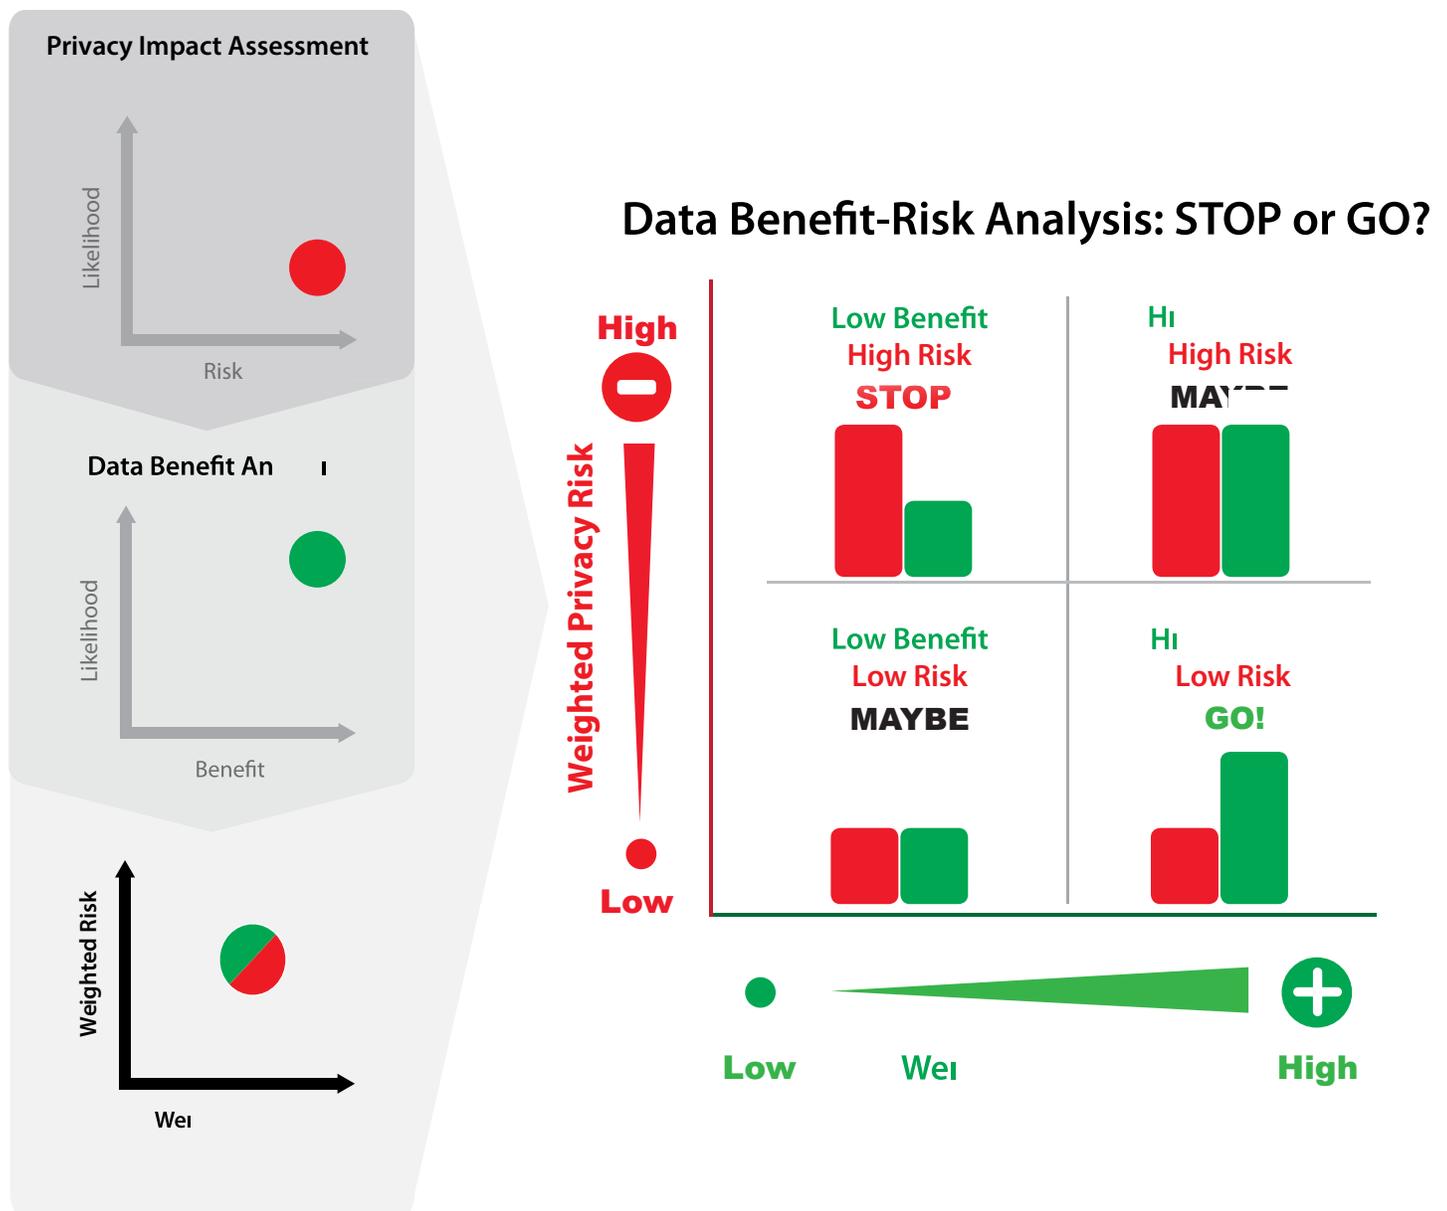
For example, as data are further de-identified, societal benefit is reduced but so is attendant privacy risk.

Mapping Benefits against Risks

Once an organization has a better understanding of a project's benefits, it can map the **discounted benefit value** against privacy risks identified through a PIA. By doing so, it can now visualize a complete picture to inform decision-making weighing both benefits and risks.

By mapping benefits against risks, an organization evaluates the merits of a big data project. To do so, an organization must elucidate where a project falls on the a risk-benefit continuum.

Mapped in this way, a contemplated project is placed on a continuum ranging from projects that the FTC and the Article 29 Working Party may view as unfair to project that the regulators regard as being within the legitimate interest of the organization or the public at large.



Who Decides?

While some of the assessments proposed in this framework can be standardized and quantified, others require value judgments and input from experts other than privacy professionals or data regulators. For example, assessing the scientific likelihood of capturing a benefit in a specialized area cannot be made solely based on privacy expertise.

Furthermore, this framework cannot achieve mathematical accuracy given the inherent degree of subjectivity in assessing the relative merits of various benefits. However, this has not stopped policymakers in other arenas from proposing structured processes to measure project benefits against risks. For example, the OMB states, “Although net present value is not always computable ... efforts to measure it can produce useful insights even when the monetary values of some benefits or costs cannot be determined.”³¹

This highlights the importance of determining *who* will be tasked with undertaking the Data Benefit Analysis. Moving forward, organizations will need to create or expand accountable data ethics review processes to engender trust and address privacy concerns. Many companies have already laid the groundwork to address these decision-making challenges by appointing Chief Privacy Officers or building internal ethical review programs. Further efforts are needed to understand the most effective structures for different organizations and different types of data. Models may range from a formal Institutional Review Board-type process to empowering Chief Privacy Officers through cross-functioning privacy committees, or involve building structures such as external advisory boards or opportunities for policy maker or regulator input.

What is an Institutional Review Board?

Institutional Review Boards (IRB)³² emerged as the chief regulatory response to concerns about ethical abuse in the use of human subjects for research. IRBs are therefore charged with balancing the potential risks and benefits arising from any project involving human subject research. Policy guidance on IRBs recognizes that research benefits fall into different categories, including acquiring new knowledge, improving drug safety, promoting technological advances, or providing better healthcare.

IRBs must have at least five members, encompassing a wide variety of backgrounds and professional expertise. Boards that review research involving specific categories of human subjects, such as children, pregnant women, or the mentally disabled, must include members who have special experience with those groups.

An IRB’s final assessment of a project depends on prevailing community standards and subjective determinations of risks and benefits. While there are limits on the risks that individuals should ethically be asked to accept for the potential benefit of others, IRBs are generally directed to not be overprotective.

