

Judged by the Tin Man: Individual Rights in the Age of Big Data

Omer Tene and Jules Polonetsky*

"How about my heart?" asked the Tin Woodman.

*'Why, as for that,' answered Oz, 'I think you are wrong to want a heart. It makes most people unhappy. If you only knew it, you are in luck not to have a heart.'*¹

Introduction

Big data, the enhanced ability to collect, store and analyze previously unimaginable quantities of data in tremendous speed and with negligible costs,² delivers immense benefits in marketing efficiency, healthcare, environmental protection, national security and more.³ While some privacy advocates may dispute the merits of sophisticated behavioral marketing practices or debate the usefulness of certain data sets to efforts to identify potential terrorists,⁴ few remain indifferent to the transformative value of big data analysis for government, science and society at large.⁵ At the same time, even big data evangelists should recognize the potentially ominous social ramifications of a surveillance society governed by heartless algorithmic machines.⁶

In this essay, we present some of the privacy and non-privacy risks of big data as well as directions for potential solutions. In a previous paper, we argued that the central tenets of the current privacy framework, the principles of data

* Omer Tene is Vice Dean of the College of Management Haim Striks School of Law, Israel and a Senior Fellow at the Future of Privacy Forum; Jules Polonetsky is Co-chair and Executive Director of the Future of Privacy Forum.

¹ L. Frank Baum, *THE WONDERFUL WIZARD OF OZ*, Chapter 15 (1900).

² McKinsey Global Institute defines big data as data for which the "size is beyond the ability of typical database software tools to capture, store, manage, and analyze." McKinsey Global Institute, *BIG DATA: THE NEXT FRONTIER FOR INNOVATION, COMPETITION AND PRODUCTIVITY*, June 2011, http://www.mckinsey.com/~media/McKinsey/dotcom/Insights%20and%20pubs/MGI/Research/Technology%20and%20Innovation/Big%20Data/MGI_big_data_full_report.ashx. Big data, however, is typically characterized not only by volume but also by velocity (speed of processing) and variety (the capability to link diverse data sets and process unstructured data) (together referred to as the "three v's"). See Gartner Press release, Gartner Says Solving 'Big Data' Challenge Involves More Than Just Managing Volumes of Data, <http://www.gartner.com/it/page.jsp?id=1731916>.

³ Omer Tene & Jules Polonetsky, *Privacy in the Age of Big Data: A Time for Big Decisions*, 64 *STAN. L. REV. ONLINE* 63 (2012).

⁴ See Julia Angwin, *U.S. Terrorism Agency to Tap a Vast Database of Citizens*, *WSJ*, December 13, 2012, <http://online.wsj.com/article/SB10001424127887324478304578171623040640006.html>; R. Jeffrey Smith, *Senate Report Says National Intelligence Fusion Centers Have Been Useless*, *FOREIGN POLICY*, October 3, 2012, http://www.foreignpolicy.com/articles/2012/10/03/senate_report_says_national_intelligence_fusion_centers_have_been_useless.

⁵ For some examples see World Economic Forum, *PERSONAL DATA: THE EMERGENCE OF A NEW ASSET CLASS*, 2011, http://www3.weforum.org/docs/WEF_ITTC_PersonalDataNewAsset_Report_2011.pdf.

⁶ Paul Ohm, *Don't Build a Database of Ruin*, *HARVARD BUSINESS REVIEW*, August 23, 2012, http://blogs.hbr.org/cs/2012/08/dont_build_a_database_of_ruin.html.

minimization and purpose limitation, are severely strained by the big data technological and business reality.⁷ Here, we assess some of the other problems raised by pervasive big data analysis. To highlight the ethical and moral dilemmas, we sometimes refer to big data algorithms as “the machine” (which is more elegant than “zombie”, though less animated than the “tin man” in the title).⁸ In their book, “*A Legal Theory for Autonomous Artificial Agents*”, Samir Chopra and Larry White note that “as we increasingly interact with these artificial agents in unsupervised settings, with no human mediators, their seeming autonomy and increasingly sophisticated functionality and behavior, raises legal and philosophical questions”.⁹ In this article we argue that the focus on the machine is a distraction from the debate surrounding data driven ethical dilemmas, such as privacy, fairness and discrimination. The machine may exacerbate, enable, or simply draw attention to the ethical challenges, but it is humans who must be held accountable.

Instead of vilifying machine-based data analysis and imposing heavy-handed regulation, which in the process will undoubtedly curtail highly beneficial activities,¹⁰ policymakers should seek to devise agreed-upon guidelines for ethical data analysis and profiling. Such guidelines would address the use of legal and technical mechanisms to obfuscate data; criteria for calling out unethical, if not illegal, behavior; categories of privacy and non-privacy harms; and strategies for empowering individuals through access to data in intelligible form.

Human subject research unfettered

Big data has expanded the scope of human subject research far beyond anything envisaged by social science or medical researchers just a few years ago. Today, everyone – including businesses, governments, private citizens and platform operators – is a “researcher”, analyzing the data exhaust produced by individuals’ daily lives to identify useful patterns and correlations. In most cases, these research activities are not tempered by the procedural and ethical safeguards, which were traditionally required to conduct human subject research. To the contrary, the machine is often driven by entrepreneurs, app developers, or data scientists who seek innovation at any cost. Although in many large companies Chief Privacy Officers and legal teams play an oversight role, today’s start-up app developers can rapidly amass vast amounts of data with little if any oversight. This type of research impacts not only the privacy of individuals whose data is examined, but also the rights of those subject to social sorting as a consequence.¹¹

⁷ Omer Tene & Jules Polonetsky, *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 *NW J. TECH & IP* (forthcoming 2013).

⁸ Cf. Chopra & Whites’ use of “artificial agents”; Samir Chopra & Laurence White, *A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS* (University of Michigan Press 2011).

⁹ Chopra & White, *ibid*, at p. 2.

¹⁰ See discussion *infra*, notes 54 to 60 and accompanying text.

¹¹ For example, if research demonstrates that men between age 40 and 50 who smoked for 10 years have a high instance of heart disease, the insurance premiums charged to an individual who meets these criteria will rise regardless of whether or not his data was in the original dataset.

