1

Consumer Privacy and the Future of Society

Jules Polonetsky, Omer Tene, and Evan Selinger

In the course of a single day, hundreds of companies collect massive amounts of information from individuals. Sometimes they obtain meaningful consent. Often, they use less than transparent means. By surfing the web, using a cell phone and apps, entering a store that provides Wi-Fi, driving a car, passing cameras on public streets, wearing a fitness device, watching a show on a smart TV or ordering a product from a connected home device, people share a steady stream of information with layers upon layers of hardware devices, software applications, and service providers. Almost every human activity, whether it is attending school or a workplace, seeking healthcare or shopping in a mall, driving on a highway or watching TV in the living room, leaves behind data trails that build up incrementally to create a virtual record of our daily lives. How companies, governments, and experts should use this data is among the most pressing global public policy concerns.

Privacy issues, which are at the heart of many of the debates over data collection, analysis, and distribution, range extensively in both theory and practice. In some cases, conversations about privacy policy focus on marketing issues and the minutiae of a website's privacy notices or an app's settings. In other cases, the battle cry for privacy extends to diverse endeavors, such as the following: calls to impose accountability on the NSA's counterterrorism mission;[1] proposals for designing safe smart toys;[2] plans for enabling individuals to scrub or modify digital records of their pasts;[3] pleas to require database holders to inject noise into researchers' queries to protect against leaks that disclose an individuals' identity;[4] plans to use crypto currencies[5] or to prevent criminals and terrorists from abusing encryption tools;[6] proposals for advancing medical research

[1] Richard Clarke, Michael Morell, Geoffrey Stone, Cass Sunstein & Peter Swire, The NSA Report: Liberty and Security in a Changing World (The President's Review Group on Intelligence and Communications Technologies, Princeton University Press, 2014).

[2] *Kids and the Connected Home: Privacy in the Age of Connected Dolls, Talking Dinosaurs, and Battling Robots* (Future of Privacy Forum and Family Online Safety Institute, Dec. 2016), https://fpf.org/wp-content/uploads/2016/11/Kids-The-Connected-Home-Privacy-in-the-Age-of-Connected-Dolls-Talking-Dinosaurs-and-Battling-Robots.pdf.

[3] Case C-131/12 Google Spain v. Agencia Española de Protección de Datos (AEPD) and Mario Costeja González, ECLI: EU:C:2014:317.

[4] Cynthia Dwork, Frank McSherry, Kobbi Nissim & Adam Smith, *Calibrating Noise to Sensitivity in Private Data Analysis*, in Proceedings of the 3rd Theory of Cryptography Conference, 265–284 (2006).

[5] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller & Steven Goldfeder, Bitcoin and Cryptocurrency Technologies (Princeton University Press, 2016).

[6] In re Order Requiring Apple, Inc. to Assist in the Execution of a Search Warrant Issued by This Court, No. 15-mc-1902 (JO) (E.D.N.Y. Feb. 29, 2016).

and improving public health without sacrificing patients' control over their data;[7] and ideas for how scientists can make their data more publicly available to facilitate replication of studies without, at the same time, inadvertently subjecting entire populations to prejudicial treatment, including discrimination.[8]

At a time when fake news influences political elections, new and contentious forms of machine-to-machine communications are emerging, algorithmic decision-making is calling more of the shots in civic, corporate, and private affairs, and ruinous data breaches and ransomware attacks endanger everything from financial stability to patient care in hospitals, "privacy" has become a potent shorthand. Privacy is a boundary, a limiting principle, and a litmus test for identifying and adjudicating the delicate balance between the tremendous benefits and dizzying assortment of risks that insight-filled data offers.

DIVERSE PRIVACY PERSPECTIVES

The wide scope of perspectives found in this collection reflects the very diversity of privacy discourse.

Since privacy is front-page news, politicians regularly weigh in on it. Some politicians make privacy their signature issue by submitting legislative proposals, convening committee hearings, and sending letters to technology companies as they launch and test new tools. Interestingly, in the United States, privacy can be a bipartisan issue that brings together coalitions from opposite sides of the aisle. For example, on questions of national security surveillance, right wing libertarians side with left wing civil rights activists in opposing government powers and advocating for robust oversight mechanisms. However, in the consumer privacy space, traditional roles are often on display as supporters of regulation spar with free market activists on issues ranging from telecom regulation to the legitimacy of the data broker industry. In Europe, left wing parties, such as the German Greens or the Scandinavian Pirate Party, have played important roles in privacy advocacy by embracing an expansive reading of data protection principles. Conservatives, by contrast, have sought to balance data protection against economic interests and free trade. This political tension manifests itself in the twin, often conflicting objectives of the European data protection regime, which instructs Member States to "protect the fundamental rights and freedoms of natural persons, and in particular their right to privacy with respect to the processing of personal data," while, at the same time, "neither restrict[ing] nor prohibit[ing] the free flow of personal data between Member States."

Industry interest in privacy often aligns with businesses uniformly vying for more data use and less regulation. Even so, opinions still splinter across a broad spectrum. Some publishers believe that stronger limits on ad-tracking will advantage them to collect ad revenue that is earned today by advertising technology companies or large platforms. Other companies believe that new data portability rules will enable them to leverage data now held by platforms to better compete or to launch new services. Nevertheless, incumbents in many sectors worry that new regulations and more extensive liability will impede their digital strategies.