1. TRILATERAL RESEARCH & CONSULTING, *PRIVACY IMPACT ASSESSMENT AND RISK MANAGEMENT, A REPORT FOR THE INFORMATION COMMISSIONER'S OFFICE* (May 4, 2013), http://ico.org.uk/about_us/consultations/~media/documents/library/Corporate/Research_and_reports/pia-and-risk-management-full-report-for-the-ico.pdf.
2. Omer Tene & Jules Polonetsky, *Privacy in the Age of Big Data: A Time for Big Decisions*, 64 STAN. L. REV. ONLINE 63 (2012).
3. Civil Rights Principles for the Era of Big Data, <http://www.civilrights.org/press/2014/civil-rights-principles-big-data.html> (last visited March 15, 2014).
4. Article 29 Data Protection Working Party, Opinion 06/2014 on the notion of legitimate interests of the data controller under Article 7 of Directive 95/46/EC (Apr. 9, 2014), https://www.huntonprivacyblog.com/files/2014/04/wp217_en.pdf.
5. WHITE HOUSE, EXECUTIVE OFFICE OF THE PRESIDENT, *BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES* (May 2014), http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.
6. Office of Management & Budget, Circular No. A-94 Revised, Memorandum for Heads of Executive Departments and Establishments, Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs (Oct. 29, 1992), http://www.whitehouse.gov/omb/circulars_a094.
7. INFORMATION COMMISSIONER'S OFFICE, *CONDUCTING PRIVACY IMPACT ASSESSMENTS CODE OF PRACTICE* (February 2014), http://ico.org.uk/for_organisations/data_protection/topic_guides/~media/documents/library/Data_Protection/Practical_application/pia-code-of-practice-final-draft.pdf.
8. DEPARTMENT OF HOMELAND SECURITY, *PRIVACY IMPACT ASSESSMENTS, THE PRIVACY OFFICE OFFICIAL GUIDANCE* (June 2010), https://www.dhs.gov/xlibrary/assets/privacy/privacy_pia_guidance_june2010.pdf.
9. ICO *PRIVACY IMPACT ASSESSMENT HANDBOOK* (2010) (describing a five-part PIA process including (1) a preliminary phase, (2) preparation, (3) consultation and analysis, (4) documentation, and (5) review and audit. See also <http://www.piafproject.eu> (last visited March 15, 2014).
10. Ryan Calo, *The Boundaries of Privacy Harm*, 86 IND. L.J. 1131 (2011); SILICON FLATIRONS CENTER, *THE NEW FRONTIERS OF PRIVACY HARM*, January 2014, <http://www.siliconflatirons.org/events.php?id=1381>.
11. Daniel J. Solove, *A Taxonomy of Privacy*, 154 U. PA. L. REV. 477 (2006). *Information collection* includes surveillance and interrogation; *information processing* includes aggregation, identification, insecurity, secondary use and exclusion; *information dissemination* includes breach of confidentiality, disclosure, exposure, increased accessibility, blackmail, appropriation and distortion; and *invasion* includes intrusion and decisional interference.
12. Julie Brill, Commissioner, Fed. Trade Comm'n, Remarks at Fordham University School of Law: Big Data, Big Issues (Mar. 2, 2012) (transcript available at <http://ftc.gov/speeches/brill/120228fordhamlawschool.pdf>). Commissioner Brill said: "Big Data's impact on privacy is requiring some new and hard thinking by all of us."
13. Omer Tene & Jules Polonetsky, *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 NW. J. TECH. & INTELL. PROP. 239 (2013).
14. *Id.* at 271.
15. Press Release, Polaris Project Launches Global Human Trafficking Hotline Network, Polaris Project (Apr. 9, 2013), <http://www.polarisproject.org/media-center/news-and-press/press-releases/767-polaris-project-launches-global-human-trafficking-hotline-network>.
16. Nicholas Tatonetti et al., *Detecting Drug Interactions From Adverse-Event Reports: Interaction Between Paroxetine and Pravastatin Increases Blood Glucose Levels*, 90 CLINICAL PHARMACOLOGY & THERAPEUTICS 133 (2011); Nicholas Tatonetti et al., *A Novel Signal Detection Algorithm for Identifying Hidden Drug-Drug Interactions in Adverse Event Reports*, 12 J. AM. MED. INFORMATICS ASS'N 79 (2011).
17. Derrick Harris, *Google Explains How More Data Means Better Speech Recognition*, GIGAOM (Oct. 31, 2012), <http://gigaom.com/2012/10/31/google-explains-how-more-data-means-better-speech-recognition/>.
18. Press Release, Alabama's Largest School District Turns to IBM to Build a Smarter Education System, IBM (July 13, 2009), <http://www-03.ibm.com/press/us/en/pressrelease/27984.wss>.
19. Antonio Regalado, *Unsubscribing? The New York Times Wants to Predict That*, MIT TECH. REV. (Feb. 12, 2014), <http://www.technologyreview.com/news/524716/unsubscribing-the-new-york-times-wants-to-predict-that/>.
20. Haomiao Huang, *Calling All Cars: Cell Phone Networks and the Future of Traffic*, ARSTECHNICA (Feb. 24, 2011), <http://arstechnica.com/gadgets/2011/02/calling-all-cars-measuring-traffic-using-cell-phone-data/>.
21. David Carr, *Giving Viewers What They Want*, N.Y. TIMES (Feb. 24, 2013), http://www.nytimes.com/2013/02/25/business/media/for-house-of-cards-using-big-data-to-guarantee-its-popularity.html?pagewanted=all&_r=1&.
22. Katherine Noyes, *For the Airline Industry, Big Data Is Cleared for Take-Off*, FORTUNE (June 19, 2014), <http://fortune.com/2014/06/19/big-data-airline-industry/>.
23. 15 U.S.C. Sec. 45(n) (emphasis added).
24. Article 7(f) of Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data, 1995 O.J. (L 281) 31 (Nov. 23, 1995), available at <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:1995:281:0031:0050:EN:PDF>.
25. Article 6(1)(f) of Proposal for a Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), COM(2012) 11 final (Jan. 25, 2012), available at http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf.
26. Opinion 06/2014 on the Notion of Legitimate Interests of the Data Controller Under Article 7 of Directive 95/46/EC, 844/14/EN WP 217, at 30 (Apr. 2014).
27. In its guidance, the OMB calls for discounting both costs and benefits, stating that "Estimates of benefits and costs are typically uncertain because of imprecision in both underlying data and modeling assumptions. Because such uncertainty is basic to many analyses, its effects should be analyzed and reported."
28. Office of Management & Budget, Circular No. A-94 Revised, *supra* note 6.
29. *Id.* at 6, *supra* note 6.
30. *Id.* at 9(c), *supra* note 6.
31. *Id.* at 5(a), *supra* note 6.
32. 45 CFR § 46. See also U.S. DEP'T OF HEALTH & HUMAN RESOURCES, *INSTITUTIONAL REVIEW BOARD GUIDEBOOK* (1993), available at http://www.hhs.gov/ohrp/archive/irb/irb_chapter3.htm#e1.



Appendix B

BIG DATA:
**A Tool for Fighting Discrimination and
Empowering Groups**

Discussion Draft 7-31-14



PREFACE

In May 2014, the Executive Office of the President concluded its 90-day study of Big Data and privacy and released a report entitled *Big Data: Seizing Opportunities, Preserving Values*. The report highlighted certain positive uses of Big Data, such as identifying health risks at an early stage, creating efficiencies in energy distribution, and uncovering fraud through predictive analysis. However, it also concluded that Big Data analytics' could facilitate discrimination in housing, credit, employment, health, education, and a range of other markets. These potential benefits and drawbacks underscore the need to better understand how Big Data will shape our lives in years to come.

Recognizing the 50th anniversary of the Civil Rights Act and the challenges to fighting discrimination in the 21st century, the case studies included in this report show how businesses, governments, and civil society organizations are leveraging Big Data (and other data sets) to protect and empower vulnerable groups, including by providing access to job markets, uncovering discriminatory practices, and creating new tools to improve education and assist those in need. While by no means an exhaustive list of Big Data's potential to uncover and fight discrimination, we offer these examples to show how Big Data already is redefining efforts to ensure equal opportunity for all.

We would like to thank the Anti-Defamation League for its partnership in preparing this report, Jared Bomberg and Julian Flamant at Hogan Lovells US LLP for providing essential research and drafting support, and members of the FPF Advisory Board for reviewing drafts and providing guidance.

We hope these examples will contribute to discussions about Big Data's impact on discrimination.

Jules Polonetsky
Executive Director and Co-Chair
Future of Privacy Forum

Christopher Wolf
Founder and Co-Chair
Future of Privacy Forum

TABLE OF CONTENTS

I. Seeing Beyond Bias to Provide New Opportunities	1
Case Study 1: Workplace Diversity (Entelo)	1
Case Study 2: Smarter Hiring (Google)	2
Case Study 3: Allocation of Public Works (Cedar Grove Institute).....	3
Case Study 4: Demographics of Health (State of New York)	4
II. Transparency is a Necessary Disinfectant	5
Case Study 5: Auditing Anti-Semitism (Anti-Defamation League).....	5
Case Study 6: <i>McClesky v. Kemp</i>	6
Case Study 7: Discrimination Complaint Data (EEOC)	7
Case Study 8: <i>United States v. Sterling</i>	8
III. Long-Term Problems Require New Solutions	9
Case Study 9: Education for All (NSBA)	9
Case Study 10: Tracking Migratory Patterns (United Nations)	10
Case Study 11: Economic Development and Equality (OECD).....	11
Case Study 12: Finding Missing and Exploited Children (Palantir/NCMEC)	12
Case Study 13: Human Trafficking (Palantir/Polaris Project)	13

I. SEEING BEYOND BIAS TO PROVIDE NEW OPPORTUNITIES

Case Study 1: Workplace Diversity (Entelo)

Entelo Diversity, a candidate recruiting platform launched in April 2014, is improving workplace diversity by empowering recruiters to search for job candidates from within underrepresented segments of the population. Using a proprietary algorithm, this workplace diversity tool sifts through publicly available data—pulled from social media platforms—to match recruiters with candidates who hold necessary qualifications, but also meet particular diversity requirements. The tool can filter candidates based on gender, race, and military history in five categories: Female, African American, Asian, Hispanic, and Veteran.

Example of an Entelo Search

The screenshot displays the Entelo search interface. On the left, a search sidebar includes a search bar with the keyword 'PHP', location filters for 'San Francisco, CA 50mi', and various filters such as 'Email available', 'Years of experience', 'Recently updated', and 'Diversity'. The 'Diversity' filter is expanded, showing 'Female' selected. The main area shows '3,807 candidates found' and a list of three candidates:

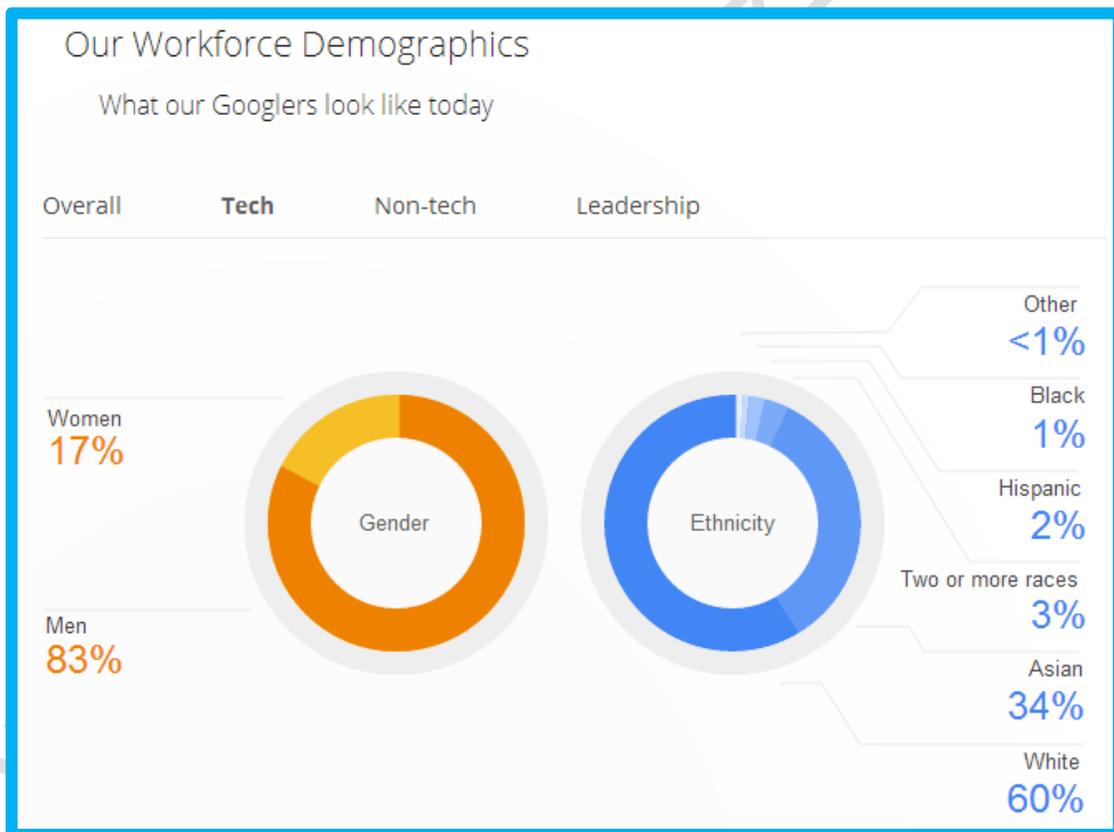
- Eve Killaby**: SAN FRANCISCO, CALIFORNIA. Back-end Engineer at Bloodhound - about 1 year. Senior Web Engineer at DocuSign, Inc. - 9 months. State University of New York at Buffalo, Bachelors of Science, Computer Science. Skills: mysql, node.js, software engineering, web development, PHP, web design, entrepreneurship, 25 more.
- Lia Napolitano**: SAN FRANCISCO BAY AREA. Siri Interaction Designer at Apple Inc. - over 1 year. User Experience Designer at Apple Inc. - almost 2 years. Wellesley College, BA, Media Arts and Sciences. Skills: user interface design, user interface, mac, user-centered design, PHP, flash, iOS, 31 more.
- Julia Kryzstofiak-Szopa**: SAN FRANCISCO BAY AREA, US. Founder at Wellfitting.com - about 1 year. Captain Ivanova at Blackbox Accelerator, LLC - about 1 year. Katholieke Universiteit Leuven, Philosophy, AI. Skills: git, html 5, web development, social media, PHP, user experience, Django, 51 more.