Like any interpretative process, big data analysis is prone to error and far from objective. Data crunching may appear to be an exact science; yet it is laden with subjective input from researchers who decide which data to analyze, questions to examine, and purposes to pursue. As danah boyd put it: “Do numbers speak for themselves? The answer, we think, is a resounding ‘no’. All researchers are interpreters of data.”¹² The same numbers tell different stories depending on the methodologies and theories of those who set the research agenda. Furthermore, the machine is not immune to error based on inaccurate input, skewed samples, or faulty algorithms. Some commentators go so far as arguing that most current published research findings are false or at best accurate measures of the prevailing bias.¹³ When viewing ads tailored by intricate behavioral tracking infrastructures, we often sense that the machine gets us all wrong. Eli Pariser called this “a bad theory of you”, based on there being no single set of criteria that describes who we are.¹⁴ While relatively benign when the decision is whether to show a web surfer a sports ad or a fashion ad, erroneous results may have profound adverse effects on individuals’ lives in other contexts, such as healthcare, credit, employment or law enforcement.¹⁵

Even more fundamentally, big data analysis is inapposite to traditional methods of scientific research (*i.e.*, define a research question; gather information; form an explanatory hypothesis; test the hypothesis; etc.) While these earlier paradigms were characterized by experimentation and reasoning, big data analysis is driven by the availability of data at an unprecedented scale as well as the computational resources enabling rapid value extraction. As Julie Cohen observes, “the idea of scientific research as a series of limited data collections for the purpose of testing and possibly falsifying a particular hypothesis.”¹⁶ Some regard this challenge to traditional scientific method a groundbreaking revolution, heralding a “fourth paradigm” of scientific research.¹⁷ Others question the rigor of scientific investigations that are both open-ended and ongoing.¹⁸ One commentator notes that “Relaxed practices regarding the communication of computational details is creating a credibility crisis in computational science, not only among scientists, but as a basis for policy decisions and in the public mind.”¹⁹

¹² danah boyd & Kate Crawford, Six Provocations for Big Data, A Decade in Internet Time: Symposium on the Dynamics of the Internet and Society, September 2011.

¹³ John Ioannidis, Why Most Published Research Findings Are False, 2(8) PLOS MEDICINE e124(2005); also see David Freedman, Lies, Damned Lies, and Medical Science, THE ATLANTIC (November 2010).

¹⁴ Eli Pariser, THE FILTER BUBBLE: WHAT THE INTERNET IS HIDING FROM YOU (Penguin, 2011).

¹⁵ See, e.g., Nassim Taleb, Beware the Big Errors of ‘Big Data’, Wired, February 8, 2013, <http://www.wired.com/opinion/2013/02/big-data-means-big-errors-people>.

¹⁶ Julie Cohen, What Privacy is For, 126 HARV. L. REV. ___ (forthcoming 2013).

¹⁷ See, e.g., Gordon Bell, Tony Hey & Alex Szalay, Beyond the Data Deluge, 323 SCIENCE 1297 (2009).

¹⁸ Mark Birkin, Big Data Challenges for Geoinformatics 1(1) GEOINFOR GEOSTAT: AN OVERVIEW (2012), <http://www.scitechnol.com/GIGS/GIGS-1-e101.pdf>.

¹⁹ Stodden ; also see David Berry, The Computational Turn: Thinking About the Digital Humanities, 12 CULTURE MACHINE 1 (2011); David Donoho *et al*, Reproducible Research in Computational Harmonic Analysis, 11(1) COMPUTING IN SCIENCE & ENGINEERING 8 (2009).

Finally, attention must be given to the accessibility of big data sets to the research community at large.²⁰ Traditionally, when scientists published their research, they also made the underlying data available so that other scientists could verify the results. Yet with big data, it is often only the employees of certain organizations that benefit from access, conducting analysis and publishing results without making the underlying data publicly available.²¹ Such scientists may argue, first, that the data are a proprietary asset of their business. Indeed, they may claim that disclosing the data could infringe customers' privacy.²² Who gets access to big data sets; for what purposes; in what contexts; and with what constraints – are fundamental questions that must be addressed by future research.²³

Discrimination – telling right from wrong

Significantly, big data analysis allows for granular distinctions to be made between individual characteristics, preferences and activities. Whether such distinctions are made for the sake of personalization, research or public planning, they facilitate discrimination based on a wide (in fact, infinite) spectrum of characteristics. We refer here to “discrimination” in a value-neutral sense; *i.e.*, drawing distinctions between individuals and treating them differently based on such distinctions.²⁴ To assess the ethical implications of discrimination, we need to unpack the meaning of the term, which is, of course, highly charged. Discrimination could be socially desired (*e.g.*, treating minors as children and not as adults); generally acceptable (*e.g.*, applying Amazon's recommendation system to enhance consumers' shopping experience); or morally reprehensible (*e.g.*, not hiring individuals of a certain age or race). In our daily life, we draw distinctions (*i.e.*, discriminate) all the time. A person sitting next to us on a plane is tall or short, agitated or relaxed, attractive or unattractive, young or old – there is an endless list of such adjectives; and our attitudes and actions towards that person will vary accordingly.

²⁰ See John Markoff, Troves of Personal Data, Forbidden to Researchers, NY TIMES, May 21, 2012, <http://www.nytimes.com/2012/05/22/science/big-data-troves-stay-forbidden-to-social-scientists.html>.

²¹ See Lev Manovich, Trending: The Promises and the Challenges of Big Social Data, in DEBATES IN THE DIGITAL HUMANITIES (Matthew Gold, Editor, The University of Minnesota Press, 2012), http://www.manovich.net/DOCS/Manovich_trending_paper.pdf, claiming that “only social media companies have access to really large social data – especially transactional data. An anthropologist working for Facebook or a sociologist working for Google will have access to data that the rest of the scholarly community will not”.

²² See Bernardo Huberman, Sociology of Science: Big Data Deserve a Bigger Audience, 482 NATURE 308 (2012), warning that privately held data was threatening the very basis of scientific research, and complaining that “[m]any of the emerging 'big data' come from private sources that are inaccessible to other researchers. The data source may be hidden, compounding problems of verification, as well as concerns about the generality of the results”.

