



fpf.org

1400 I Street NW | Suite 450 | Washington, DC 20005 · 202-768-8950 ·

October 27, 2017

Federal Trade Commission
600 Pennsylvania Avenue NW
Washington, DC 20530

Re: Comments in Advance of December 12, 2017 Informational Injury Workshop P175413

Dear Colleagues,

On behalf of the Future of Privacy Forum, we are pleased to submit these comments in response to the Federal Trade Commission (FTC)'s Request for Comment in advance of the December 12, 2017 Informational Injury Workshop. The Future of Privacy Forum (FPF) is a Washington, DC based non-profit organization that serves as a catalyst for privacy leadership and scholarship, and advances principled data practices in support of emerging technologies.¹ We thank the FTC for requesting comments on the broad range of legal and policy issues related to informational injury. Our comments focus on two topics we regard as particularly important: describing the harms that can arise from automated decision-making, and highlighting existing risk-based privacy analyses.

In recent years, the FTC has been at the forefront of national conversations around big data, algorithmic decision-making, and the implications of innovative and emerging technologies.² In these conversations, the benefits of data collection and analysis can be compelling, including revolutionary advances in health care, improving educational outcomes, increasing access to employment, and opening the door to evidence-based policymaking that can lead to more insightful and justice-oriented public policy.³ Nonetheless, as automated decision-making becomes more common, it is increasingly important to closely examine the legal and ethical issues related to the use of sensitive data for hiring, policing, public benefits decisions, marketing, and other purposes.

Categorizing Algorithmic Harms

Analysis of personal data can be used to improve services, promote inclusion, and combat discrimination.⁴ However, such analysis can also create valid concerns about differential treatment of individuals or disparate impacts on vulnerable communities. These concerns can be amplified when algorithmic decision-making uses sensitive data (such as race, gender, or familial status), impacts protected classes, or affects individuals'

¹ The views herein do not necessarily reflect those of the Advisory Board or supporters of the Future of Privacy Forum.

² See, e.g., FEDERAL TRADE COMMISSION, *BIG DATA: A TOOL FOR INCLUSION OR EXCLUSION?* (Jan. 2016), <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf>; Press Release, *FTC to Host Fall Seminar Series on Emerging Consumer Technology Issues* (March 31, 2016), <https://www.ftc.gov/news-events/press-releases/2016/03/ftc-host-fall-seminar-series-emerging-consumer-technology-issues>.

³ See EXEC. OFF. OF THE PRESIDENT, *BIG DATA: A REPORT ON ALGORITHMIC SYSTEMS, OPPORTUNITY, AND CIVIL RIGHTS* (May 2016) (at 10-22) (discussing opportunities and challenges using data to expand access to credit; reduce employment discrimination; increase educational opportunities; and build trust in law enforcement). See also REPORT OF THE COMMISSION ON EVIDENCE-BASED POLICYMAKING: *THE PROMISE OF EVIDENCE-BASED POLICYMAKING* (Sept 2017), <https://www.cep.gov/content/dam/cep/report/cep-final-report.pdf>.

⁴ FUTURE OF PRIVACY FORUM & ANTI-DEFAMATION LEAGUE, *BIG DATA: A TOOL FOR FIGHTING DISCRIMINATION AND EMPOWERING GROUPS* (2014), <https://fpf.org/wp-content/uploads/Big-Data-A-Tool-for-Fighting-Discrimination-and-Empowering-Groups-FINAL.pdf>

eligibility for housing, employment, or other core services. Recent controversies have highlighted legal and ethical issues raised by organizations using sensitive data for marketing and other purposes.⁵

Despite the increasing relevance of these concerns, conversations around the potential harms of algorithmic decision-making often become mired in definitional challenges. Analyses often fail to separate harms from causes and solutions, and sometimes conflate digital causes with human biases that may already shape both input and algorithmic design. To inform discussions on this topic, FPF conducted a review of existing literature with the goal of categorizing the harms that leading thinkers have identified as potentially resulting from automated decision-making. Please see **Appendix A** for a literature review that informs this taxonomy. We hope that a clear articulation of the universe of proposed harms will empower those seeking solutions to identify the harms that may warrant legal recognition, those that may be *de minimus* or easy to mitigate, or those for which mitigation could have deleterious effects by limiting economic competition or freedom of speech.

In our preliminary review of the relevant literature and public policy regarding automated decision-making, we found that the concerns identified by leaders in this space fall into four broad categories of potential harms: (1) loss of opportunity; (2) economic loss; (3) social stigmatization; and (4) loss of liberty. Examples of “loss of opportunity” harms include informational injuries related to employment, insurance and social benefits, housing, and education. “Economic loss” harms relate to credit, differential pricing, and narrowing of choice. “Social stigmatization” harms arise from network and filter bubbles, dignitary harms and stereotype reinforcement, the constraints of bias, and confirmation bias. “Loss of liberty” harms include the effects of surveillance, suspicion, and incarceration. Depending on the context and circumstances, we determined that each of these categories of harms can accrue to individuals, groups, or society as a whole. Notably, not all harms described in existing literature will necessarily be legally cognizable – although they may be widely considered unfair – while some may already be illegal under existing laws.