Source: [Entelo](#)

Case Study 2: Smarter Hiring (Google)

A challenge for the technology industry is ensuring diversity in the workplace. Twitter has recently reported that 90% of its global “tech” employees are male and Google admits “[it’s] not where [it] wants to be when it comes to diversity,” with only 1% black and 17% women among its tech workforce.

Recognizing the value of a diverse workforce, [Google](#) is leveraging its data analytics capabilities to help change those numbers. Through analytics, the company identified issues with its traditional hiring process, which favors certain groups and may provide little insight into a candidate’s “Googliness.” Based on this analysis, the company recently implemented a new interview process focused on structured behavioral questions, rather than “brainteasers” and academic G.P.A.’s. Google also now uses employee panels to evaluate candidates through several rounds of review. The result of Google’s intensive analytical study and the adoption of new hiring methods is that the company is better able to evaluate candidates on an individualized basis, where questions like race and gender do not hold sway.

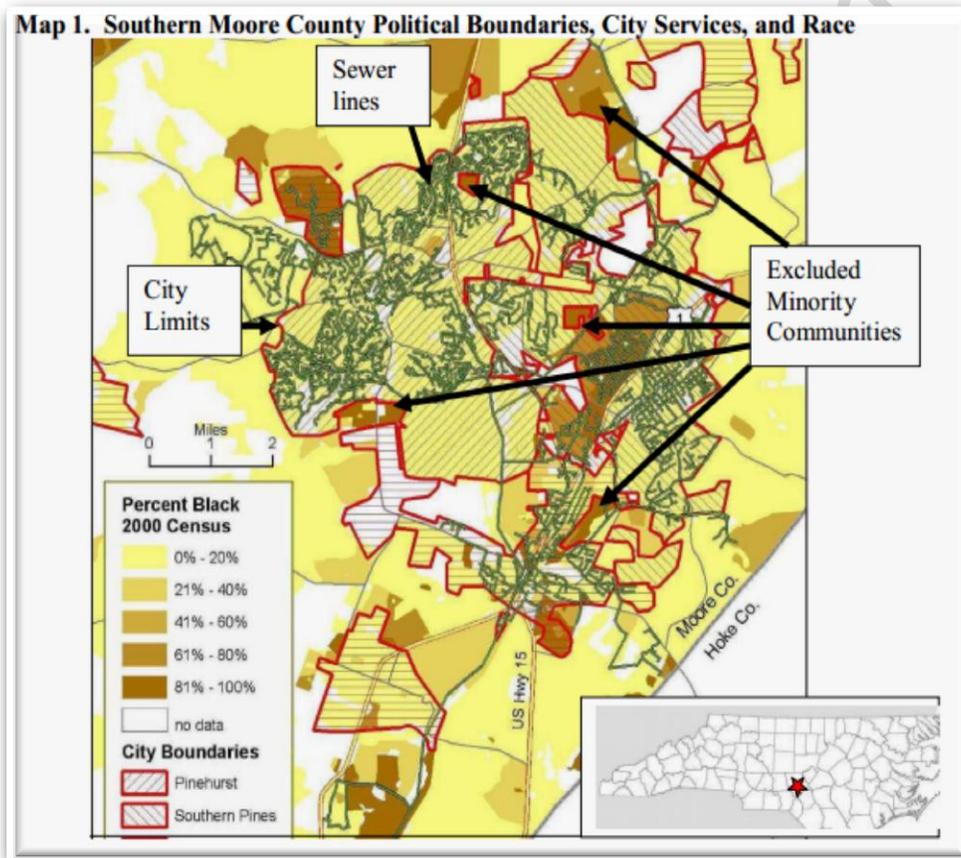


Source: [Google Diversity](#)

Case Study 3: Allocation of Public Works (Cedar Grove Institute)

The Cedar Grove Institute for Sustainable Communities is a non-profit organization that leverages open data to explore disparities in allocation of geographic boundaries and public works of various communities. Cedar Grove uses a combination of demographic analysis, contextual investigation, housing and economic analysis, and geographic information systems to explore potentially discriminatory implications of public policy decisions. In a 2005 study, [Segregation in the Modern South: A Case Study of Southern Moore County](#), Cedar Grove Institute combined census data and other publicly available surveys and demographic information to explore the impact on community development of land annexation policies of Moore County North Carolina.

Map Showing Excluded Minority Communities



Source: [Cedar Grove Institute for Sustainable Communities](#)

Case Study 4: Demographics of Health (State of New York)

In 2011, the Institute of Medicine, the health arm of the National Academy of Sciences, [reported](#) that lesbian, gay, bisexual, and transgender (LGBT) individuals have unique health experiences and needs, but as a nation, we do not know exactly what these needs are. The IOM also reported that clinicians and researchers are faced with incomplete information regarding the health status of LGBT individuals and that current research has not adequately examined subpopulations, particularly racial and ethnic groups and peoples' health needs based on age.

In response, the State of New York has [launched](#) a coordinated, multi-agency effort to strengthen data collection regarding LGBT individuals in New York. The campaign will rely on data collected on a self-reporting basis by the New York's Department of Health, Department of Corrections and Community Supervision, Office for the Aging, Office of Mental Health, Office of Alcohol and Substance Abuse Services, Office of Temporary and Disability Assistance, Office of Children and Family Services, and Office for People with Developmental Disabilities. The data collected will be shared among the eight agencies to create a comprehensive method for identifying the needs of the lesbian, gay, bisexual, and transgender community. It is hoped that stronger data sets will empower the State and others to create more tailored approaches to reduce health disparities impacting the LGBT community.

News Article

Wednesday, July 23, 2014

GOVERNOR CUOMO ANNOUNCES MULTI-AGENCY STATE EFFORT TO ADDRESS LGBT DISPARITIES

New York becomes first state in the nation with coordinated statewide strategy to improve LGBT data collection

Governor Andrew M. Cuomo announced that New York State is undertaking a coordinated, multi-agency effort to strengthen data collection for lesbian, gay, bi-sexual and transgender (LGBT) New Yorkers. Outlined in the first report by the State's Interagency LGBT Task Force, this statewide effort to include sexual orientation and gender identity information in data collections will allow the state to better tailor services to meet LGBT needs, ultimately improving the health and lives of thousands of New Yorkers. This effort makes New York the first state in the nation to employ a coordinated strategy to develop its data collection procedures for the LGBT community.

"New York State has a long history of advancing progressive ideals, and today we are continuing to lead the nation by identifying new ways to improve services and better meet the needs of the LGBT community," Governor Cuomo said. "By being more inclusive with how state agencies monitor the demographics of those they serve, we can address health and financial disparities, safety concerns, and a myriad of other issues that impact LGBT New Yorkers. This is another step forward for an important community in New York, and our administration will continue standing up for all New Yorkers, regardless of their sexual orientation or gender identity."

The Institute of Medicine in its March 2011 report, *The Health of Lesbian, Gay, Bisexual, and Transgender People: Building a Foundation for Better Understanding*, emphasized the need for collection of gender identity

Source: [State of New York](#)

II. TRANSPARENCY IS A NECESSARY DISINFECTANT

Case Study 5: Auditing Anti-Semitism (Anti-Defamation League)

Anti-Semitism is a persisting global problem. Each year the Anti-Defamation League publishes an *Audit of Anti-Semitic Incidents*, which reports on both criminal and non-criminal acts of harassment and intimidation, including distribution of hate propaganda, threats and slurs. In 2013 alone, there were 751 anti-Semitic incidents across the United States. The annual report combines data on anti-Semitic assaults, vandalism, and harassment events as collected by 27 regional offices across the U.S. While the data is broken down to detail each event and show granular state-wide trends, it is also aggregated to provide a comprehensive analysis of anti-Semitism on a nationwide basis. In 2013, total anti-Semitic incidents declined by 19 percent across the U.S., while “violent” incidents increased by 17 percent.

In recent years, the ADL Audit has placed anti-Jewish incidents into one of three categories: assault, vandalism or harassment. In 2013 the numbers broke down as follows:

- *Assaults*: 31 incidents reported in 2013, compared with 17 in 2012;
- *Vandalism*: 315 incidents in 2013, compared with 440 in 2012;
- *Harassment, threats and events*: 405 incidents in 2013, compared with 470 in 2012.

State-By-State Totals

While the number of anti-Semitic incidents in the U.S. decreased in 2013, ADL recorded increases in some states when compared with 2012 figures, such as in Connecticut, Michigan, Ohio and Texas. Additionally, continuing a consistent trend for many years, the states with the highest totals were those with large Jewish populations. Once again, New York and California, while each experiencing declines compared to the previous year, topped the list:

- *New York State*, with 203 incidents in 2013, down from 245 in 2012;
- *California*, with 143 incidents, down from 185;
- *New Jersey*, with 78 incidents, down from 173;
- *Florida*, with 68 incidents, down from 88;
- *Massachusetts*, with 46 incidents, up from 38;
- *Pennsylvania*, with 43 incidents, up from 37.

Source: [Anti-Defamation League](#)

Case Study 6: *McClesky v. Kemp*

McCleskey, an African American man, was sentenced to death after being convicted of armed robbery and the murder of a white police officer. In a writ of habeas corpus, McCleskey argued that the Georgia capital sentencing process was administered in a racially discriminatory manner in violation of the Eighth and Fourteenth Amendments. In support of the claim, McCleskey offered a statistical study (the Baldus study) to show disparities in the imposition of the death sentence in Georgia based on the murder victim's race and the defendant's race. The study was based on over 2,000 murder cases that occurred in Georgia during the 1970's, and involved data relating to the victim's race, the defendant's race, and the various combinations of such persons' races. The study found a consistent pattern of discrimination in the use of the death penalty against defendants who were charged with killing white victims compared to defendants who were charged with killing African American victims.

While the court ultimately found against McCleskey, the case has been described as a turning point in the debate over the death penalty in the United States. The Baldus study has been replicated in numerous jurisdictions with similar findings. Race is now a powerful issue in debates over the death penalty because of studies like the Baldus study, which show that race can affect death penalty decisions.