²³ boyd & Crawford, *supra* note 17, at p. 12.

²⁴ Miriam-Webster defines the intransitive verb “discriminate” as “1. to make a distinction; 2. to make a difference in treatment or favor on a basis other than individual merit”. Black's Law Dictionary defines “discrimination” as “1. The effect of a law or established practice that confers privileges on a certain class or that denies privileges to a certain class because of race, age, sex, nationality, religion, or disability. 2. Differential treatment; esp., a failure to treat all persons equally when no reasonable distinction can be found between those favored and those not favored”.

The machine can instantly make millions of such distinctions working with vast pool of personal data. But an ethical assessment of machine-driven distinctions requires a coherent theory of discrimination. The machine is incapable of determining whether a distinction is ethical or not. Unless we come up with a comprehensive theory of discrimination that can be represented algorithmically, we have no rigorous way to distinguish between ethical and non-ethical machine-based discrimination.²⁵ We certainly should not expect the machine to make moral decisions that *we* have yet to make.

Have we decided why it is legitimate to market to pregnant women in one context (*e.g.*, based on subscription to a magazine) but morally distasteful to do so in another (*e.g.*, Target's compilation of a "pregnancy score" for shoppers)?²⁶ Can an employer ethically decline to interview a job candidate because they see a picture of them drinking a beer on a social media site?²⁷ Is price discrimination, that is the offering of different prices to different people based on their perceived willingness to pay, good or bad? Does it favor the wealthy²⁸ or the less privileged?²⁹ Is it fair that companies can exploit price sensitivity on an individualized basis, thereby usurping the entire value surplus available in a transaction by pricing goods or services as close as possible to an individual's reservation price?³⁰ What is the fault line between legitimate (or at least not illegal) price discrimination and price discrimination that effectively excludes entire groups of individuals (who are viewed as not being "worth enough" to bother with) from the market? And what if the makeup of such excluded groups is positively correlated with racial or gender bias?

²⁵ There have been attempts of statistical testing for discrimination in big data analysis. *See, e.g.*, Salvatore Ruggieri, Dino Pedreschi & Franco Turini, Data Mining for Discrimination Discovery, 4(2) ACM TRANSACTIONS ON KNOWLEDGE DISCOVERY FROM DATA, Article 9 (May 2010); Binh Thanh Luong, Salvatore Ruggieri & Franco Turini, k-NN as an Implementation of Situation Testing for Discrimination Discovery and Prevention, 17 ACM INTERNATIONAL CONFERENCE ON KNOWLEDGE DISCOVERY AND DATA MINING 502 (August 2011). These efforts too must first coalesce around an agreed upon delineation of legitimate vs. illegitimate discrimination.

²⁶ Charles Duhigg, How Companies Learn Your Secrets, NY TIMES MAGAZINE, February 16, 2012, <http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.

²⁷ Jeffrey Rosen, The Web Means the End of Forgetting, NY TIMES, July 21, 2010, <http://www.nytimes.com/2010/07/25/magazine/25privacy-t2.html>.

²⁸ *See* Jennifer Valentino-Devries, Jeremy Singer-Vine & Ashkan Soltani, Websites Vary Prices, Deals Based on Users' Information, WSJ, December 24, 2012, http://online.wsj.com/article_email/SB1000142412788732377204578189391813881534-1MyQjAxMTAyMDIwNDEyNDQyWj.html#12, reporting that "areas that tended to see the discounted prices had a higher average income than areas that tended to see higher prices"; *also see* Omer Tene, Privacy: For the Rich or for the Poor?, CONCURRING OPINIONS, July 26, 2012, <http://www.concurringopinions.com/archives/2012/07/privacy-for-the-rich-or-for-the-poor.html>.

²⁹ Dana Mattioli, On Orbitz, Mac Users Steered to Pricier Hotels, WSJ, August 23, 2012, <http://online.wsj.com/article/SB10001424052702304458604577488822667325882.html>.

³⁰ In a prior article we likened transacting with a big data platform to a game of poker where one of the players has his hand open and the other keeps his cards close. The online company knows the preferences of the transacting individual inside out, perhaps better than the individual knows him or herself. Tene & Polonetsky, *supra* note 7.

Difficult enough to decide which forms of discrimination are illegal, deciding whether discrimination that is not illegal is unethical or morally undesired may become daunting. Robert Fullinwider explains: “Many may be led to the false sense that they have actually made a moral argument by showing that the practice discriminates (distinguishes in favor of or against). The temptation is to move from ‘X distinguishes in favor of or against’ to ‘X discriminates’ to ‘X is wrong’ without being aware of the equivocation involved”.³¹ Should we preempt any form of discrimination by requiring companies to mail Porsche catalogs to everyone regardless of income? Should Victoria Secret or Pampers be required to target all shoppers regardless of gender or age? Or perhaps offers should always be *available* to all but not *promoted* to all? But then again, that may deny the benefit of the bargain to those who do not know about it.

Some of our ethical and moral criteria are so fragile, nuanced, and culturally dependent that it is not clear that the machine will *ever* be capable of appropriately weighing them. Indeed, it is far from clear that we would even *want* the machine to obtain the ability to distinguish right from wrong. Such an anthropomorphized machine – a “technological singularity”³² – would likely cause more privacy and moral angst than the current dumbed-down version. Artificial intelligence has yet to produce systems that approach human-cognition.³³ Far from it, the only morality that can currently be attributed to the machine is what Ian Kerr calls “slave morality”, the proclivity to fulfill human orders to inhuman perfection. Bruce Boyden recently argued that it is precisely the inhuman nature of the machine that allays privacy concerns in the context of machine-based communications surveillance.³⁴ “What people who worry about privacy are trying to prevent is changed beliefs about themselves, changed behavior by other people, or changed attributions of social status resulting from a disclosure of private information—in other words, changed mental states.” For the sake of privacy, it may be best to leave the tin man without a heart.

Don’t blame the machine

A complicating factor is that the machine’s unrestricted ability to identify patterns in endless piles of data facilitates the masking of illegitimate or illegal discrimination behind layers upon layers of mirrors and proxies.³⁵ A clever

³¹ Robert Fullinwider, THE REVERSE DISCRIMINATION CONTROVERSY 11–12 (1980).