[7] Salil Vadhan, David Abrams, Micah Altman, Cynthia Dwork, Paul Kominers, Scott Duke Kominers, Harry Lewis, Tal Moran & Guy Rothblum, Comments on Advance Notice of Proposed Rulemaking: Human Subjects Research Protections: Enhancing Protections for Research Subjects and Reducing Burden, Delay, and Ambiguity for Investigators, Docket ID No. HHS-OPHS-2011–0005 (2011), https://privacytools.seas.harvard.edu/publications/comments-advance-notice-proposed-rulemaking-human-subjects-research.

[8] Daniel Goroff, Jules Polonetsky & Omer Tene, *Privacy Protective Research: Facilitating Ethically Responsible Access to Administrative Data*, 65 Ann. Am. Acad. Pol. & Soc. Sci. 46–66 (2018).

Regulators chase the flurry of market developments with carrots and sticks. Approaches vary, with some regulators, such as the UK Information Commissioner's Office, offering advice, best practices, and compliance tools. Others, such as the Canadian Federal Privacy Commissioner, enhance limited enforcement powers by actively engaging with the media to "name and shame" alleged violations of privacy laws. Some European data protection regulators are known to levy stiff fines and penalties even for technical violations of local statutes. The compliance risks for businesses will escalate sharply with the imposition of formidable sanctions under the General Data Protection Regulation. The Federal Trade Commission (FTC), the main federal privacy regulator in the United States, has developed a complex privacy and security regulatory approach that is built on two pillars. On the one hand, it includes a string of settlements referred to by Daniel Solove and Woodrow Hartzog as a "common law" of privacy.[9] On the other hand, the FTC issues a line of policy guidelines through workshops and reports on cutting-edge issues ranging from connected vehicles and consumer genetics to the sharing economy.

Privacy academics are a heterogeneous group who occupy a central place in policy debates. Some are data optimists. They see a bright future in data-intensive technologies and seek to facilitate their adoption while respecting individuals' rights. Others are data pessimists. They warn against the disruptive risk of data technologies and in extreme cases even see an inevitable decline toward a "database of ruin."[10] More traditionally, academics can be loosely categorized according to their disciplines. Law and policy scholars explore issues such as the Fourth Amendment, privacy legislation such as the Health Insurance Portability Act, the Family Educational Rights and Privacy Act, the Fair Credit Reporting Act, the Children's Online Privacy Protection Act, and the FTC's body of privacy law. Computer scientists deal with issues such as security and privacy in online, mobile operating systems and software, network security, anonymity, human–machine interaction, and differential privacy. Engineers work on network security, values in design, privacy by design, blockchain, and privacy-enhancing technologies. Economists assess the value and markets for data, as well as such issues as the value of privacy, privacy incentives and nudges, data-based price discrimination, privacy in credit and health markets, the behavioral economics of privacy, and more. Design schools innovate privacy messaging, information schools explore the role of privacy in media and culture, psychologists experiment on individuals' responses to incentives in cyber and real-world spaces, and ethicists weigh in on all of this.

CONSUMER PRIVACY

This book brings together academics, policy makers, and industry leaders to critically address the subset of issues that are raised in the context of *consumer privacy*. It purposefully sets aside the fateful dilemmas raised by government surveillance. This includes the continuing fallout from Edward Snowden's revelations about the prevalence of government access to private communications data. And it extends to newly emerging challenges, such as deploying military drones to assassinate suspected terrorists, using data-driven software for criminal sentencing, and monitoring people awaiting trial and serving court-mandated sentences in the seclusion of their homes. Yet, even narrowed to consumer privacy, this book still addresses a rich spectrum of issues triggered by an exceedingly broad swath of activities. While consumer privacy once was limited

[9] Daniel Solove & Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUM. L. REV. 583 (2014).
[10] Paul Ohm, *Don't Build a Database of Ruin*, HARV. BUS. REV, Aug. 23, 2012, https://hbr.org/2012/08/dont-build-a-database-of-ruin.

to the realm of online tracking for targeted advertising,[11] the topic now extends to wearable technologies and implantable medical devices, smart homes and autonomous vehicles, facial recognition and behavioral biometrics, and algorithmic decision-making and the Internet of Things.[12] As companies collect massive amounts of data through the Internet, mobile communications, and a vast infrastructure of devices and sensors embedded in healthcare facilities, retail outlets, public transportation, social networks, workplaces, and homes, they use the information to test new products and services, improve existing offerings, and conduct research.

Given the wide scale and scope of consumer privacy, the topic can't be easily distinguished from government surveillance. With companies amassing huge warehouses of personal information, governments can swoop in when necessary to access the data through procurement, legal process, or technological capabilities. As Chris Hoofnagle observed more than a decade ago, "Accumulations of information about individuals tend to enhance authority by making it easier for authority to reach individuals directly. Thus, growth in society's record-keeping capability poses the risk that existing power balances will be upset."[13]

Since each new space and field of activity raises weighty policy, legal, ethical, economic, and technological questions and challenges, input on privacy is needed from experts across the disciplines. Philosophers, social scientists, legal theorists, geneticists, mathematicians, computer scientists, and engineers all have important roles to play. The pressing debates require a careful balancing of diverse values, interests, rights, and considerations. In many cases, individual benefits are pitted against the public good, and this tension tests the contours of autonomy and fundamental human rights in a constantly shifting techno-social environment.