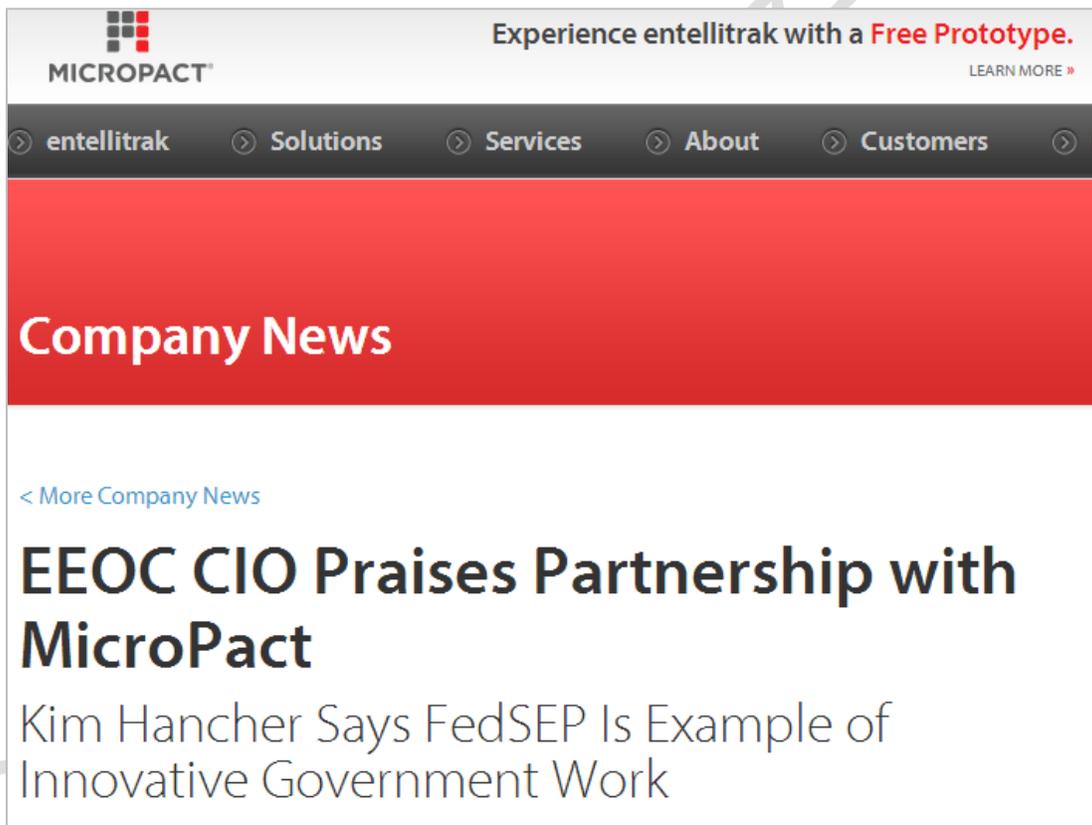


Source: *McClesky v. Kemp*, 481 U.S. 279 (1987).

Case Study 7: Discrimination Complaint Data (EEOC)

The Equal Employment Opportunity Commission (EEOC) is responsible for enforcing federal laws that make it illegal to discriminate against a job applicant or an employee because of the person's race, color, religion, sex (including pregnancy), national origin, age (40 or older), disability, or genetic information.

In March 2013 the EEOC unveiled FedSEP, an electronic portal through which more than 325 federal agencies interact with the EEOC relative to their workforce and complaint data. This gateway provides each agency's staff with a single point of access to EEOC data collection systems and provides a new tool to collect and analyze government-agency data on workplace discrimination charges. By streamlining the submission process for so many documents and aggregating data from many sources, FedSEP allows EEO professionals greater ability to spot trends and uncover discrimination across the federal government.



Source: [Entellitrak](#)

Case Study 8: *United States v. Sterling*

In 2006, Donald Sterling, his wife, and their family-trust real-estate company, *Beverly Hills Properties*, were accused of engaging in discriminatory practices in violation of the Fair Housing Act and Title VIII of the Civil Rights Act of 1968. The United States alleged, *inter alia*, that the defendants refused to rent portions of their 27-building development in the “Koreatown” neighborhood of Los Angeles to non-Koreans (*e.g.*, *Hispanics and blacks*). A study by Dr. Shelley Lapkoff used a database of tenant information released by the defendants to show that the number of Korean tenants across the 27 buildings had increased “significantly,” from 64 percent to 83 percent, within the year following acquisition of those buildings by the defendants. Dr. Lapkoff’s report also included an analysis of census data to determine whether the defendant’s claim that the changing demographic distribution of Koreatown could explain the decreasing diversity of its tenants.

Dr. Lapkoff’s analysis found that the overall demographics of Koreatown remained relatively stable during the period in question, and that Hispanics remained the dominant race in the area. The report concluded that, absent external changes, the increase in Korean tenants was consistent with the United States’ allegation of housing discrimination. A later study by Dr. Lapkoff also used census data to show that there were no major shifts in household income or Korean households in the area that could explain the increase of Korean renters.

The studies helped lead to a settlement agreement that included a number of measures aimed at ending discriminatory renting practices in the 27 buildings owned by Beverly Hills Properties and required the defendants to pay \$2,625,000 to be disbursed among aggrieved persons and a \$100,000 civil penalty.

Table 3
Percentage of New Tenants Who Were Korean, Before and After Acquisition

Building #	Year Before Acquisition			Year After Acquisition		
	Total	Korean	% Korean	Total	Korean	% Korean
Buildings Where Less than 80 percent of New Tenants are Korean at Time of Acquisition						
101	12	0	0%	15	4	25%
112	4	0	0%	6	6	100%
93	14	1	7%	18	13	72%
96	9	1	11%	11	9	82%
102	25	3	12%	22	8	36%
105	8	2	25%	7	5	71%
108	20	5	25%	30	28	93%
85	12	4	33%	14	13	93%
106	20	7	35%	18	11	61%
98	7	4	57%	11	11	100%
88	43	26	60%	41	34	83%
97	8	5	63%	10	8	80%
82	29	20	69%	27	25	93%
95	33	24	73%	44	38	86%
104	19	15	79%	28	25	89%
Subtotal	263	117	44%	303	238	79%
Buildings Where More than 90 percent of New Tenants are Korean at Time of Acquisition						
87	40	36	90%	45	43	96%
89	12	11	92%	13	13	100%
83	11	11	100%	9	9	100%
84	18	18	100%	17	14	82%
90	8	8	100%	19	17	89%
91	9	9	100%	4	4	100%
94	5	5	100%	10	10	100%
103	14	14	100%	18	16	89%
107	14	14	100%	10	9	90%
109	14	14	100%	20	18	90%
110	12	12	100%	7	6	86%
Total	157	152	97%	172	159	92%

Building # 86 Had No Rentals the Year Before Acquisition for Comparison

Source: *United States v. Sterling*, No. 2:06CV04885, (C.D. Cal. Nov. 12, 2009).

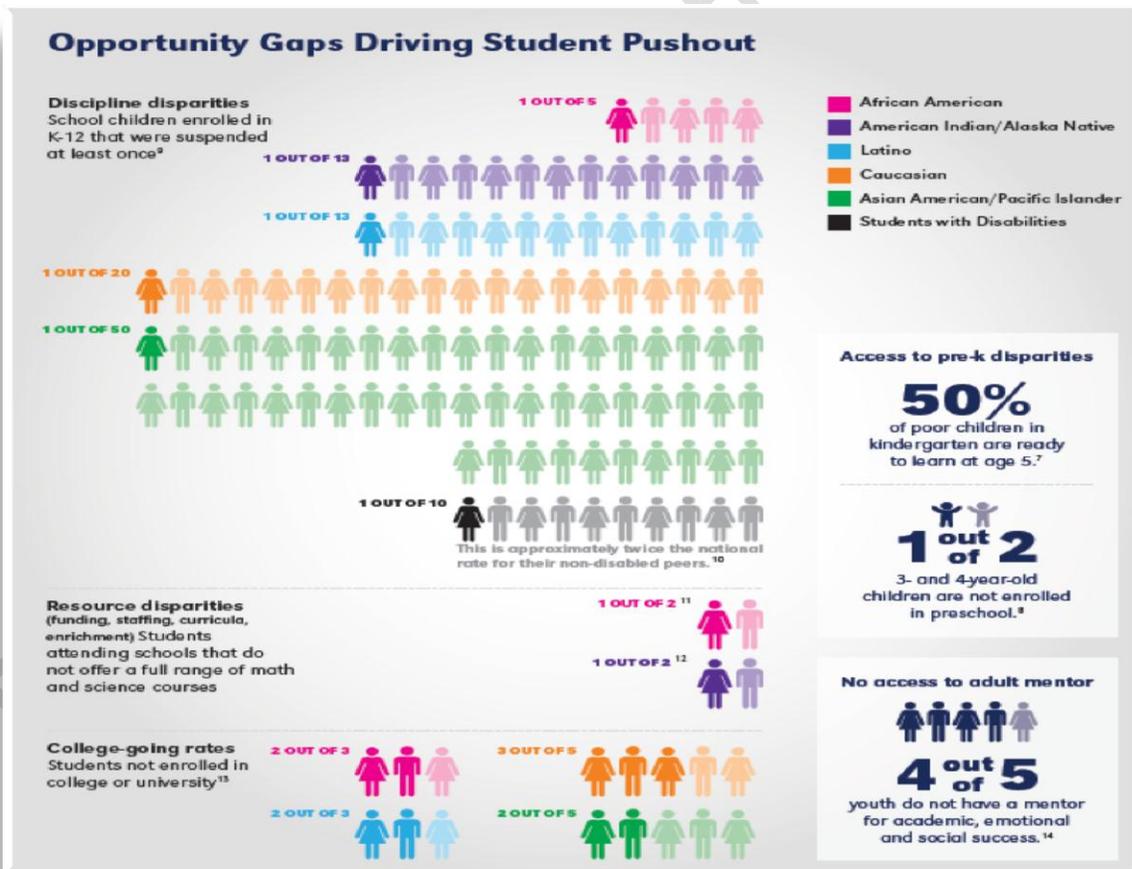
III. LONG-TERM PROBLEMS REQUIRE NEW SOLUTIONS

Case Study 9: Education for All (NSBA)

A recent report by National School Boards Association (NSBA) offers novel policy solutions for increasing education rates in America. The report, *Partnerships, not Pushouts*, combines census data with data collected by various organizations to identify factors—known as “pushouts”—that may be responsible for driving young people away from education. Pushout factors can be more common among different segments of the population. For example, school suspensions—considered a major “pushout” factor, affect one out of five African American students and only one out twenty Caucasian students, which may partly explain the large discrepancy between graduation rates of those two groups.

To increase education levels among American youth, the NSBA proposes a variety student-centered “Personal Opportunity Plans” (POPs). To be effective, POPs are tailored to meet the needs of students on an individualized basis, addressing the pushout factor(s) most threatening to a particular student’s academic success.