Even for companies seeking to avoid these harms, mitigation efforts are often hindered by the myriad of challenging issues involved. Companies, regulators, enforcement agencies, and individuals can bring different mitigation strategies to bear. To capture the universe of potential approaches, we surveyed the literature for solutions that experts and policymakers have proposed to mitigate or prevent these harms.

These strategies generally fall into one of four categories: (1) algorithmic design solutions; (2) business process solutions; (3) legal and policy solutions; or (4) data methods solutions. “Algorithmic design solutions” include approaches to matching data sets accurately, validating correlative and causal links, and training algorithms to avoid disparate mistreatment. “Business process solutions” can include the creation of ethical processes or reviews for big data research, and other opportunities for human oversight of algorithms. “Legal and policy solutions” could include “rights to explanation” of automated decision-making, regulation of data brokers, or data protection impact assessments (DPIAs). “Data methods solutions” focus on, for example, ensuring that data being used for inputs to these algorithms is accurate, and not biased from the start at collection—whether from selection bias or historical bias.

As with harms, these potential solutions describe the universe of proposals rather than specific recommended solutions. It is also important to recognize that proposed solutions may sometimes impact other important values, such as freedom of speech or economic competition. Their use may need to be considered on a case-by-case basis and by a balancing of the benefits and risks of intervention.

Risk-Based Analyses in Law and Policy

The challenges of conceptualizing informational injury are increasingly relevant as risk-based privacy analyses become more common in law, policy, and internal business practices. One long-standing legal basis for

⁵ See, e.g., Charles Duihigg, *How Companies Learn Your Secrets*, NEW YORK TIMES (Feb. 16, 2012), <http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>; Jennifer Valentino-DeVries et al., *Websites Vary Prices, Deals Based on Users' Information*, WALL STREET JOURNAL (Dec. 24, 2012), <https://www.wsj.com/articles/SB1000142412788732377204578189391813881534>.

processing data in the European Union is the “legitimate interests” framework,⁶ which has similarities to the FTC’s unfairness analysis under Section 5 of the FTC Act.⁷ Under this basis for lawful processing, companies may engage in lawful data processing if their legitimate interests are not “overridden by the interests or fundamental rights and freedoms of the data subject.”⁸

In addition, under the General Data Protection Regulation (GDPR) that will come into effect in May 2018, companies are required to carry out a data protection impact assessment if data processing is “likely to result in a high risk to the rights and freedoms of natural persons.”⁹ The determination of whether processing is likely to result in “high risk” can be challenging, and recent guidance from the Article 29 Data Protection Working Party provides detailed recommendations and real-world examples of potentially “high risk” processing, including automated decision-making with legal or other similar significant effects.¹⁰

In each of these benefit-risk analyses, the underlying risk relies on an accurate assessment of the nature of informational injuries. Risks can be defined broadly, as can be seen in the expansive view of privacy harms that largely prevails in the EU, or more narrowly, as is often seen in the United States. Meanwhile, because existing risk mitigation strategies often focus exclusively on privacy risks – not accounting for a tremendous variance in potential anticipated benefits – they may not always be helpful to companies or policymakers attempting to decide whether privacy risks are truly outweighed by countervailing benefits to individuals or society.¹¹

Conclusion

We see a promising set of solutions arising in literature and regulatory conversations on the topic of automated decision-making and risk-based analyses, and we look forward to a robust conversation on these issues at the upcoming FTC workshop. The Commission is well positioned to lead these conversations, and to encourage stakeholders to embrace new processes for addressing these challenges while ensuring the benefits of big data and automated decision-making. FPF thanks the Commission for this opportunity to submit comments and we look forward to future engagement on this topic.

Sincerely,

Lauren Smith, Policy Counsel, lsmith@fpf.org
Stacey Gray, Policy Counsel, sgray@fpf.org

Future of Privacy Forum
www.fpf.org

⁶ See ARTICLE 29 WORKING PARTY, OPINION 06/2014 ON THE NOTION OF LEGITIMATE INTERESTS OF THE DATA CONTROLLER UNDER ARTICLE 7 OF DIRECTIVE 95/46/EC (adopted April 9, 2014), <http://www.dataprotection.ro/servlet/ViewDocument?id=1086>.

⁷ 15 U.S.C. §45.

⁸ E.U. General Data Protection Regulation, Article 6(1)(f) (2016). See also *supra*, note 6 (at 4-8) (describing how the concepts of lawfulness and legal grounds for processing, including legitimate interests, have developed in the European Union prior to the General Data Protection Regulation).

⁹ E.U. General Data Protection Regulation, Article 35(1) (2016).

¹⁰ ARTICLE 29 WORKING PARTY, GUIDELINES ON DATA PROTECTION IMPACT ASSESSMENT (DPIA) AND DETERMINING WHETHER PROCESSING IS “LIKELY TO RESULT IN A HIGH RISK” FOR THE PURPOSES OF REGULATION 2016/679 (as last Revised and Adopted on 4 October 2017).