The impact of technology on the economy and global markets cannot be overstated. Several of the most highly valued companies are data-driven innovators. That is why companies such as Apple, Google, Microsoft, Amazon, and Facebook, alongside traditional technology powerhouses, such as Intel, IBM and AT&T, and new upstarts, including Uber and Snap, are the focus of heated consumer discussion and regulatory debate.[14] This trend goes beyond the United States and, more broadly, the Western world. Chinese tech giants, such as Baidu, Alibaba, JD.com, and surging new entrants – notably, Didi Chuxing, and Lu.com – are shaking up the Asian economy and gaining a global footprint.[15] These companies have profound impacts our lives. Every day, they confront a host of complex value-laden choices when designing products that collect, analyze, process, and store information about every aspect of our behavior. Realizing the magnitude of these decisions, companies have begun to create ethical review processes, employ data ethicists and philosophers, and seek guidance from academics, think tanks, policymakers, and regulators.[16] The role of the chief privacy officer, once the domain of only a handful of

[11] 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).

[12] Woodrow Hartzog & Evan Selinger, *The Internet of Heirlooms and Disposable Things*, 17 N. C. J. L. & Tech. 581 (2016).

[13] Chris Jay Hoofnagle, *Big Brother's Little Helpers: How ChoicePoint and Other Commercial Data Brokers Collect, Process, and Package Your Data for Law Enforcement*, 29 N. C. J. Int'l L. & Com. Reg. 595 (2004).

[14] Farhad Manjoo, Tech's *"Frightful 5" Will Dominate Digital Life for Foreseeable Future*, N.Y. Times, Jan. 20, 2016, https://www.nytimes.com/2016/01/21/technology/techs-frightful-5-will-dominate-digital-life-for-foreseeable-future.html.

[15] Brendon Kochkodin, *Chinese Big Five Tech Companies Gain on U.S. Counterparts*, Bloomberg Businessweek, June 22, 2017, https://www.bloomberg.com/news/articles/2017-06-23/chinese-big-five-tech-companies-gain-on-u-s-counterparts.

[16] Jules Polonetsky, Omer Tene & Joseph Jerome, *Beyond the Common Rule: Ethical Structures for Data Research in Non-Academic Settings*, 13 Colo. Tech. L. J. 333 (2015); also see Ryan Calo, *Consumer Subject Review Boards: A Thought Experiment*, 66 Stan. L. Rev. Online 97, 102 (2013); Evan Selinger & Woodrow Hartzog, *Facebook's*

technology leaders, has emerged as a strategic C-suite position.[17] Within a decade, privacy has matured into a full-fledged profession with a body of knowledge, professional certifications, and formal legal status.[18]

Increasingly, not only companies but also government entities are transforming into data service providers for consumers. Consider smart cities, where local governments have become hubs of data that is collected through growing networks of sensors and connected technologies to generate actionable, often real-time information.[19] By relying on ubiquitous telecommunications technologies to provide connectivity to sensor networks and set actuation devices into operation, smart cities are increasingly collecting information on cities' air quality, temperature, noise, street and pedestrian traffic, parking capacity, distribution of government services, emergency situations, and crowd sentiments, among other data points. This information can now be cheaply aggregated, stored, and analyzed to draw conclusions about the intimate affairs of city dwellers. The more connected a city becomes, the more it will generate steady streams of data from and about its citizens and the environment they live in.[20]

The urban data revolution enables cities to better manage traffic congestion, improve energy efficiency, expand connectivity, reduce crime, and regulate utility flow. By analyzing data trends and auditing the performance of schools, public transportation, waste management, social services, and law enforcement, smart cities can better identify and respond to discriminatory practices and biased decision-making, empowering weakened populations and holding institutions to account. At the same time, the specter of constant monitoring threatens to upset the balance of power between city governments and city residents. At the extreme, it might destroy the sense of anonymity that has defined urban life over the past century. As Kelsey Finch and Omer Tene observe, "There is a real risk that, rather than standing as 'paragons of democracy,' [smart cities] could turn into electronic panopticons in which everybody is constantly watched."[21]

Smart community policy also highlights the tension between the push for open data mandates and public records acts and the desire citizens have for privacy. On the one hand, the transparency goals of the open data movement serve important social, economic, and democratic functions. Open and accessible public data can benefit individuals, companies, communities, and government by fueling new social, economic, and civic innovations, and improving government accountability and transparency. On the other hand, because the city collects and shares information about its citizens, public backlash over intrusive surveillance remains an ever-present possibility.[22] Due to these competing concerns, the consumer privacy discussion requires aligning potentially conflicting interests: maximizing transparency and accountability without forsaking individual rights.

Emotional Contagion Study and the Ethical Problem of Co-Opted Identity in Mediated Environments Where Users Lack Control, 12 Research Ethics 35 (2016).

[17] Andrew Clearwater & J. Trevor Hughes, *In the Beginning . . . An Early History of the Privacy Profession*, 74 Ohio St. L. J. 897 (2013).