Discipline Disparities



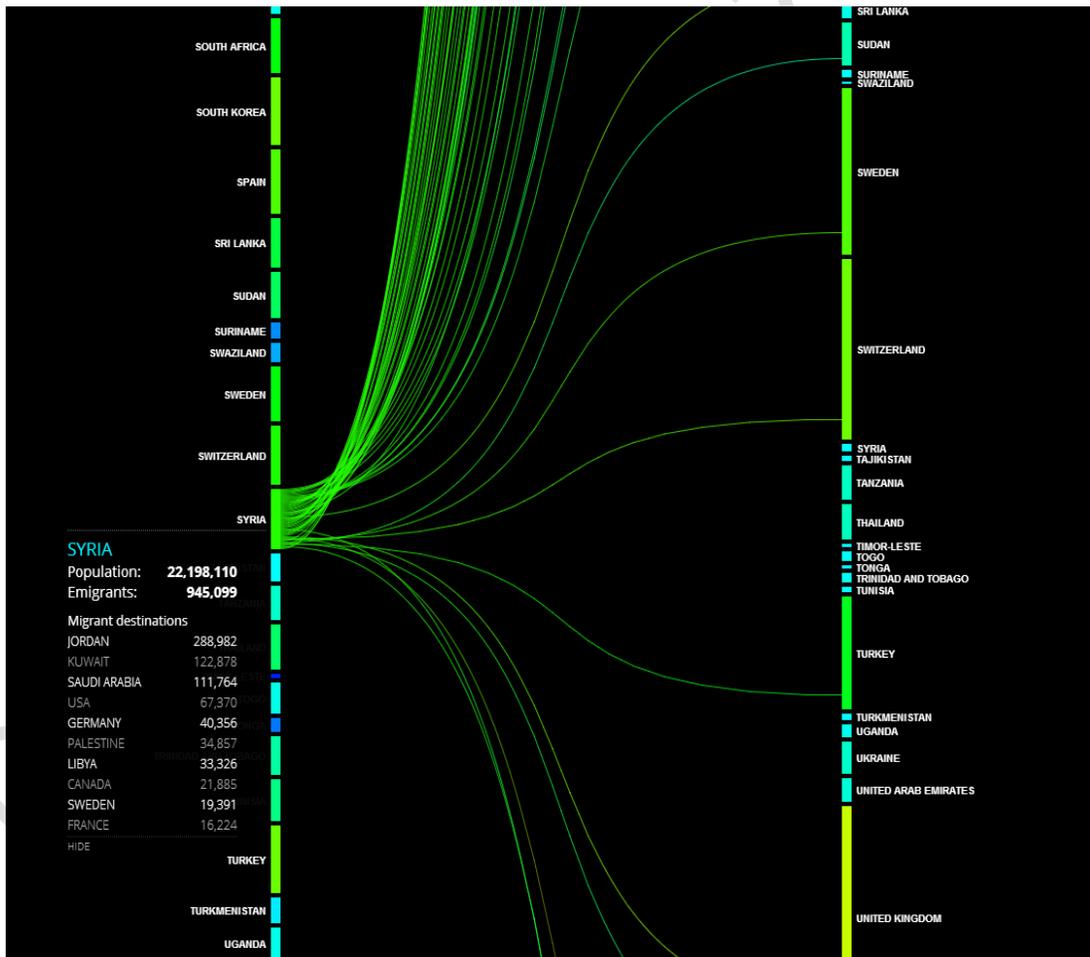
Source: [National School Board Association](#)

Case Study 10: Tracking Migratory Patterns (United Nations)

The [United Nations has highlighted the social benefits of tracking migratory patterns](#) of diverse peoples. For example, tracking the movement of displaced populations can empower humanitarian groups to provide better aid to those populations. As a new project under the *UN Global Pulse* banner, the organization is exploring new ways to track displaced populations using Big Data. The organization cites “significant shortcomings” with traditional methods of migratory benchmarking such as censuses, demographic and thematic surveys and administrative registers, which quickly become outdated.

In its review, Global Pulse highlights a number of studies that rely on Big Data collected from social media sites or open data initiatives to draw important conclusions about population movements. In the example below, PeopleMovin, repurposes “open” migration, refugee and asylum, and world population data to create an interactive tool allowing users to quickly identify international movement patterns and identify where relief efforts are most valuable.

Migration Patterns from Syria:



Source: http://peplemov.in/#f_SY

Case Study 11: Economic Development and Equality (OECD)

Since 2009, the Organization for Economic Co-Operation and Development has offered the publicly available *Gender, Institutions and Development Database*. The database compiles gender-discrimination data from 160 countries to provide researchers and policymakers with an analysis of 60 detailed variables, ranging from factors like “Discriminatory Family Code” to “Restricted Civil Liberties,” that are likely to impact women’s engagement in society and the economy.

A defining feature of the GID-DB is that, in addition to traditional quantitative analyses, the database uses an innovative scoring system to evaluate discriminatory institutional features. For example, while traditional studies of “early marriage” analyze rates of marriage among various age groups, the GID-DB has created a scaled system that combines rates of early marriage with an analysis of legal, traditional and religious customs to provide a much deeper look at gender discrimination. An example of the scaled system is provided below:

- 0: The law on the minimum age of marriage does not discriminate against women.
- 0.5: The law on the minimum age of marriage discriminates against some women, for example through customary, traditional and religious law.
- 1: The law on the minimum age of marriage discriminates against all women or there is no law on the minimum age of marriage.

Gender, Institutions and Development Database (GID-DB)

The screenshot displays the OECD StatExtracts interface for the Gender, Institutions and Development Database (GID-DB). The interface includes a search bar, navigation tabs, and a data table for the year 2012. The table lists variables such as 'Legal Age of Marriage', 'Early Marriage', 'Parental Authority', 'Inheritance', 'Violence Against Women (Laws)', 'Female Genital Mutilation', 'Domestic Violence', and 'Reproductive Integritty' across various countries including Afghanistan, Albania, Algeria, Angola, and Argentina.

Variable	Discriminatory Family Code				Restricted Physical Integrity				Social Institutions				
	Legal Age of Marriage	Early Marriage	Parental Authority	Inheritance	Violence Against Women (Laws)	Female Genital Mutilation	Domestic Violence	Reproductive Integritty	Attitudes	Lifetime Prevalence			
Afghanistan	1	0.17	1	1	0.5	0.5	0.75	0.5	0.25	0	0.57	..	0.23
Albania	0	0.1	0	0	0.5	0	0.25	0.5	0.25	0	0.3	0.08	0.13
Algeria	0	0.02	1	0	0.5	0.5	0.5	0.75	0.25	0	0.68	..	0.25
Angola	0.5	0.36	0	0	0.5	0.5	0.25	1	1	0.78	0.27
Argentina	0	0.11	0	0	0	0	0.25	0.25	..	0	0.06

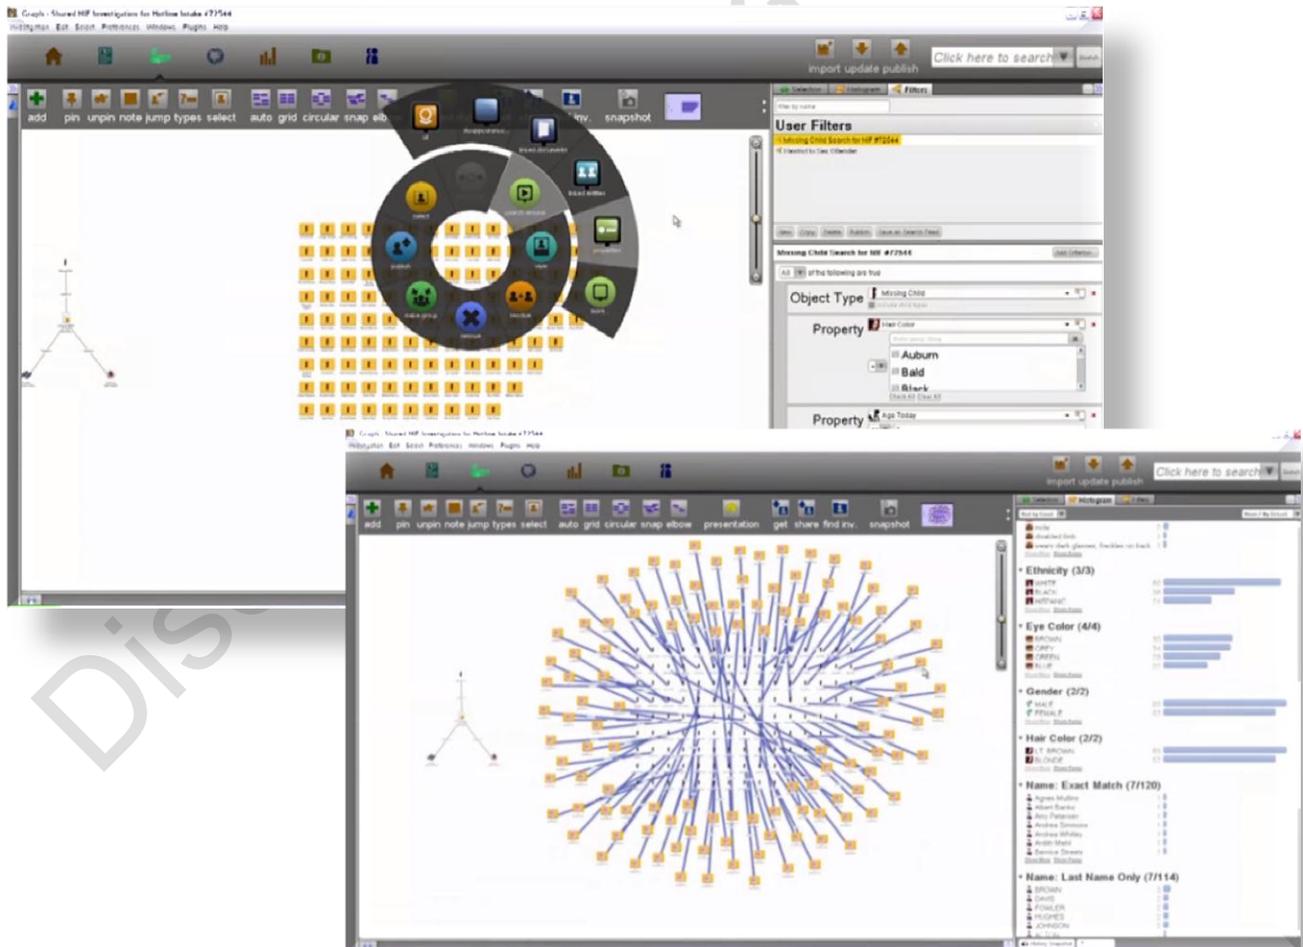
Source: [OECD, Gender, Institutions and Development Database \(GID-DB\)](#)

Case Study 12: Finding Missing and Exploited Children (Palantir/NCMEC)

The National Center for Missing & Exploited Children (NCMEC) compiles a wide variety of information from law enforcement, social media, and proprietary databases. Much of this information has traditionally been stored in siloed databases, requiring analysts to manually query each database when investigating a case. The Big Data analytics tool, developed in 2010 by Palantir, empowers NCMEC analysts to query a range of databases simultaneously.

The below case illustrates how the NCMEC uses Big Data to save children:

A 17-year-old girl was reported missing and suspected of being a victim of sex trafficking. Through various searches, a NCMEC analyst was able to find multiple posts online that advertised this missing child for sex. Through information in the ads, the analyst was able to tie them to other posts from the same pimp. The analysis included over 50 advertisements, 9 different females, and a trail covering 5 states. A Link Analysis graph was created using Palantir that helped law enforcement to easily see the large scope of the ring. This insight helped law enforcement link the pimp to a multitude of other crimes and other girls that he victimized.



Source: [Palantir, Using Palantir to Uncover Hidden Links In Missing And Exploited Children Cases](#)

Case Study 13: Human Trafficking (Palantir/Polaris Project)

Human Trafficking is a global problem. According to the *Polaris Project*, each year 21 million people are enslaved worldwide to generate a profit of \$32 billion for their captors. To combat this global problem, organizations like Polaris Project maintain extensive databases of information collected from various public and private sources. The organization reports that it may collect up to 170 different quantitative and qualitative variables per case record, including first-hand data obtained through its National Human Trafficking Resource Center Hotline. In 2013, the NHTRC received 31,945 phone calls, 1,488 e-mails, 1,669 tips from online form submissions, and 787 SMS threads. In order to leverage this vast amount of data, the organization uses the Palantir Gotham analytics platform to track trafficking rings, quickly identify discrete human-trafficking events, and mobilize appropriate response units.