³² Ray Kurzweil, THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY (Penguin 2006).

³³ Note Harry Surden’s observation that “these statistical and probability-based machine-learning models (often combined with logical-knowledge based rules about the world) often produce high-quality and effective results (not quite up to the par of nuanced human translators at this point), without any assertion that the computers are engaging in profound understanding with the underlying “meaning” of the translated sentences or employing processes whose analytical abilities approach human-level cognition”. Harry Surden, Autonomous Agents and Extension of Law: Policymakers Should be Aware of Technical Nuances, CONCURRING OPINIONS, February 16, 2012, <http://www.concurringopinions.com/archives/2012/02/autonomous-agents-and-extension-of-law-policymakers-should-be-aware-of-technical-nuances.html>.

³⁴ Bruce Boyden, Can a Computer Intercept Your Email?, 34 Cardozo L. Rev. 669 (2012).

³⁵ See, e.g., Carter Jernigan & Behram Mistree, Gaydar: Facebook friendships expose sexual orientation, 14(10) FIRST MONDAY (October 2009), <http://www.firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/2611/2302>

programmer can embed bias in a complex algorithm such that discrimination will be very difficult to detect.³⁶ The machine can find strong correlations, which result in discriminatory outcomes that are based on neutral factors. It is wrong to discriminate based on race; yet it will be exceedingly difficult to detect such discrimination if it is based on a dozen factors that through big data analysis have been found to be positively correlated to race. And sometimes it will be difficult to discern whether the category used for profiling is legitimate or just a façade for another, less wholesome agenda.³⁷ This just goes to say that the machine can be a powerful tool for discrimination, just as it is a potent tool for healthcare research, environmental sustainability and economic efficiency. It does not discriminate any more or less legitimately than the people who use it.

There is nothing new about the fact that people discriminate based on unethical criteria, some of which are not illegal. For example, employers may (or may not) prefer to hire attractive job candidates.³⁸ That today they can satisfy such bias by sifting through candidates' Facebook profiles is a phenomenon that has nothing to do with the morality of the technology itself. To be sure, the machine enables the scaling of such discrimination to entire populations. But should we outlaw distinctions drawn by the machine in cases where those same distinctions are legal (albeit subject to moral disdain) if drawn by individuals? Some laws aspire to resolve machine-based discrimination by requiring the involvement of a human operator at certain decisionmaking junctures.³⁹ However, it is far from clear that human intervention mitigates discrimination risk; in fact the opposite may be true. Indeed, when technical risk-based profiles were first introduced in the mortgage industry, they were hailed as a definitive answer to the unequal treatment loan officers give borrowers.⁴⁰

(demonstrating a method for accurately predicting the sexual orientation of Facebook users by analyzing friendship associations).

³⁶ The Supreme Court has ruled that under the Civil Rights Act of 1964, a policy that was fair in form but discriminatory in impact is illegal. *See Griggs v. Duke Power Co.*, 401 U.S. 424, 431, 91 S.Ct. 849, 28 L.Ed.2d 158 (1971) ("Congress has now provided that tests or criteria for employment or promotion may not provide equality of opportunity merely in the sense of the fabled offer of milk to the stork and the fox").

³⁷ Consider the Federal Reserve Board Report asserting that credit card companies adjusted consumers' rates and credit limits based in part on where they shopped, what they bought, and whom they bought from. What could such criteria be correlated to or disguise? *See Board of Governors of the Federal Reserve System, REPORT TO THE CONGRESS ON REDUCTIONS OF CONSUMER CREDIT LIMITS BASED ON CERTAIN INFORMATION AS TO EXPERIENCE OR TRANSACTIONS OF THE CONSUMER*, May 2010, <http://www.federalreserve.gov/BoardDocs/RptCongress/creditcard/2009/consumercreditreductions.pdf>.

³⁸ *Cf.* Don't hate me because I'm beautiful, *THE ECONOMIST*, March 31, 2012, <http://www.economist.com/node/21551535>; *with* Attractiveness discrimination: Hiring hotties, *THE ECONOMIST*, July 21, 2012, <http://www.economist.com/node/21559357>.

³⁹ *See* the European Data Protection Directive's provision on automated individual decisions; Article 15 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), <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:1995:281:0031:0050:EN:PDF>.

⁴⁰ *See* U.S. Department of Justice, Civil Rights Division, Fair Lending Enforcement Program (2001).

Consider, for example, recent research by Latanya Sweeney demonstrating that a greater percentage of ads having the word “arrest” in their text appears for searches on Google and Reuters.com for black identifying first names (such as DeShawn, Darnell and Lakisha) than for white identifying first names (such as Brad, Justin and Jill). Surely it is not the machine that independently decided to discriminate on a first name basis; rather as Sweeney observes, online ad delivery is a “socio-technical construct”. In her research findings, Sweeney could not determine whether the documented bias was caused by advertisers providing ad templates suggestive of arrest disproportionately to black-identifying names, or by the Google Ad Sense algorithm simply reflecting *society’s* bias by preferring to place ads that obtain higher clickthrough rates. Sweeney posits that “technology can do more to thwart discriminatory effects and harmonize with societal norms”. Hence, she calls for “fairness by design” to complement the increasingly prevalent requirement for “privacy by design”.⁴¹ Indeed, Cynthia Dwork and others suggest innovative ways to bake fairness into algorithms to prevent overt or covert discrimination.⁴² At the same time, if we believe certain distinctions are worthy of legal restriction, law should bar their use in decisions regardless of whether they are made by human or machine. In other words, as long as humans continue to be biased and discriminating, machine-made decisions will reflect (and may very well amplify) such discrimination.

Fragmentation of public discourse

An additional and somewhat related problem, which was exposed by Joe Turow,⁴³ Cass Sunstein⁴⁴ and others, concerns the risks to free speech and democratic discourse that are inherent in the fragmentation of the information commons. Personalization technologies channel content into “filter bubbles”, enabling platform providers and inevitably governments to “divide and conquer” by manipulating public opinion.⁴⁵ For example, during the last U.S. Presidential elections, political campaigns were “micro-targeted” delivering individualized messages to potential voters based on their narrow interests, causes and fears.⁴⁶ The data that support this micro-targeting are increasingly being merged with information about the online identities and behavior of voters. These practices raise concerns about loss of voter anonymity, political speech, freedom of association and the transparency of the political process. “This means that campaigns can develop narrow appeals based on ideology and self-interest and

⁴¹ PbD sources.... Ira Rubinstein, *Regulating Privacy by Design*, 26 BERK. TECH. L.J. 1409 (2012)

⁴² Cynthia Dwork et al, *Fairness Through Awareness*, CoRR abs/1104.3913 (2011).