¹¹ See FUTURE OF PRIVACY FORUM, BENEFIT-RISK ANALYSIS FOR BIG DATA PROJECTS (Sept 2014), at 4 (“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. . . . What is needed is a more thorough vocabulary of big data benefits.”), https://fpf.org/wp-content/uploads/FPF_DataBenefitAnalysis_FINAL.pdf.

Appendix A

Reviewed Literature

The alphabetized list below captures the literature FPF has reviewed to date for this effort. We welcome suggestions for further materials to review.

- Aaron Reike, *Don't let the hype over "social media scores" distract you*, Equal Future (2016).
- Alessandro Acquisti & Christina Fong, *An Experiment in Hiring Discrimination via Online Social Network*, presented at Privacy Law Scholars Conference (2016).
- Alethea Lange et al., *A User-Centered Perspective on Algorithmic Personalization*, presented at the Fed. Trade Comm'n PrivacyCon Conference (2017).
- Allan King & Marko Mrkonich, *"Big Data" and the Risk of Employment Discrimination*, 68 Okla. L. Rev. 555 (2016).
- Andrew Tutt, *An FDA for Algorithms*, 67 Admin. L. Rev. 1 (2016).
- Aniko Hannak et al., *Bias in Online Freelance Marketplaces: Evidence from TaskRabbit*, presented at the Workshop on Data and Algorithmic Transparency (Nov. 2016).
- Cathy O'Neil, *Weapons of Math Destruction* (2016).
- Christian Sandvig et al., *Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms*, presented at the Int'l Comm'n Ass'n Conference on Data and Discrimination: Converting Critical Concerns into Productive Inquiry (2014).
- Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 Wash. L. Rev. 1 (2014).
- Exec. Off. of the President, *Big Data: Seizing Opportunities, Preserving Values* (2014).
- Exec. Off. of the President, *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights* (2016).
- Federal Trade Commission, *Big Data: A Tool for Inclusion or Exclusion?* (Jan. 2016).
- Frank Pasquale & Danielle Keats Citron, *Promoting Innovation While Preventing Discrimination: Policy Goals for the Scored Society*, 89 Wash. L. Rev. 1413 (2014).
- Jennifer Valentino-Devries, Jeremy Singer-Vine, Ashkan Soltani, *Websites Vary Prices, Deals Based on Users' Information*, Wall St. J. (Dec. 24, 2012).
- Joshua Kroll et al., *Accountable Algorithms*, 165 U. Penn. L. Rev. 633 (2016).
- Juhi Kulshrestha et al., *Quantifying Search Bias: Investigating Sources of Bias for Political Searches in Social Media*, presented at the Workshop on Data and Algorithmic Transparency (2016).
- Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C.L. Rev. 93 (2014).
- Latanya Sweeney, *Discrimination in Online Ad Delivery*, Commc'ns of the Ass'n of Computing Machinery (2013).
- Lee Rainie & Jana Anderson, *Code-Dependent: Pros and Cons of the Algorithm Age*, Pew Research Center (2017).
- Mark MacCarthy, *Student Privacy: Harm and Context*, 21 Int'l Rev. of Info. Ethics 11 (2014).
- Mary Madden, Michele Gilman, Karen Levy & Alice Marwick, *Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans*, Wash. U. L. Rev. __ (forthcoming) (Mar. 2017).
- Megan Garcia, *How to Keep Your AI From Turning Into a Racist Monster*, Wired (2017).
- Moritz Hardt, Eric Price & Nathan Srebro, *Equality of Opportunity in Supervised Learning*, presented at the Conference on Neural Info. Processing Sys. (2016).
- Motahhare Eslami et al., *Reasoning about Invisible Algorithms in the News Feed*, presented at the Ass'n of Computing Machinery Special Interest Gp. on Computer-Human Interaction (2015).
- Muhammad Zafar et al., *Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment*, presented at the Int'l World Wide Web Conference (2017).
- Nanette Byrnes, *Why We Should Expect Algorithms to be Biased*, MIT Technology Review (2016).
- New America & Open Tech. Inst., *Data and Discrimination: Collected Essays* (S.P. Gangadharan, Ed. 2014).
- Pam Dixon & Robert Gellman, *The Scoring of America: How Secret Consumer Scores Threaten Your Privacy and Your Future*, World Privacy Forum (2014).
- Pauline Kim, *Data-Driven Discrimination at Work*, 59 William & Mary L. Rev. __ (2017).
- Peter Swire, *Lessons From Fair Lending Law for Fair Marketing and Big Data* (2014)
- ProPublica, *Machine Bias Investigative Series*, <https://www.propublica.org/series/machine-bias>
- Sandra Wachter, Brent Mittelstadt, & Luciano Floridi, *Why a right to explanation of automated decision making does not exist in the General Data Protection Regulation* (2016).
- Solon Barocas & Andrew Selbst, *Big Data's Disparate Impact*, 104 Calif. L. Rev. 671 (2016).
- Upturn, *Civil Rights, Big Data, and Our Algorithmic Future* (2014).