[18] J. Trevor Hughes & Cobun Keegan, *Enter the Professionals: Organizational Privacy in a Digital Age* (see Chapter 22).

[19] Kelsey Finch & Omer Tene, *Welcome to Metropticon: Protecting Privacy in a Hyperconnected Town*, 41 Fordham Urban L. J. 1581 (2015).

[20] Kelsey Finch & Omer Tene, *The City as a Platform: Enhancing Privacy and Transparency in Smart Communities* (see Chapter 7).

[21] Finch & Tene, *supra* note 16, at 1583.

[22] Ben Green, Gabe Cunningham, Ariel Ekblaw, Paul Kominers, Andrew Linzer & Susan Crawford, *Open Data Privacy: A Risk-Benefit, Process-Oriented Approach to Sharing and Protecting Municipal Data* (Berkman Klein Center for Internet & Society Research Publication, 2017), https://dash.harvard.edu/bitstream/handle/1/30340010/OpenData Privacy.pdf.

BEYOND PRIVACY

As we have been suggesting, arguments about privacy have become proxy debates for broader societal choices about fairness, equity, and power. Since data is central to economic activity across every sector – government, non-profit, and corporate – the privacy debate has spilled over to adjacent areas. Educational technology is a prime example.

Long confined to using textbooks, blackboards, and pencil-and-paper testing, schools now use new applications, hardware, and services. This includes online curricula and tools, social media and cloud applications for file sharing and storage, note taking, and collaboration platforms, and a variety of connected tablets and workstations. Student performance data is driving next-generation models of learning and measurements for teacher effectiveness. And connected learning is fast becoming a path for access to knowledge and academic achievement.

New educational technology offers many advantages for educators, teachers, parents, and students. Education has become more interactive, adaptive, responsive, and even fun. Parents can stay apprised of their child's performance, accomplishments, and difficulties without weighing down teachers' limited time resource. Teachers can connect to sophisticated learning management systems, while school administrations can obtain rich, measurable inputs to better calibrate resources to needs.[23]

However, from a privacy perspective, the confluence of enhanced data collection that contains highly sensitive information about children and teens also makes for a combustive mix. New data flows raise questions about who should have access to students' data and what are the legitimate uses of the information. Should a developer of a math app be authorized to offer high-performing students a version that covers more advanced material, or would that be considered undesirable marketing to children? Should an educational social network be permitted to feature a third-party app store for kids? Or, if an education service detects a security vulnerability on a website that is available for schools to use, should it be able to leverage its knowledge to protect schools as well as clients outside of the educational sector? And what about education technology developers who want to use the data they extract from students to develop software for the general market?

It is clear that when it comes to education, privacy means different things to different people and traditional privacy problems are only the tip of the policy iceberg. Activists have challenged data collection and use to debate school reform, common core curricula, standardized testing, personalized learning, teacher assessments, and more. Some critics even consider efforts to ramp up education technology misguided altogether, labeling them as the work of "corporate education reformers" who seek profit at the expense of public education. Ultimately, then, the challenge for educational technology entails differentiating problems that can be remedied with privacy solutions from problems that require other resolutions because they are, at bottom, proxies for conflicts about education policy.

Complex conversations also surround smart cars and autonomous vehicles. On the one hand, collecting data in cars is old hat. Vehicles have had computerized data systems since the 1960s. On the other hand, things are profoundly changing now that vehicles are becoming data hubs that collect, process, and broadcast information about drivers' performance, geolocation, telematics, biometrics, and even media consumption. Furthermore, vehicle-to-vehicle (V2V)

[23] Jules Polonetsky & Omer Tene, *Who is Reading Whom Now: Privacy in Education from Books to MOOCs*, 17 VAND. J. ENT. & TECH. L. 927 (2015); *also see* Jules Polonetsky & Omer Tene, *The Ethics of Student Privacy: Building Trust for Ed Tech*, 21 INT'L REV. INFO. ETHICS 25 (2014).

technology introduces a new way for smart cars to seamlessly receive and analyze information about other vehicles. This capability is essentially transforming public thoroughfares into a seamless network of information about each vehicle's position, direction of travel, speed, braking, and other variables that telematics studies.[24]

Smart car data collection raises all kinds of issues. Consumers and advocates are concerned about cars extracting personal data that can be shared with government and law enforcement. Security experts are anxious about self-driving cars being vulnerable to hacking. At the same time, under the banner of privacy concerns, critics also discuss ethics, labor markets, insurance premiums, and tradeoffs between safety and autonomy. For example, while smart cars and autonomous vehicles can reduce traffic accidents, they will also need to make decisions with moral implications, such as choosing to prioritize the safety of passengers or pedestrians. Coding algorithms to make momentous moral choices is a formidable challenge that transcends the guidance traditional privacy frameworks offer.