⁴³ Joseph Turow, *THE DAILY YOU: HOW THE NEW ADVERTISING INDUSTRY IS DEFINING YOUR IDENTITY AND YOUR WORTH* (Yale University Press, 2011).

⁴⁴ Cass Sunstein, *REPUBLIC.COM* (Princeton University Press, 2001).

⁴⁵ Pariser, *supra* note 14.

⁴⁶ Daniel Kreiss, *Yes We Can (Profile You): A Brief Primer on Campaigns and Political Data*, 64 STAN. L. REV. ONLINE 70 (2012); Natasha Singer & Charles Duhigg, *Tracking Voters’ Clicks Online to Try to Sway Them*, NY TIMES, October 27, 2012, <http://www.nytimes.com/2012/10/28/us/politics/tracking-clicks-online-to-try-to-sway-voters.html>.

direct them to different groups of voters, appearing to be all things to all people.”⁴⁷

Not only political speech but also artistic and creative freedoms may be affected in a big data environment. For example, recent reports describe how Netflix harvests data from millions of users to produce content that best fits their tastes.⁴⁸ On the one hand, the television market has for many years thrived on a ratings system assessing the popularity of shows based on eyeballs. On the other hand, the ability to amass granular information regarding individuals’ viewing habits of and target specially tailored content at them raises concerns over siloization and narrowcasting. As Joseph Turow puts it, “the industrial logic behind the[se] activities makes clear that the emerging marketplace will be far more an inciter of angst over social difference than a celebration of the ‘American salad bowl’”.⁴⁹

Quite disturbing in this context,⁵⁰ is the fact that the machine is covered by an opaque veil of secrecy, which is backed by corporate claims of trade secrecy and intellectual property. In the analogue world, we could typically understand the logic underlying political advertising, credit or employment decisions; whereas in the big data environment, we are cowed into submission by a powerful data infrastructure, a “surveillant assemblage”,⁵¹ delivering practically uncontested results. This sense of being judged by the tin man, a heartless machine that operates based on incomprehensible criteria, is troubling.⁵² It raises the specter of vulnerability and helplessness that accompanied Franz Kafka’s anti-hero Joseph K, who was confounded by an opaque, logically baffling bureaucracy trying him for an unknown charge.⁵³ And while perhaps tolerable when restricted to the marketing context, such opaque decisionmaking tools threaten to pose a risk to democracy and free speech when introduced into the political sphere.

(Lack of) Regulatory Reform

Against the backdrop of these challenges, policymakers have struggled to come up with a coherent regulatory response. Over the past two years, the OECD, EU and US have launched extensive processes for comprehensive reform of their

⁴⁷ Kreiss, *ibid*; also see Daniel Kreiss & Philip N. Howard, New Challenges to Political Privacy: Lessons from the First U.S. Presidential Race in the Web 2.0 Era, 4 INT’L J. COMM. 1032 (2010).

⁴⁸ Andrew Leonard, How Netflix is turning viewers into puppets, SALON, February 1, 2013, http://www.salon.com/2013/02/01/how_netflix_is_turning_viewers_into_puppets.

⁴⁹ Joseph Turow, NICHE ENVY: MARKETING DISCRIMINATION IN THE DIGITAL AGE (MIT Press 2008), 2.

⁵⁰ When asked if they wanted “political advertising tailored to your interests,” 86% of Americans surveyed said they did not; 64% said their support for a candidate would decrease if they found out a candidate was micro-targeting them differently than their neighbor. Allison Brennan, Microtargeting: How campaigns know you better than you know yourself, CNN, November 5, 2012, <http://edition.cnn.com/2012/11/05/politics/voters-microtargeting/index.html>.

⁵¹ Cohen, *supra* note 16.

⁵² The Wall Street Journal reports that “It is difficult for online shoppers to know why, or even if, they are being offered different deals from other people. Many sites switch prices at lightning speed in response to competitors’ offerings and other factors, a practice known as “dynamic pricing.” Valentino-Devries et al, *supra* note 28.

⁵³ Franz Kafka, THE TRIAL (1925).

privacy frameworks.⁵⁴ Yet the result of these processes remains strongly anchored in the existing policy framework, which is rooted in an architecture dating back to the 1970s.⁵⁵ The major dilemmas and policy choices for informational privacy in the age of big data remain unresolved.

Specifically, privacy and data protection laws are premised on individual control over information and on principles such as data minimization and purpose limitation. Yet it is not clear that minimizing information collection is always a practical approach to privacy in the age of big data. To the contrary, data minimization appears inimical to the very concept of big data. And the discussion over individual control, which is closely linked (through the consent requirement) to principle of purpose limitation, too often transforms into an arena for highly charged polemics between industry and privacy advocates over what the public “really” wants.⁵⁶ The recent legislative reform proposals in Europe, which all but outlaw data based profiling, appear detached from technological and business realities and impossible to operationalize.⁵⁷

⁵⁴ For a description of the reform processes, *see* Omer Tene, Privacy law’s midlife crisis: A critical assessment of the second wave of global privacy laws, ___ OHIO ST. L. J. ___ (forthcoming 2013). *Also see* Terms of Reference for the review of the OECD Guidelines Governing the Protection of Privacy and Transborder Data Flows of Personal Data in OECD, THIRTY YEARS AFTER THE OECD PRIVACY GUIDELINES, 2011, <http://www.oecd.org/sti/interneteconomy/49710223.pdf>; The White House, CONSUMER DATA PRIVACY IN A NETWORKED WORLD: A FRAMEWORK FOR PROTECTING PRIVACY AND PROMOTING INNOVATION IN THE GLOBAL DIGITAL ECONOMY, February 2012, <http://www.whitehouse.gov/sites/default/files/privacy-final.pdf>; Federal Trade Commission Report, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: RECOMMENDATIONS FOR BUSINESSES AND POLICYMAKERS, March 2012, <http://ftc.gov/os/2012/03/120326privacyreport.pdf>; 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), Brussels, January 25, 2012, COM(2012) 11 final, http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf.