Source: [Polaris Project, Launching the Global Human Trafficking Hotline Network](#)



Appendix C



August 5, 2014

Via e-mail: privacyrfc2014@ntia.doc.gov

National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Avenue N.W., Room 4725
Washington, DC 20230

Re: Privacy RFC 2014

The Future of Privacy Forum (FPF) is a think tank seeking to advance responsible data practices and includes leaders in business, academia, and consumer advocacy. FPF appreciates the opportunity to provide these Comments in response to NTIA's June 5, 2014 Request for Comment (RFC), seeking public input on how developments related to "big data" impact the Consumer Privacy Bill of Rights.¹

Unlocking the value of data and instituting responsible data practices go hand-in-hand, and both have been an important focus of FPF's work since our founding in 2008. FPF recognizes the enormous potential benefits to consumers and to society from big data analytics,² and FPF also understands that taking advantage of big data will require traditional privacy principles to evolve. The Consumer Privacy Bill of Rights likewise endorses a flexible approach to the use of different practices and tools to protect privacy. With respect to big data, flexibility in the application of privacy protections is essential.

FPF supports efforts by the NTIA to consider how the Consumer Privacy Bill of Rights can support uses of big data, and we have previously recognized the need for further conversations about how privacy can be protected in the age of big data.³ Turning to specific questions posed by the RFC:

(2) Should any of the specific elements of the Consumer Privacy Bill of Rights be clarified or modified to accommodate the benefits of big data? Should any of those elements be clarified or modified to address the risks posted by big data?

The Consumer Privacy Bill of Rights recognizes that protecting privacy requires implementing practices that are "comprehensive, actionable, and flexible."⁴ The general principles put forward by

¹ NTIA, Big Data and Consumer Privacy in the Internet Economy, Request for Public Comment, 79 Fed. Reg. 32,714 (June 5, 2014).

² Omer Tene & Jules Polonetsky, *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 NW. J. TECH. & INTELL. PROP. 239, 243-51 (2013).

³ In September 2013, FPF and the Stanford Center for Internet & Society held a day long symposium exploring how to address privacy issues raised by big data. A collection of papers prepared in advance of the event is available at <http://www.futureofprivacy.org/big-data-privacy-workshop-paper-collection/>.

the Administration's privacy framework explicitly afford companies discretion in how they implement them. This flexibility was designed both to promote innovation and to "encourage effective privacy protections by allowing companies, informed by input from consumers and other stakeholders, to address the privacy issues that are likely to be most important to their customers and users, rather than requiring companies to adhere to a single, rigid set of requirements."⁵ For example, while the framework calls for offering individual control of information, it recognizes that this may be impractical in some cases. Instead, organizations are encouraged to embrace other elements of the Consumer Privacy Bill of Rights and augment internal practices in order to adequately protect consumer privacy.⁶

Such flexibility can be accommodated within the established framework provided by the Fair Information Practice Principles (FIPPs). Data innovation requires that the FIPPs are not inflexible "one size fits all" rules but rather must be adapted to the circumstances of data collection and use.⁷ For example, use limitation is a long-standing, valuable FIPP, but it is strained in an age of big data. If use limitation is delimited by the purposes specified at the time of collection, we may never be able to obtain the unexpected benefits that are promised by big data. As the Consumer Privacy Bill of Rights recognizes, privacy is promoted when data are used in ways that respect the context of collection. Context changes over time. Respecting the context of collection must allow room for innovative uses of data that deliver unexpected benefits. And respect for context must be framed in a way that reflects the dynamic nature of social and cultural norms and the subjective nature of consumer trust.

FPF submits that the Consumer Privacy Bill of Rights should be clarified to make explicit the need for flexibility with respect to big data. In particular, and as discussed below and endorsed by the recent White House Big Data Report⁸, big data warrants a shift toward greater emphasis on responsible use. A use-based approach is one that respects the *context* in which information is collected and used, and implements stronger accountability and enforcement mechanisms. As discussed further below, the context principle should encompass factors such as the type of data being used and how the data are being used, including the broader societal benefits that could result and consideration of relevant risks.

Pragmatic contextual considerations can also inform how other existing privacy principles, including notice, choice, and data minimization, are applied in practice. Again, flexibility will be key, and the principles within the Consumer Privacy Bill of Rights should be properly calibrated both to protect privacy and encourage innovative uses of data. Longstanding practices emphasizing limits on collection and use impose rigid rules of the road that dismiss the flexibility inherent in calls for increased transparency and individual control. A flexible approach will allow organizations to

⁴ WHITE HOUSE, CONSUMER DATA PRIVACY IN A NETWORKED WORLD: A FRAMEWORK FOR PROTECTING PRIVACY AND PROMOTING INNOVATION IN THE GLOBAL DIGITAL ECONOMY 9 (2012), <http://www.whitehouse.gov/sites/default/files/privacy-final.pdf> [hereinafter WHITE HOUSE BLUEPRINT].

⁵ *Id.* at 2.

⁶ *Id.* at 9.

⁷ *See, e.g.*, Comments of the Future of Privacy Forum RE: Internet of Things, Project No. P135405, at 3-5 (Jan. 2014), available at http://www.futureofprivacy.org/wp-content/uploads/FPF-IoT-Comments_January-2014.pdf.

⁸ WHITE HOUSE EXEC. OFFICE OF THE PRESIDENT, BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES 56 (May 2014), http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf [hereinafter BIG DATA REPORT].

implement benefit-risk analyses and accountability mechanisms that will be conducive to a considered approach to privacy and big data.

At the same time, big data also highlights the importance of data security. One recent survey suggested that 80% of consumers are more worried that information they share will be “hacked” or stolen, while 16% were more concerned that their information will be shared for commercial purposes or used to target advertising to them.⁹ According to the Federal Trade Commission, identity theft is regularly the top consumer worry,¹⁰ and FTC Chairwoman Edith Ramirez has cautioned that big data will exacerbate the threat posed by data breaches.¹¹ The Consumer Privacy Bill of Rights recognizes the importance of security, but it must remain a top priority moving forward.

(3) Should a responsible use framework, as articulated in Chapter 5 of the Big Data Report, be used to address some of the challenges posed by big data? If so, how might that framework be embraced within the Consumer Privacy Bill of Rights? Should it be? In what contexts would such a framework be most effective? Are there limits to the efficacy or appropriateness of a responsible use framework in some contexts? What added protections do usage limitations or rules against misuse provide to users?

New technologies, including big data and the emerging Internet of Things, are increasingly better suited to a responsible use framework.¹² In particular, one of biggest challenges – and benefits – posed by big data is that much of the new value from data is being discovered in surprising ways.¹³ Accordingly, it is not always possible to provide precise notice in advance of collection on subsequent use of data. Calls for a responsible use framework are a response to concerns that traditional applications of existing privacy principles both constrain valuable uses of data and do not accurately reflect how consumers use new technologies.¹⁴

A responsible use framework is reflected throughout the Consumer Privacy Bill of Rights. Specifically, the Consumer Privacy Bill of Rights calls for respecting the context in which information is collected and used. The principle of respect for context makes companies stewards of consumer data.¹⁵ Context reflects and respects how individuals think about their personal data,¹⁶ and it provides a framework for organizations to ensure data use is in line with consumer expectations.

⁹ Heather Greenfield, Computer & Communication Industry Association, *Major Study Sheds Light On Online Privacy, Security Values, Behavior* (Dec. 20, 2013), <http://www.ccianet.org/2013/12/major-study-sheds-light-online-privacy-security-values-behavior/>.

¹⁰ Colleen Tressler, Fed. Trade Comm'n, *Identity Theft Tops List of Consumer Complaints for 14th Consecutive Year* (Feb. 27, 2014), <http://www.consumer.ftc.gov/blog/identity-theft-tops-list-consumer-complaints-14th-consecutive-year>.

¹¹ Edith Ramirez, Chairwoman, Fed. Trade Comm'n, Keynote Address at the Technology Policy Institute Aspen Forum, *The Privacy Challenges of Big Data: A View from the Lifeguard's Chair* (Aug. 19, 2013), available at www.ftc.gov/public-statements/2013/08/privacy-challenges-big-data-view-lifeguard's-chair.

¹² Christopher Wolf & Jules Polonetsky, *An Updated Privacy Paradigm for the “Internet of Things”* (2013), <http://www.futureofprivacy.org/wp-content/uploads/Wolf-and-Polonetsky-An-Updated-Privacy-Paradigm-for-the-“Internet-of-Things”-11-19-2013.pdf>.

¹³ E.g., VIKTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, *BIG DATA: A REVOLUTION THAT WILL TRANSFORM HOW WE LIVE, WORK, AND THINK* (2013).

¹⁴ FRED CATE & VIKTOR MAYER-SCHÖNBERGER, *DATA USE AND IMPACT GLOBAL WORKSHOP* (Dec. 1, 2013), http://cacr.iu.edu/sites/cacr.iu.edu/files/Use_Workshop_Report.pdf.

¹⁵ WHITE HOUSE BLUEPRINT, *supra* note 4, at 16.

¹⁶ Christopher Wolf, *The Role of Data Use Analysis in Measuring and Protecting Against Privacy Harms 5* (June 5, 2014) (unpublished manuscript, on file with Future of Privacy Forum).

When personal information is used in ways that individuals would reasonably expect, privacy concerns are avoided. Context can be broken down by examining: 1) the type of data being used; 2) the service relationship in which data are collected; 3) how the information is used; 4) the type of device on which the transaction is being made; 5) whether the data was collected passively or actively; and 6) the potential benefits and harms that could result.¹⁷ In addition, the direct use of data on a first-party basis may raise different privacy considerations than data that is shared with third parties or disclosed publicly. The context principle rests on a number of subjective and fact-specific variables, including an individual's level of trust in an organization and his or her perceptions of the value he or she derives from the use of the information.¹⁸ Organizations and policy makers will need to acknowledge and remain aware of ever-shifting social and cultural norms in evaluating contextual use.¹⁹

Respect for context is only part of a responsible use framework. The Big Data Report recognizes that contextual principles may complement emerging use-based approaches to thinking about data.²⁰ There are uses of data that are outside of the traditionally-understood bounds of context that may also have high societal value. Ideally, a use-based approach aspires to consensus around broadly acceptable data uses—or on data uses that are deemed illegitimate, allowing organizations and policy makers to focus on managing the risks associated with middle of the road cases.²¹

When data are used in ways that may be out of context, the results may still be acceptable – and beneficial to consumers. Increasingly, the data collected from a range of apps as well as wearable and other “smart” devices could be used in ways that are out of context but useful to consumers. For example, smart store technologies can leverage the interaction of store Wi-Fi networks and mobile phones. Stores can learn how long consumers are waiting in line to check out and can understand how many consumers are repeat shoppers or which window displays are successful at bringing consumers into the store or to a register.²² Increasingly, retailers hope to use real-time analytics data to better serve customers, make the shopping experience more enjoyable, and better compete with the convenience of simply shopping online.²³ FPF worked with a number of stakeholders to ensure that smart store technologies can be deployed in a way that protects consumer privacy and promote responsible uses of this information.²⁴

Moving forward, when data are used out of context or otherwise in a manner that challenges the principles in the Consumer Privacy Bill of Rights, organizations should demonstrate that their big

¹⁷ *See id.*

¹⁸ Carolyn Nguyen, Director, Microsoft Technology Policy Group, Contextual Privacy, Address at the FTC Internet of Things Workshop (Nov. 19, 2013) (transcript available at http://www.ftc.gov/sites/default/files/documents/public_events/internet-things-privacy-security-connected-world/final_transcript.pdf).