Insurance companies are vigorously embracing the growth in vehicle-generated data by developing usage-based applications to harness information emanating from onboard diagnostic systems. These applications provide insurers with information on how a vehicle is driven, and they factor in this information when making decisions about safe driver programs and personalized insurance rates. While the Fair Credit Reporting Act applies to the process of using data to make insurance decisions, its standards cannot address all of the questions that are starting to arise. Concern is being expressed over allocations of risk and the process of creating categories of drivers who are uninsurable due to traits and tendencies that potentially can be correlated with health, genetics, race, and ethnicity. Also, within a generation, autonomous vehicles will fundamentally upend labor markets. Ostensibly consumers will benefit from increased fleet efficiency and huge savings in labor costs. At the same time, the economic changes seem poised to dramatically affect employment prospects, especially for the millions of taxi and truck drivers in the United States and beyond.[25] These policy issues clearly extend digital and cyber privacy debates into new realms and possibly transform them as well.

THE FUTURE OF SOCIETY

The upshot of the dynamics and processes highlighted here is that the chapters in this book are about much more than consumer privacy – which is to say, they go far beyond consumer privacy construed as a niche topic. Contributors fundamentally advance conversations about what paths should be paved in order to create flourishing societies in the future. With every aspect of human behavior being observed, logged, analyzed, categorized, and stored, technology is forcing legislatures, regulators, and courts to deal with an incessant flow of weighty policy choices. These debates have long spilled over from the contours of privacy, narrowly defined as a right to anonymity, seclusion and intimacy – a right to be let alone[26] – to a discussion about power and democracy, social organization, and the role humans should occupy in technologically mediated spaces. These tough discussions are about matters such as exposure, profiling and discrimination, self-expression, individual autonomy, and the relative roles of humans and machines.

[24] Lauren Smith & John Verdi, *Comments from the Future of Privacy Forum to the Federal Trade Commission and U.S. Department of Transportation* (National Highway Traffic Safety Administration, May 1, 2017), https://fpf.org/wp-content/uploads/2017/05/Future-of-Privacy-Forum-Comments-FTC-NHTSA-Workshop.pdf.

[25] *See, e.g.*, *The Future of Jobs: Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution* (World Economic Forum, Jan. 2016), http://www3.weforum.org/docs/WEF_Future_of_Jobs.pdf.

[26] Samuel Warren and & Louis Brandeis, *The Right to Privacy*, 4 Harv. L. Rev. 193 (1890).

Consider what happened when a teacher was fired after a picture was posted on Facebook of her dressed as a drunk pirate. It was hard to know if the ensuing public debate was about privacy settings on the social network or the limits of assessing behavior in a world where every action is documented, tagged, and presented to the public to judge.[27] Similarly, it is hard to pinpoint what parents and teachers are concerned about when they recoil against ephemeral cyberbullying messages on apps such as Snapchat. Is it dismay about the software's privacy settings? Or might it be sadness over the cruel experiences of childhood being exposed and augmented through a new medium?[28] And what about autonomous vehicles engineers who design a real-life response to the longstanding trolley problem? Are they dealing with fair information practice principles or ethical challenges that have occupied philosophers from Aristotle to Immanuel Kant and John Stuart Mill?[29]

Advances in artificial intelligence and machine learning keep raising the stakes. Developers deploy artificial intelligence to improve organizations' performance and derive predictions in almost every area of the economy. This happens in domains ranging from social networks, autonomous vehicles, drones, precision medicine, and the criminal justice system. And it includes such processes as speech and image recognition, universal translators, and ad targeting, to name a few. Organizations leverage algorithms to make data-based determinations that impact individuals' rights as citizens, employees, seekers of credit or insurance, and so much more. For example, employers use algorithms to assess prospective employees by offering neuroscience-based games that are said to measure inherent traits. Even judges turn to algorithms for sentencing and parole decisions. They use data to predict a person's risk of recidivism, violence, or failure to appear in court based on a complicated mix of behavioral and demographic characteristics.[30]

Daniele Citron has written about the importance of creating appropriate standards of algorithmic due process that include transparency, a right to correct inaccurate information, and a right to appeal adverse decisions.[31] Unfortunately, this goal might be incredibly difficult to meet. Thanks to machine learning, sophisticated algorithmic decision-making processes arguably have become inscrutable, even to their programmers. The emergent gap between what humans and machines know has led some critics, such as Frank Pasquale, to warn against the risks of a *Black Box Society*[32] driven by what Cathy O'Neil dubs *Weapons of Math Destruction*.[33]

At the same time, breakthroughs in artificial intelligence have enabled disenfranchised groups to speak the truth to power by identifying biases and inequities that were previously hidden in

[27] Jeffrey Rosen, *The Web Means the End of Forgetting*, N.Y. Times, July 21, 2010, http://www.nytimes.com/2010/07/25/magazine/25privacy-t2.html.

[28] J. Mitchell Vaterlaus, Kathryn Barnett, Cesia Roche and & Jimmy Young, "*Snapchat is more personal*": *An Exploratory Study on Snapchat Behaviors and Young Adult Interpersonal Relationships*, 62 Computers Hum. Behav. 594 (2016); *also see* Evan Selinger, Brenda Leong & Bill Fitzgerald, *Schools Fail to Recognize Privacy Consequences of Social Media*, Christian Sci. Monitor, Jan. 20, 2016, https://www.csmonitor.com/World/Passcode/Passcode-Voices/2016/0120/Opinion-Schools-fail-to-recognize-privacy-consequences-of-social-media.