⁵⁵ Tene, *ibid*; *also see* Omer Tene, There is no new thing under the sun, CONCURRING OPINIONS, July 30, 2012, <http://www.concurringopinions.com/archives/2012/07/there-is-no-new-thing-under-the-sun.html>.

⁵⁶ Consider the discussion in the W3C Tracking Protection Working Group over the “correct” default setting for the “do not track” mechanism: Natasha Singer, Do Not Track? Advertisers Say ‘Don’t Tread on Us’, NY TIMES, October 13, 2012, <http://www.nytimes.com/2012/10/14/technology/do-not-track-movement-is-drawing-advertisers-fire.html>.

⁵⁷ The Draft EU General Data Protection Regulation, which was submitted by the European Commission in January 2012 as part of the reform of the 1995 Data Protection Directive, would impose strict restrictions on “profiling”. *See* Articles 19-20 of the General Data Protection Regulation, *supra* note 54. These restrictions would be further tightened according to the draft submitted by the European Parliament Rapporteur, which adds to Article 4 of the General Data Protection Regulation a definition of “profiling”: “any form of automated processing of personal data intended to evaluate certain personal aspects relating to a natural person or to analyze or predict in particular that natural person’s performance at work, economic situation, location, health, personal preferences, reliability or behavior”. Draft Report on the proposal for a regulation of the European Parliament and of the Council on the protection of individual with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) (COM(2012)0011 – C7 0025/2012 – 2012/0011(COD)), Committee on Civil Liberties, Justice and Home Affairs, Rapporteur: Jan Philipp Albrecht, http://www.europarl.europa.eu/meetdocs/2009_2014/documents/libe/pr/922/922387/922387en.pdf (noting “a general ban is introduced on profiling as defined in Article 4 and it is only permissible where provided for by law, *i.e.*, either by means of the data subject’s consent or a statutory provision”. *Ibid*, Justification to Amendment 38).

When trying to solve the big data conundrum, it is easy to swing to extremes ranging from techno-utopianism on the one hand⁵⁸ to alarmist fear mongering on the other hand. Alas, technological, business, social and ethical realities will force us to more carefully tread a path towards a nuanced reconciliation of big data benefits with individual rights. Clearly, the principles of privacy and data protection must be balanced against additional societal values such as public health, national security and law enforcement, environmental protection, and economic efficiency. Despite the heated rhetoric,⁵⁹ this remains true regardless of whether privacy is viewed as consumer protection issue, as is often the case in the United States, or as a fundamental human right, as in Europe. Even fundamental rights are seldom absolute and often need to accommodate competing rights and interests.⁶⁰ In this part, we lay out several potentially useful directions for progress, focusing on empowering individuals by enhancing transparency and accountability.

Obscurity – in praise of fuzziness

One promising path is the concept of obscurity, allowing individuals to hide in plain sight. Individuals are far less troubled by data analysis processes that do not *single them out* from a group. Stutzman and Hartzog note that “for an individual to be obscure, an observer must not possess critical information that allows one to make sense of the individual.”⁶¹ In the context of big data, this can be achieved through various means of de-identification, preventing the metaphorical camera lenses from focusing on a particular individual. Indeed, this approach can be viewed as a reconceptualization of Warren and Brandeis’ “right to be let alone”.⁶² One forceful technique is *differential privacy*, which allows researchers to draw lessons and derive valuable conclusions from a data set without being able to determine whether or not such conclusions are based on the personal data of any given individual.⁶³ Hence, differential privacy emphasizes not whether an individual can be directly *associated* with a particular revealed value; but rather the extent to which any revealed value

⁵⁸ Julie Cohen criticizes: “Some of the claims on behalf of Big Data, those framed in terms of a ‘singularity’ waiting in our soon-to-be-realized future, sound quasi-religious, conjuring up the image of throngs of dyed-in-the-wool rationalists awaiting digital rapture.” *Supra* note 16.

⁵⁹ See, e.g., James Fontanella-Khan, Brussels fights US data privacy push, FINANCIAL TIMES, February 10, 2013, <http://www.ft.com/intl/cms/s/0/903b3302-7398-11e2-bcbd-00144feabdc0.html#axzz2KmNKUWab> (noting: “Europe’s most senior justice official is adamant she will fight US attempts to water down a proposed EU data protection and privacy law that would force global technology companies to obey European standards across the world. Viviane Reding, EU commissioner for justice, said that the EU was determined to respond decisively to any attempts by US lobbyists – many working for large tech groups such as Google and Facebook – to curb the EU data protection law”).

⁶⁰ John Morijn, Balancing Fundamental Rights and Common Market Freedoms in Union Law: Schmidberger and Omega in the Light of the European Constitution, 12 EUROP. L. J. 15 (2006).

⁶¹ Woodrow Hartzog & Frederic Stutzman, The Case for Online Obscurity, 101 CAL. L. REV. ___ (forthcoming 2013).

⁶² Samuel Warren & Louis Brandeis, The Right to Privacy, 4 HARV. L. REV. 193 (1890).

⁶³ Cynthia Dwork, Differential Privacy, in 33RD INTERNATIONAL COLLOQUIUM ON AUTOMATA, LANGUAGES AND PROGRAMMING (ICALP 2006), http://www.dbis.informatik.hu-berlin.de/fileadmin/lectures/SS2011/VL_Privacy/Differential_Privacy.pdf.

depends on an individual's data. Another technique is *k-anonymity*, which requires that the data for each person contained in a data release cannot be distinguished from at least $k-1$ individuals whose information also appears in the dataset.⁶⁴ In a previous article, we have argued that there are limits to de-identification in the context of big data.⁶⁵ While we realize that de-identification is not a panacea, we recognize that there is a broad range of situations where it can be a mitigating precaution.