¹⁹ *Id.*

²⁰ The White House Big Data Report recognizes this tension. *See* BIG DATA REPORT, *supra* note 8, at 56.

²¹ SCOTT CHARNEY, MICROSOFT, TRUSTWORTHY COMPUTING NEXT 22 (2012); *see also* Scott Charney, *The Evolving Pursuit of Privacy*, HUFFINGTON POST (Apr. 10, 2014 3:04 PM), http://www.huffingtonpost.com/scott-charney/the-evolving-pursuit-of-p_b_5120518.html.

²² *See* SmartStorePrivacy.com (last visited July 1, 2014).

²³ *See* Cindy Waxer, *Brick and Mortar Retailers Enlist Real-Time Analysis to Counter Online*, DATA INFORMED (Apr. 29, 2014), <http://data-informed.com/brick-mortar-retailers-enlist-real-time-analysis-counter-online/>.

²⁴ Press Release, The Future of Privacy Forum and Sen. Schumer Announce Important Agreement to Ensure Consumers Have Opportunity to “Opt-Out” Before Stores Can Track Their Movement Via Their Mobile Devices (Oct. 22, 2013), <http://www.futureofprivacy.org/2013/10/22/schumer-and-tech-companies-announce-important-agreement-to-ensure-consumers-have-opportunity-to-opt-out-before-stores-can-track-their-movement-via-their-cell-phones/>.

data projects have undergone a benefit-risk analysis that weighs the issues at hand. FPF believes that Chief Privacy Officers, privacy boards, and broader ethics panels can provide a practical outlet to preserve innovation and demonstrate accountability.

In addition to providing a foundation for responsible use, establishing effective enforcement and well-defined accountability mechanisms are also essential to increase the global interoperability of privacy laws.²⁵ The White House's "Privacy Blueprint," which includes the Consumer Privacy Bill of Rights,²⁵ specifically flags the use of multi-stakeholder processes to develop voluntary but enforceable codes of conduct to implement its principles across a range of sectors.²⁶ FPF supports the continued development of these codes and the pathway to interoperability as described in the Administration's Blueprint. Global interoperability is particularly important in the context of big data.²⁷ If big data is balkanized into geographic regions, its value is diminished. Without a common understanding of what constitutes the responsible use of big data, conflicting approaches at the national level can and will take root, further diminishing the economic gains that can be realized through innovation. FPF encourages the promotion and maintenance of existing frameworks that promote interoperability. These frameworks include the U.S.-EU Safe Harbor and the APEC Cross Border Privacy Rules System. Additionally, FPF encourages the development and consideration of new frameworks, such as those that might emerge in ongoing trade negotiations.

(4) What mechanisms should be used to address the practical limits to the “notice and consent” model noted in the Big Data Report? How can the Consumer Privacy Bill of Rights’ “individual control” and “respect for context” principles be applied to big data? Should they be? How is the notice and consent model impacted by recent advances concerning “just in time” notices?

Notice is often considered the most “fundamental” principle of privacy protection.²⁸ Yet there is wide acknowledgement that a privacy framework based solely on notice and choice has significant limitations.²⁹ The vast majority of consumers do not read privacy policies,³⁰ and further, studies have shown that consumers make privacy decisions not based on policies but rather on the context in which they are presented by a use of their data.³¹ More detailed privacy policies should not be the sole solution, and policy makers should be wary of inundating users with more and more notices. Instead, in the age of big data, notice and choice mechanisms will need to be supplemented by additional tools to effectively protect privacy. As a practical matter, connected devices or other “smart” technologies that will not be equipped with interactive screens or other easily accessible user interfaces will prove challenging for traditional implementations of notice and choice. Information

²⁵ WHITE HOUSE BLUEPRINT, *supra* note 4, at 31.

²⁶ *Id.* at 25.

²⁷ The Big Data Report recognizes that “[t]he benefits of big data depend on the global free flow of information.” BIG DATA REPORT, *supra* note 8, at 63.

²⁸ FEDERAL TRADE COMMISSION, PRIVACY ONLINE: A REPORT TO CONGRESS 7 (1998).

²⁹ Fred Cate, *Looking Beyond Notice and Choice*, PRIVACY & SECURITY LAW REPORT (Mar. 29, 2010),

[http://www.hunton.com/files/Publication/f69663d7-4348-4dac-b448-3b6c4687345e/Presentation/PublicationAttachment/dfdad615-e631-49c6-9499-](http://www.hunton.com/files/Publication/f69663d7-4348-4dac-b448-3b6c4687345e/Presentation/PublicationAttachment/dfdad615-e631-49c6-9499-ead6c2ada0c5/Looking_Beyond_Notice_and_Choice_3.10.pdf)

[ead6c2ada0c5/Looking_Beyond_Notice_and_Choice_3.10.pdf](http://www.hunton.com/files/Publication/f69663d7-4348-4dac-b448-3b6c4687345e/Presentation/PublicationAttachment/dfdad615-e631-49c6-9499-ead6c2ada0c5/Looking_Beyond_Notice_and_Choice_3.10.pdf) (citing Former FTC Chairman Jon Liebowitz conceding

that the “notice and choice” regime offered by the FIPPs hasn’t “worked quite as well as we would like.”).

³⁰ Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 I/S: J. L. & POL’Y FOR INFO. SOC’Y 543 (2008).

³¹ See, e.g., Alessandro Acquisti et al., What Is Privacy Worth? 27-28 (2010) (unpublished manuscript), *available at* <http://www.heinz.cmu.edu/~acquisti/papers/acquisti-ISR-worth.pdf>.

collected in “public” spaces and used for data analytics also suggest the need for novel approaches to privacy protection.

Flexibility will be especially important with respect to the concepts of notice and consent. Organizations need to think creatively about how to provide consumers with meaningful insight into commercial data practices when necessary, and regulators and policymakers should encourage these efforts.³² Techniques to inform consumers of data practices might include symbols, short phrases, colors, diagrams, or any of the tools otherwise available to designers seeking to provide users with an engaging user experience. In some cases, just in time notices may be most effective. In others, it may be important to explain to consumers the benefit of the bargain in a document they can examine in advance. In the end, design features that “communicate” information to users may be more helpful than traditional notice models. The Consumer Bill of Rights should continue to express support for these innovative developments through flexible multi-stakeholder processes, as opposed to establishing overly restrictive rules that limit the collection and use of data.

Public-facing efforts to inform consumers about big data may improve awareness and understanding more than rigid applications of the notice and consent model. The NTIA should recognize and encourage efforts by companies to engage consumers in meaningful conversations where both parties’ interests and expectations can be aligned. FPF has previously described the benefits of data “featurization,” transforming data analysis into a consumer-side application by granting individual access to their personal data in intelligible, machine-readable forms.³³ Services such as personal clouds or data stores will allow individuals to contract with third-parties who would get permission to selectively access certain categories of their data to provide further analysis or value-added services.³⁴ “Featurization” could allow individuals to declare their own policies, preferences and terms of engagement, and do it in ways that can be automated both for them and for the companies they engage.³⁵

11) How significant are the privacy risks posed by re-identification of de-identified data? How can de-identification be used to mitigate privacy risks in light of the analytical capabilities of big data? Can particular policy safeguards bolster the effectiveness of de-identification? Does the relative efficacy of de-identification depend on whether it is applied to public or private data sets? Can differential privacy mitigate risks in some cases? What steps could the government or private sector take to expand the capabilities and practical application of these techniques?

³² Testimony Before the California State Assembly Joint Committee Hearing on Privacy (Dec. 12, 2103) (statement of Jules Polonetsky, Executive Director, Future of Privacy Forum, at 3), *available at* http://www.futureofprivacy.org/wp-content/uploads/CA-Assembly-Hearing-Privacy-Policies_Does-Disclosure-Transparency-Adequately-Protect-Consumers-Privacy-Final.pdf (citing Susanna Kim Ripken, *The Dangers and Drawbacks of the Disclosure Antidote: Toward a More Substantive Approach to Securities Regulation*, 58 BAYLOR L. REV. 139, 147 (2006) (calling for regulators to “lay aside the gospel of disclosure in favor of more substantive laws that regulate conduct directly”).

³³ See, e.g., Omer Tene & Jules Polonetsky, *A Theory of Creepy: Technology, Privacy and Shifting Social Norms*, 16 YALE J.L. & TECH. 59 (2013).

³⁴ Tene & Polonetsky, *supra* note 2, at 263-70.

³⁵ The rise of privacy and reputation management vendors points to a future where organizations will be able to unlock additional value by collaborating with their users. One interesting first step is the launch of “About the Data” by Acxiom, the nation’s largest data broker. “About the Data” is a consumer-facing tool that gives individuals control over certain categories of information (such as personal characteristics, interests, and finances) gathered by Acxiom. The site allows consumers to correct information, suppress any data they see, or opt-out of Acxiom’s marketing profile system altogether via an approachable user interface.

As the President's Council of Advisors on Science and Technology Big Data Report recognizes, de-identification remains a useful tool to protect privacy in an age of big data.³⁶ However, the report remains considerably skeptical about the technical ability to ensure de-identified data cannot be re-identified.³⁷ These concerns result from the fact that researchers have routinely discovered ways to re-identify de-identified datasets that have been released publicly. But these examples often overstate the failure of de-identification.³⁸ De-identification must be assessed within the context of the type of data involved and how it is being used, and not the merely the mathematical possibility that data may be re-identified.³⁹

Prominent examples of successful re-identification demonstrate the challenges of perfectly de-identifying datasets that are *released publicly*. Once data are made publicly available, there is an opportunity for any attacker with the time, resources, or technological capability can attempt to re-identify information. While perfect de-identification may not be technically achievable, the actual risk of re-identification of individuals from properly de-identified data is incredibly low.

Much of our discourse around de-identification focuses on the technical possibility of re-identification and the assumption that all data will be made publicly available.⁴⁰ This emphasis does not describe the entire universe of data that exists today. It is essential to differentiate the risks between privately-held and publicly-available information. Tying discussions of de-identification to its effectiveness with public datasets dismisses its value for safeguarding data that is either kept internally or only shared with a limited audience. To be sure, when de-identified data are disclosed publicly with only minimal technical controls under controlled conditions in place to mask identifiers, the risks of identification are much greater.

The successful re-identification of Netflix users is an example of how low the actual risk of re-identification is with respect to publicly released data.⁴¹ In a recent study, researchers compared data widely released by Netflix with records available on the Internet Movie Database in order to uncover the identities of Netflix users. However, only two out of 480,189 Netflix users were successfully identified with this method. That equates to a 0.0004 percent rate of re-identification. That is a significantly lower risk than an individual's lifetime risk of being struck by lightning, which is approximately one in seven thousand.⁴² Though one should not ignore the invasion of privacy

³⁶ PRESIDENT'S COUNCIL OF ADVISORS ON SCIENCE & TECH., BIG DATA AND PRIVACY: A TECHNOLOGICAL PERSPECTIVE 38 (May 2014), http://www.whitehouse.gov/sites/default/files/microsites/ostp/PCAST/pcast_big_data_and_privacy_-_may_2014.pdf.

³⁷ *Id.*

³⁸ *E.g.* ANN CAVOUKIAN & DANIEL CASTRO, BIG DATA AND INNOVATION, SETTING THE RECORD STRAIGHT: DE-IDENTIFICATION DOES WORK (June 2014), <http://www2.itif.org/2014-big-data-deidentification.pdf>.