[29] *Why Self-Driving Cars Must Be Programmed to Kill*, MIT Tech. Rev., Oct. 22, 2015, https://www.technologyreview.com/s/542626/why-self-driving-cars-must-be-programmed-to-kill/.

[30] Omer Tene & Jules Polonetsky, *Taming the Golem: Challenges of Ethical Algorithmic Decision Making*, 19 N. C. J. L. & Tech. (forthcoming 2019).

[31] Danielle Keats Citron, *Technological Due Process*, 85 Wash. U. L. Rev. 1249 (2008).

[32] Frank Pasquale, The Black Box Society (Harvard University Press, 2015).

[33] Cathy O'Neil, Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy (Crown, 2016).

opaque databases or behind faceless human bureaucracies.[34] New uses of data can also save lives. For example, the United Nations Global Pulse project uses data from cell phones in developing countries to detect pandemics, relieve famine, and fight human trafficking.[35] New policy initiatives have been started that recognize the mixed blessings of artificial intelligence and the inevitable trade-offs created by using it. The Partnership on Artificial Intelligence, for example, is set "to address such areas as fairness and inclusivity, explanation and transparency, security and privacy, values and ethics, collaboration between people and artificial intelligence systems, interoperability of systems, and the trustworthiness, reliability, containment, safety, and robustness of the technology."[36]

In an ideal world, due process would be secured and every company would follow all of the technical privacy rules all of the time. But even in such a utopia, consumers and commentators probably still would be unsettled by "creepy" technologically mediated behavior. Attributions of "creepy" revolve around activities where people believe that harm is occurring even though privacy settings are not circumvented and data use technically remains within the scope of its intended purposes. These are instances where new technology further erodes cherished values, such as obscurity, or new uses of existing technologies produce novel outcomes, such as unexpected data use or customization.[37] People are rattled by such threats to traditional social norms and the prospect that unsettling new practices will be normalized. In these moments, they wonder why engineers and marketers fail to anticipate problems. Sometimes, they hold these groups accountable.

All of this suggests that, far from what a first glance at the title of this volume might lead readers to expect, the *Cambridge Handbook of Consumer Privacy* critically explores core issues that will determine how the future is shaped. To do justice to the magnitude and complexity of these topics, we have asked contributors to address as many parts and perspectives of the consumer privacy debate as possible. How we, all of us, collectively grapple with these issues will determine the fate of technology and course of humanity.[38]

CHAPTER SUMMARIES[39]

The Pervasiveness and Value of Tracking Technologies

In Chapter 2, "Data Brokers – Should They Be Reviled or Revered," Jennifer Barrett Glasgow defines the various types of data brokers as they exist today in the United States. She discusses where they get their data and how much of it is aggregated from multiple sources. Glasgow also describes how data brokers deliver data to the marketplace and who buys data from a data broker. She covers how data brokers are regulated by law or self-regulation and how they interact with consumers. Finally, Glasgow outlines the risks that data brokers pose, and briefly poses some thoughts about their future.

[34] *Big Data: A Tool for Fighting Discrimination and Empowering Groups* (Future of Privacy Forum and Anti-Defamation League Report, 2014), https://fpf.org/wp-content/uploads/Big-Data-A-Tool-for-Fighting-Discrimination-and-Empowering-Groups-FINAL.pdf.

[35] *The State of Mobile Data for Social Good Report* (United Nations Global Pulse, June 2017), http://unglobalpulse.org/sites/default/files/MobileDataforSocialGoodReport_29June.pdf.

[36] Goals statement, Partnership on AI, https://www.partnershiponai.org/.

[37] Woodrow Hartzog & Evan Selinger, *Surveillance as Loss of Obscurity*, 72 Wash. & Lee L. Rev. 1343 (2015).

[38] *See, e.g.*, Evan Selinger & Brett Frischmann, *Utopia?: A Technologically Determined World of Frictionless Transactions, Optimized Production, and Maximal Happiness*, 64 UCLA L. Rev. Disc. 372 (2016).

[39] To ensure all of the chapters are fairly summarized, we asked contributors to provide their own. What follows are versions of their summaries, in some cases verbatim.

In Chapter 3, "In Defense of Big Data Analytics," Mark MacCarthy argues that big data analytics, including machine learning and artificial intelligence, are natural outgrowths of recent developments in computer technology such as the availability of massive data sets, vast increases in computing power, and breakthroughs in analytical techniques. These techniques promise unprecedented benefits for consumers, workers, and society at large, but they also pose challenges for privacy and fairness. MacCarthy's chapter contains a short summary of the range of potential benefits made possible by these new analytic techniques and then discusses privacy and fairness challenges. Principles of privacy policy requiring data minimization and restricting secondary data use need to be reformulated to allow for both the successful delivery of big data benefits and effective privacy protection. Ubiquitous re-identification risks and information externalities reduce the ability of individuals to control the disclosure of information and suggest less reliance on notice and choice mechanisms. Big data analytics can pose fairness challenges, but these techniques are not exempt from existing antidiscrimination and consumer protection laws. Regulatory agencies and courts need to enforce these laws against any abuses accomplished through big data analysis. Disclosure of source code is not an effective way to respond to the challenges of designing and using unbiased algorithms. Instead, enterprises should develop and implement a framework for responsible use of data analytics that will provide for fairness by design and after-the-fact audits of algorithms in use. Such a framework will need to adopt standards of fairness and appropriate remedies for findings of disparate impact. This will require moving beyond technical matters to address sensitive normative issues where the interests of different groups collide and moral intuitions diverge. A collaborative effort of businesses, governments, academics, and civil rights and public interest groups might sharpen the issues and allow sharing of information and best practices in a way that would benefit all.