A more proactive approach, referred to by Stutzman and Hartzog as “obscurity by design” would mask personal information behind a veil of obscurity through means such as pseudonymization, restricted access policies and limited searchability.⁶⁶ This would allow information to be shared usefully while at the same time minimizing privacy risks. Similarly, privacy enhancing measures can be integrated into new technologies to minimize their privacy footprint. For example, Snapchat is a mobile application that enables users to share photos or videos that vanish several seconds after being viewed by recipients.⁶⁷ If a recipient manages to capture a screenshot of a flickering photo, the sender is promptly notified by the app. Hence, Snapchat manages to achieve by nimble design what the European legislators seek to impose by regulation, namely a “right to be forgotten” or to erase one’s digital trail.⁶⁸

Access and Transparency

A second promising path entails empowering individuals by granting them access to their personal data in intelligible, machine-readable form. Individuals would thus become active participants in the big data economy, analyzing their own information to improve their health, finances, career prospects, traffic management and more. Through mechanisms such as personal clouds or data stores, individuals could contract with third parties who would get permission to selectively access certain categories of their data to provide analysis, value-added services and mash-ups. We have called this the “featurization” of big data,⁶⁹ making data analysis a consumer-side application and unleashing a wave

⁶⁴ Latanya Sweeney, *k-anonymity: a model for protecting privacy*, 10 (5) INT’L J. UNCERTAINTY, FUZZINESS AND KNOWLEDGE-BASED SYSTEMS 557 (2002).

⁶⁵ Tene & Polonetsky, *supra* note 7; the seminal article advocating de-identification skepticism is Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701 (2010); see discussion in Felix Wu, *Privacy and Utility in Data Sets*, ___ U. COLO. L. REV. ___ (forthcoming 2013).

⁶⁶ Fred Stutzman & Woodrow Hartzog, *Obscurity by Design: An Approach to Building Privacy into Social Media*, CSCW’12, February 2012, http://fredstutzman.com/papers/CSCW2012W_Stutzman.pdf.

⁶⁷ Jenna Wortham, *A Growing App Lets You See It, Then You Don’t*, NY TIMES, February 8, 2013, <http://www.nytimes.com/2013/02/09/technology/snapchat-a-growing-app-lets-you-see-it-then-you-dont.html>.

⁶⁸ See Jeffrey Rosen, *The Right to be Forgotten*, 64 STAN. L. REV. ONLINE 88 (2012); Peter Fleischer, *Foggy thinking about the Right to Oblivion*, PETER FLEISCHER: PRIVACY...? BLOG, March 9, 2011, <http://peterfleischer.blogspot.co.il/2011/03/foggy-thinking-about-right-to-oblivion.html>.

⁶⁹ *Ibid.*

of innovation in the market for personal data applications.⁷⁰ Indeed, the thriving market for mobile apps provides ample proof that user-side installs work in real life.⁷¹ This “sharing the wealth” strategy is justified by both efficiency and fairness concerns. In addition, it will benefit not only individuals but also businesses, which will get access to higher quality data about individuals’ expressed intentions as opposed to guessing such intentions by analyzing online clues.⁷²

A third path involves a different type of transparency -- with respect to businesses’ data analysis processes. Danielle Citron set the stage for the discussion of “technological due process”, pointing-out that “automated systems jeopardize due process norms”.⁷³ It is hard to audit, challenge or amend processes that are concealed in a black box. We propose that businesses be required to reveal the *criteria* used in their decision-making processes, if not the actual algorithms that may be subject to protection of trade secrets and other intellectual property rights.⁷⁴ As Louis Brandeis once wrote, “[s]unlight is said to be the best of disinfectants”.⁷⁵ We trust that if the existence and uses of databases were visible to the public, businesses would more likely avoid unethical or socially unacceptable (albeit legal) uses of data. In certain cases, such as micro-targeting election campaigns, simply shining the light to expose different communications made to specific audiences may provide the necessary check on concerns of inappropriate pandering to constituencies. In other contexts, where the machine makes binding determinations as to individuals’ legal rights, due process requires that the subjects of such decisions are able to challenge them.

Classification of Harms

In order to tailor appropriate responses to big data problems, policymakers need to better define the risk of harm model. The regulatory toolbox to address

⁷⁰ The pioneering work in this field is by Doc Searls, *THE INTENTION ECONOMY: WHEN CUSTOMERS TAKE CHARGE* (2012); and earlier: Rick Levine, Christopher Locke, Doc Searls & David Weinberger, *THE CLUETRAIN MANIFESTO: THE END OF BUSINESS AS USUAL* (Perseus 2000).

⁷¹ iOS v Android: app revenues, downloads and country breakdowns, *THE GUARDIAN*, December 4, 2012, <http://www.guardian.co.uk/technology/appsblog/2012/dec/04/ios-android-revenues-downloads-country>; AppNation & Rubinson Partners, *How Big is the US App Economy? Estimates and Forecasts 2011-2015* (2011), <http://www.slideshare.net/joelrubinson/an3-us-app-economy20112015>.

⁷² Also see Ira Rubinstein, *Big Data: The End of Privacy or a New Beginning*, ___ INT’L DATA PRIV. L. ___ (forthcoming 2013).

⁷³ Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249 (2008).

⁷⁴ See, e.g., Article 12(a) of the European Data Protection Directive, *supra* note 39, requiring organizations to provide an individual with “knowledge of the logic involved in any automatic processing of data concerning him at least in the case of the automated decisions (...)”. Also see Section 1100F of the Dodd-Frank Wall Street Reform and Consumer Protection Act, 2010, Pub. L. 111-203, amending Section 615 of the Fair Credit Reporting Act, 15 U.S.C. 1681m, to require lenders to disclose to borrowers information used to in risk-based pricing decisions, including any numerical credit score used; the range of possible scores; and key factors that adversely affected the borrower’s credit score.

⁷⁵ Louis Brandeis, *What Publicity Can Do*, *HARPER’S WEEKLY*, December 20, 1913, http://c0403731.cdn.cloudfiles.rackspacecloud.com/collection/papers/1910/1913_12_20_What_Publicity_Ca.pdf.

privacy problems (*e.g.*, notice and choice; data retention limitations) does not necessarily answer, and in fact may exacerbate, other harms such as fairness and discrimination.⁷⁶ Given the blurry edges of the concept of privacy, privacy harms are notoriously difficult to categorize.⁷⁷ Yet without such categorization, privacy policy can become muddled with peripheral or even conflicting considerations. For example, as currently framed the European “right to be forgotten” may be viewed as affording protection for one’s reputation rather than privacy; and the right to “data portability” arguably belongs in the sphere of competition – not privacy law.