³⁹ Comments of the Future of Privacy Forum to the Federal Communications Commission, WC Docket No. 13-306 (Jan. 2014), *available at* <http://www.futureofprivacy.org/wp-content/uploads/01-17-2014-FPF-Comments-to-the-FCC.pdf>.

⁴⁰ Yianni Lagos & Jules Polonetsky, *Public vs. Nonpublic Data: The Benefits of Administrative Control*, 66 STAN. L. REV. ONLINE 103 (2013), <http://www.stanfordlawreview.org/online/privacy-and-big-data/public-vs-nonpublic-data>.

⁴¹ *Id.* at 6 (citing Arvind Narayanan and Vitaly Shmatikov, Robust De-anonymization of Large Sparse Datasets, Proceedings of the 2008 IEEE Symposium on Security and Privacy (2008)).

⁴² Joseph Jerome, *Making Perfect De-Identification the Enemy of Good De-Identification*, FUTURE OF PRIVACY FORUM (June 19, 2014), <http://www.futureofprivacy.org/2014/06/19/making-perfect-de-identification-the-enemy-of-good-de-identification/>.

experienced by these two individuals, the chief examples of de-identification's failings ultimately produced minimal privacy harms.

Moreover, the data released by Netflix could have been more effectively de-identified using methods such as the HIPAA de-identification Safe Harbor standards. In other words, the better conclusion to be drawn from the Netflix example is not that de-identification is ineffective, but rather that a more rigorous standard should have been applied to the dataset.⁴³

De-identification must be assessed in context, and when appropriate technical standards are combined with reasonable legal and administrative safeguards, de-identification is an important tool to protect the privacy of individuals. Policy makers should recognize that de-identification techniques encompass a range of different solutions that can be deployed depending upon the potential threats at stake or proposed use of data.

For information that is not made public, practical de-identification can be supported by an array of safeguards including access controls and restrictions, contractual data use restrictions, and data deletion protocols that can augment technical measures in order to protect privacy. When these tools used in combination, a bad actor must circumvent administrative restraints *and* then re-identify any data before getting any value from his malfeasance – which, in practice, is difficult to accomplish.⁴⁴

The Federal Trade Commission's reasonableness standard for de-identification is highly informed by non-technical factors, specifically the "nature of the data at issue and the purposes for which it will be used."⁴⁵ As the FTC suggests, the level of technical de-identification should take into account the sensitivity and potential value of the data at stake. The level of de-identification should be calibrated to consider the sensitivity of information, the level of identifiability of the data, the potential privacy risks at stake, *and* data utility for researchers and industry. The FTC standard also recognizes the importance of good internal processes and controls on the sharing of data with third parties. FPF believes that incorporating a benefit-risk assessment will help organizations tailor the level of de-identification necessary for particular uses of data.

13) Can accountability mechanisms play a useful role in promoting socially beneficial uses of big data while safeguarding privacy? Should ethics boards, privacy advisory committees, consumer advisory boards, or Institutional Review Boards (IRBs) be consulted when practical limits frustrate transparency and individuals' control over their personal information? How could such entities be structured? How might they be useful in the commercial context? Can privacy impact assessments and third-party audits complement the work of such entities? What kinds of parameters would be valuable for different kinds of big data analysts to consider, and what kinds of incentives might be most effective in promoting their consideration?

As discussed above, a responsible use framework will require organizations to develop new accountability mechanisms to ensure that they are managing personal information is responsible –

⁴³ CAVOUKIAN & CASTRO, *supra* note 38, at 6. Further, it is worth recognizing that the content of Netflix' records by themselves were not sufficient to re-identify anyone – it required access to information beyond the records themselves.
⁴⁴ *Id.*

⁴⁵ FED. TRADE COMM'N, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID Change 37 (Mar. 2012), <http://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf>.

and ethical – ways. The NTIA should recognize the important role that internalized and formal review processes can have in considering innovative data projects and identifying any potential risks.⁴⁶ In addition, big data “algorithmists” or other professionals could also function inside organizations to evaluate the selection of data sources, the choice of analytical tools, and the interpretation of any predictive results.⁴⁷ As organizations increasingly face interesting new proposals for using data, these professionals could operate across the public and private sectors and conduct risk-benefit analyses of data uses.

Industry increasingly faces ethical considerations over how to minimize data risks while maximizing benefits to all parties.⁴⁸ Big data projects sometimes raise issues that transcend traditional privacy concerns and implicate broader policy concerns about new forms of discrimination or filter bubbles that could chill democratic discourse.⁴⁹ As the Big Data Report recognizes, there is a potential tension between socially beneficial and privacy invasive uses of information in everything from educational technology to consumer generated health data. The advent of big data requires active engagement by both internal and external stakeholders to increase transparency, accountability and trust.

A documented review process may serve as an effective tool to infuse ethical considerations into data analysis without requiring radical changes to the business practices or innovators or industry in general.⁵⁰ Institutional review boards (IRBs), which are the chief regulatory response to decades of questionable ethical decisions in the field of human subject testing, provide a useful precedent for focusing on good process controls as a way to address potential privacy concerns. While IRBs have become a rigid compliance device and would be inappropriate for wholesale use in big data decision-making,⁵¹ they could provide a useful template for how projects can be evaluated based on prevailing community standards and subjective determinations of risks and benefits, particularly in cases involving greater privacy risks. Using an IRB model as inspiration, big data may warrant the creation of new advisory processes within organizations to more fully consider ethical questions posed by big data.

Traditional IRBs provide only one example of what this might look like, however. Privacy officers and other internal processes, such as privacy impact assessments and review boards, can provide an important outlet for evaluating certain categories of data uses. Documentation efforts could be particularly important when a proposed use might exceed the principles of the Consumer Privacy Bill of Rights. For example, a company may believe that keeping data beyond the minimum time needed may yield important benefits or may discover a use that was not specified in a notice. By carefully considering the potential risks and benefits, a company should be able to advance such a use and be consistent with any proposed law.

Moving forward, broader big data ethics panels could provide a commonsense response to public concerns about data misuse. While these institutions could provide a further expansion of the role of

⁴⁶ See Ryan Calo, *Consumer Subject Review Boards: A Thought Experiment*, 66 STAN. L. REV. ONLINE 97 (2013).

⁴⁷ SCHÖNBERGER & CUKIER, *supra* note 13, at 180.

⁴⁸ danah boyd, *What Does the Facebook Experiment Teach Us?*, MEDIUM (July 1, 2014), <https://medium.com/message/what-does-the-facebook-experiment-teach-us-c858c08e287f>.

⁴⁹ Jules Polonetsky & Omer Tene, *The Facebook Experiment: Gambling? In This Casino?*, RE/CODE (July 2, 2014 11:55 AM PDT), <http://recode.net/2014/07/02/the-facebook-experiment-is-there-gambling-in-this-casino/>.

⁵⁰ Matthew Salganik, *After the Facebook Emotional Contagion Experiment: A Proposal for a Positive Path Forward*, FREEDOM TO TINKER (July 7, 2014), <https://freedom-to-tinker.com/blog/mjs3/after-the-facebook-emotional-contagion-experiment-a-proposal-for-a-positive-path-forward/>.

⁵¹ See Calo, *supra* note 46.

privacy professionals within organizations, they might also provide a forum for a diversity of viewpoints inside and out of organizations. Ethics reviews could include members with different backgrounds, training, and experience, and could seek input from outside actors including consumer groups and regulators. It is likely that these review boards will vary from organization to organization, and standardization would present a unique set of challenges, but organizations should be encouraged to develop structures and personnel to grapple with big data and provide an important check on any data misuse.

Any successful approach to big data must recognize the different ways that data can now be used. Many uses of big data are machine-to-machine or highly aggregated and do not implicate privacy concerns. It is also the case that many new uses of data are marginal and should not require enhanced review processes. Furthermore, the process of mitigating risks is often iterative and nuanced to reflect strategies that can well address minor changes to products or services.

That said, the more challenging areas of big data should be guided by a risk-benefit analysis that takes into account exactly how the benefits of big data will be distributed. So far, our procedural frameworks are largely focused on traditional privacy risks and assessing what measures can be taken to mitigate those risks. In 2010, for example, the Department of Commerce's Internet Policy Task Force endorsed the use of privacy impact assessments (PIAs) both to help organizations decide whether it is appropriate to engage in innovative data uses and to identify alternative approaches that could reduce relevant privacy risks.⁵² However, human research IRBs also take into account anticipated benefits and even the importance of any knowledge that may result from research.⁵³

Organizations and privacy professionals have become experienced at evaluating risk, but they should also engage in a rigorous data benefit analysis in conjunction with traditional privacy risks assessments. FPF suggests that organizations could develop procedures to assess the "raw value" of a data project, which would require organizations to identify the nature of a project, its potential beneficiaries, and the degree to which those beneficiaries would benefit from the project. This value would be discounted by the probability that a big data benefit can actually be accomplished. While there are no definitive rules of what probability of benefit is needed to overcome a presumption against exposing a beneficiary to privacy risk, it is clear that the mere assertion that a product or service can be improved is inadequate; yet demanding organizations provide proof beyond any doubt would be an impossible standard. Further, any analysis would need to take into account different community standards and expectations, as well as an inherent degree of subjectivity in assessing the relative merits of various projects. FPF's initial guidance on this type of analysis is attached as Appendix A.

However, this limitation has not stopped the federal government from developing structured processes to measure project benefits against risks. The White House Office of Management and Budget, for example, requires federal agencies to engage in comprehensive cost-benefit analysis. This analysis recognizes that even where values are not computable, "efforts to measure it can produce useful insights even when the monetary values of some benefits or costs cannot be determined."⁵⁴

⁵² U.S. DEP'T OF COMMERCE, INTERNET POL'Y TASK FORCE, COMMERCIAL DATA PRIVACY AND INNOVATION IN THE INTERNET ECONOMY: A DYNAMIC POLICY FRAMEWORK 34-35 (2010), <http://www.commerce.gov/sites/default/files/documents/2010/december/iptf-privacy-green-paper.pdf>.

⁵³ 45 CFR § 46.111.

⁵⁴ OFFICE OF MANAGEMENT & BUDGET, CIRCULAR NO. A-94 REVISED, MEMORANDUM FOR HEADS OF EXECUTIVE DEPARTMENTS AND ESTABLISHMENTS, GUIDELINES AND DISCOUNT RATES FOR BENEFIT-COST ANALYSIS OF FEDERAL PROGRAMS (Oct. 29, 1992), http://www.whitehouse.gov/omb/circulars_a094.

Conclusion

Big data provides a necessary catalyst for many important conversations about privacy. From advancing practical de-identification to infusing ethical decision-making into data projects, big data will require organizations to create new processes and innovative tools to engender trust and address privacy.

The Consumer Privacy Bill of Rights offers a helpful, flexible framework to resolve these challenges. More work is needed to ensure the principles in the Consumer Privacy Bill of Rights are operationalized by organizations in a common-sense fashion, but we already have a path forward to thoughtfully address the concerns posed by big data.

FPF thanks the National Telecommunications and Information Administration for considering these Comments, and we look forward to further engagement and collaboration on the issue of big data.

Sincerely,

Jules Polonetsky
Director and Co-Chair
Future of Privacy Forum

Christopher Wolf
Founder and Co-Chair
Future of Privacy Forum

Josh Harris
Policy Director
Future of Privacy Forum

Joseph Jerome
Policy Counsel
Future of Privacy Forum