In Chapter 4, "Education Technology and Student Privacy," Elana Zeide argues that new education technology (ed tech) creates new ways to manage, deliver, and measure education that generate a previously unimaginable array and detail of information about students' actions both within and outside classrooms. She claims that data-driven education tools have the potential to revolutionize the education system – and, in doing so, provide more access to better quality, lower-cost education and broader socioeconomic opportunity. The information generated by such tools also provides fodder for more informed teacher, school, and policy decision-making. At the same time, Zeide maintains, these data practices go against traditional expectations about student privacy. The education context requires a tailored approach to data protection. Few students can opt out of school information practices, making consent-based protections potentially problematic. Maturing data subjects, Zeide cautions, raises concerns about creating modern day "permanent records" with outdated information that unfairly foreclose opportunities. Many fear for-profit providers will prioritize generating revenue over students' educational interests. Traditional student privacy regulations aren't designed for an era of the tremendous volume, variety, and velocity of big data, because they rely on privacy self-management and institutional oversight. Many newer state laws restrict commercial educational technology services to using student data only for "school purposes," but don't cover the potential unintended consequences and nuanced ethical considerations surrounding educational use of data. As a result, Zeide concludes, the responsibility rests on entities generating, collecting, and using student data to adopt best practices to meet the specific expectations and considerations of education environments.

In Chapter 5, "Mobile Privacy Expectations: How Privacy Is Respected in Mobile Devices," Kirsten Martin and Katie Shilton describe privacy challenges raised by mobile devices, explore user privacy expectations for mobile devices, and discuss developer responses to privacy

concerns. Martin and Shilton argue that mobile technologies change social practices and introduce new surveillance concerns into consumers' everyday lives. Yet, consumers, regulators, and even the developers who build mobile applications struggle to define their expectations for the privacy of this data, and consumers and developers express different privacy expectations. The authors argue firms and regulators can help mitigate the gap between developers and consumers by making privacy by design part of corporate governance and establishing privacy as a first-order concern for protecting consumer trust.

In Chapter 6, "Face Recognition, Real-Time Identification, and Beyond," Yana Welinder and Aeryn Palmer provide an overview of face recognition technology and recent efforts to regulate its use. They first explain the process by which face recognition technology operates, including recent advancements in its capabilities through the use of neural networks. They then discuss various consumer applications of the technology, such as mobile apps and social network features that can identify people in photos. Next, they survey regulatory responses to face recognition technology across the globe, highlighting new developments and previewing possible trends to come in the United States, the European Union, Canada, China, and other jurisdictions. The discussion demonstrates the lack of regulation in some areas and reveals global uncertainty about how best to control face recognition technology under the law. The chapter concludes with recommendations to two types of stakeholders. First, it addresses policy-makers, encouraging them to balance support for innovation with protection of individual privacy rights. It stresses the importance of obtaining consent from all relevant parties, and of giving special consideration to government access to privately held face recognition data. Finally, Welinder and Palmer suggest that developers leverage User Experience Design as a notice tool, collect and retain a minimal amount of data, and keep the principles of security by design at the forefront of their minds.

In Chapter 7, "Smart Cities: Privacy, Transparency, Community," Kelsey Finch and Omer Tene argue that today's cities are pervaded by growing networks of connected technologies to generate actionable, often real-time data about the city and its citizens. The more connected a city becomes, the more it will generate a steady stream of data from and about its citizens. As smart city technologies are being rapidly adopted around the globe, we must determine how communities can leverage the benefits of a data-rich society while minimizing threats to individuals' privacy and civil liberties. Just as there are many methods and metrics to assess a smart city's livability, or sustainability, or efficiency, so too there are different lenses through which cities can evaluate their privacy preparedness. This chapter lays out three such perspectives, considering a smart city's privacy responsibilities in the context of its roles as a data steward, data platform, and government authority. By considering the deployment of smart city technologies in these three lights, communities will be better prepared to reassure residents of smart cities that their rights will be respected and their data protected.

Ethical and Legal Reservations about Tracking Technologies

In Chapter 8, "Americans and Marketplace Privacy: Seven Annenberg National Surveys in Perspective," Joseph Turow sketches the growing surveillance and personalized targeting of Americans carried out by marketers as well as the public arguments used to defend these activities. At the core of their justifications is the notion that despite professed concerns over privacy, people are rationally willing to trade information for the relevant benefit that marketers provide. Drawing on seven nationally representative telephone surveys from 1999 through 2015, Turow presents findings that tend to refute marketers' justifications for increased personalized

surveillance and targeting of individuals. Contrary to the claim that a majority of Americans consent to data collection because the commercial benefits are worth the costs, he also shows that the 2015 survey supports a different explanation: a large pool of Americans feel resigned to the inevitability of surveillance and the power of marketers to harvest data. When they give up information as they shop it merely *appears* they are interested in tradeoffs. The overall message of the surveys is that legislators, regulators, and courts ought to rethink the traditional regulatory understanding of harm in the face of a developing American marketplace that ignores the majority of Americans' views and is making overarching tracking and surreptitious profiling an aspect of society taken for granted.