A harms-based approach to privacy need not be limited to pecuniary or tangible harms. A better understanding of the effect of data analysis on fairness, discrimination, siloization and narrowcasting can expand the scope of privacy harms that are subject to legal protection. Cynthia Dwork and Deirdre Mulligan refer to fairness concerns heavily weighted by issues of discrimination, including price discrimination based on location (redlining) or on knowledge of the consumer’s state of mind.⁷⁸ Jules Polonetsky and I point out that processing of personal data increasingly affects fairness, equality and other values, which are no less important than – even if theoretically distinct from core privacy interests.⁷⁹

Putting Data in Context

A final response involves the concept of context, which is based on Helen Nissenbaum’s “contextual integrity” analysis of privacy.⁸⁰ Privacy, according to Nissenbaum, is “a function of several variables, including the nature of the situation or context; the nature of information in relation to that context; the roles of agents receiving information, their relationships to information subjects; on what terms the information is shared by the subject and the terms of further dissemination.”⁸¹ This approach may require, for example, that certain categories of sensitive data (*e.g.*, genetic data) be segregated from the decisionmaking process in certain contexts (*e.g.*, employment applications). Where to draw the contextual line becomes a weighty policy question where considerations of national security or public health are involved. In these cases, involving, for example, harvesting of social networking information to detect

⁷⁶ For example, in order to comply with rules on affirmative action, certain organizations are compelled to collect and retain information about individuals’ gender or race. In these cases, data deletion, while privacy protective, would be counter-productive.

⁷⁷ For notable exceptions *see* M. Ryan Calo, *The Boundaries of Privacy Harm*, 86 *IND. L. J.* 1131 (2011); Daniel Solove, *A Taxonomy of Privacy*, 154 *U. PA. L. REV.* 477 (2006).

⁷⁸ Cynthia Dwork & Deirdre Mulligan, *Aligning Classification Systems with Social Values through Design* (manuscript on file with authors).

⁷⁹ Tene & Polonetsky, *supra* note 7.

⁸⁰ Helen Nissenbaum, *PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE* (Stanford Law Books 2009).

⁸¹ Helen Nissenbaum, *Privacy as Contextual Integrity*, 79 *WASH. L. REV.* 119 (2004). The Federal Trade Commission adopted this approach in its *REPORT ON PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: RECOMMENDATIONS FOR BUSINESSES AND POLICYMAKERS*, March 2012, <http://ftc.gov/os/2012/03/120326privacyreport.pdf>.

potential terrorist threats⁸² or analyzing search engine logs to analyze harmful drug interactions,⁸³ individuals' privacy interests may be outweighed by public policy concerns.

Moreover, as Nissenbaum recognizes, relationships and therefore context can change over time. Some argue that soliciting express consent should be a prerequisite to *any* shift in existing boundaries. In reality, however, shifting contexts are not always readily negotiated. Rather, organizations should assess the effects of any prospective change on data subject expectations; convey their policies clearly and conspicuously; and in certain cases provide data subjects with an opportunity to opt out. When a change in context is radical and transparency measures inadequate to support it, express consent can be relied upon to ensure that data subjects are willing to accept a new data use.⁸⁴

Conclusion

As we recognize the immense benefits of big data, we should avoid technological determinism that allows the machine to surge forward with disregard for evolving social norms. Instead of asking "what technology wants",⁸⁵ we should explore what it is that *we want* to achieve with technology and what price we are or are not willing to pay in privacy, social cohesion and individual rights. The lack of agreement in the effort to standardize a "Do Not Track" protocol demonstrates the challenge in seeking a technological solution when the value of the activity to be proscribed remains widely disputed.⁸⁶ Hence, we must first address the ethics and morality of the decisions that confront us. Practically, we need to devise agreed-upon guidelines for ethical data analysis and profiling, addressing such issues as obscurity by design; empowerment through useful access; transparency of decisional criteria; and categorization of potential harms. Technology innovators and data scientists will lead the way to new big data frontiers, but it is philosophers seeking "a new digital humanism"⁸⁷ who must closely follow in their footsteps.

⁸² Ryan Gallagher, Software that tracks people on social media created by defense firm, THE GUARDIAN, February 10, 2013, <http://www.guardian.co.uk/world/2013/feb/10/software-tracks-social-media-defence>.

⁸³ Nicholas Tatonetti, Guy Haskin Fernald & Russ Altman, A Novel Signal Detection Algorithm for Identifying Hidden Drug-Drug Interactions in Adverse Event Reports, J. AM. MED. INFORM. ASSOC. (2011).

⁸⁴ Jules Polonetsky & Omer Tene, IT'S NOT HOW MUCH DATA YOU HAVE, BUT HOW YOU USE IT: ASSESSING PRIVACY IN THE CONTEXT OF CONSUMER DATA INTEGRATION, Future of Privacy Forum White Paper, December 2012, [http://www.scribd.com/doc/115516310/It-s-Not-How-Much-Data-You-Have-But-How-You-Use-It-Assessing-Privacy-in-the-Context-of-Consumer-Data-Integration;_also-see-Paul-Ohm,Branding-Privacy,97-MINN.L.REV.____\(forthcoming-2013\)](http://www.scribd.com/doc/115516310/It-s-Not-How-Much-Data-You-Have-But-How-You-Use-It-Assessing-Privacy-in-the-Context-of-Consumer-Data-Integration;_also-see-Paul-Ohm,Branding-Privacy,97-MINN.L.REV.____(forthcoming-2013)).

⁸⁵ Kevin Kelly, WHAT TECHNOLOGY WANTS (Penguin 2010).

⁸⁶ Omer Tene & Jules Polonetsky, To Track or 'Do Not Track': Advancing Transparency and Individual Control in Online Behavioral Advertising, 13 MINN. J. L. SCI. & TECH. 281 (2012).

⁸⁷ Jaron Lanier, YOU ARE NOT A GADGET (Vintage 2010).