In Chapter 9, "The Federal Trade Commission's Inner Privacy Struggle," Chris Jay Hoofnagle's discusses the cultural and ideological conflicts on privacy internal to the FTC, and explains why the lawyers at the Commission are leading the privacy charge. This is because the Bureau of Economics is constitutionally skeptical of information privacy. Privacy skepticism reflects the economists' academic methods and ideological commitments. While information privacy is a deeply multidisciplinary field, the Bureau of Economics adheres to a disciplinarity that bounds its inquiry and causes it to follow a laissez faire literature. Commitments to "consumer welfare," concerns about innovation policy, lingering effects of Reagan-era leadership, the lack of a clearly-defined market for privacy, and the return of rule of reason analysis in antitrust also contribute to the Bureau of Economics' skepticism toward rights-based privacy regimes. Hoofnagle concludes with a roadmap for expanding the BE's disciplinary borders, for enriching its understanding of the market for privacy, and for a reinvigoration of the FTC's civil penalty factors as a lodestar for privacy remedies.

In Chapter 10, "Privacy and Human Behavior in the Information Age," Alessandro Acquisti, Laura Brandimarte, and George Lowenstein provide a review that summarizes and draws connections between diverse streams of empirical research on privacy behavior. They use three themes to connect insights from social and behavioral sciences: people's uncertainty about the consequences of privacy-related behaviors and their own preferences over those consequences; the context-dependence of people's concern, or lack thereof, about privacy; and the degree to which privacy concerns are malleable – manipulable by commercial and governmental interests. Organizing our discussion by these themes, the authors offer observations concerning the role of public policy in the protection of privacy in the information age.

In Chapter 11, "Privacy, Vulnerability, and Affordance," Ryan Calo argues that the relationship between privacy and vulnerability is complex. Privacy can be both a shield against vulnerability and a sword in its service. What is needed to capture this nuanced interaction is a theoretical lens rooted in the physical and social environments as they exist, but also sensitive to the differing ways people perceive and experience that environment. Calo further contends that James Gibson's theory of affordance is an interesting candidate to capture this complexity, including in the context of consumer privacy. Affordance theory, Calo demonstrates, helps generate and unify some of consumer privacy's most important questions and will perhaps one day lead to better answers.

In Chapter 12, "Ethical Considerations When Companies Study – and Fail to Study – Their Customers," Michelle N. Meyer provides an overview of the different ways in which businesses increasingly study their customers, users, employees, and other stakeholders, and the different reasons why they do so. Meyer argues, however, that a complete ethical analysis of business research requires consideration not only of the purpose, nature, and effects of such research but also of a business's choice *not* to study the effects of its products, services, and practices on stakeholders. Depending on a variety of criteria she discusses, a particular business study – even

one conducted without study-specific informed consent – can fall on a spectrum from unethical to ethically permissible to ethically laudable or even obligatory. Although business research is now ubiquitous – in many ways, happily so – the fact that individual, study-specific informed consent is usually infeasible in this context means that a careful consideration of a study's risks and expected benefits is called for. For reasons that Meyer explains, the *content* of federal regulations that govern risk-benefit analyses of most academic and some industry research – the so-called Common Rule – is not easily translated to the business setting. But she argues that companies should consider adopting something like the *process* used by institutional reviews boards (IRBs) to prospectively review and oversee research, and provides recommendations about how such company "research review boards" might operate.

In Chapter 13, "Algorithmic Discrimination vs. Privacy Law," Alvaro Bedoya addresses the intersection of two pressing debates: the desire to eliminate bias in automated decision-making systems, and the recent industry-led push to enforce privacy protections at the point of data *use*, rather than the point of data *collection*. Bedoya highlights that most proposed solutions to the problem of algorithmic bias have tended to focus on *post*-collection remedies. Honing in on a specific technology, face recognition, Bedoya argues that correcting for algorithmic bias in this way will prove to be difficult, if not impossible. Instead, he says, the most effective means to counter algorithmic discrimination may come at the beginning of the data life cycle – at the point of collection. In making this argument, he emphasizes the importance of collection controls in any comprehensive privacy protection regime.

In Chapter 14, "Children, Privacy, and the New Online Realities," Stephen Balkam discusses the extraordinary challenges we all face in staying private in our hyperconnected lives. He emphasizes the difficulties parents, platforms, and policy makers face in keeping children's data private in an age of connected toys, devices, and always-on connectivity. Balkam looks at the history and evolution of the Children's Online Privacy Protection Act (COPPA) and addresses its benefits and shortcomings. He looks at how major social media platforms, such as Facebook, have responded to COPPA as well as some of the companies that have fallen foul of the law. In addition to considering the likes of Hello Barbie and Amazon's Echo, Balkam also considers the range of potential privacy issues brought by innovations in virtual, augmented, and mixed reality devices, apps and games. He concludes with a look at the future of children's privacy in an AI-infused, constantly monitored world. Balkam suggests that solutions will have to be found across the public, private, and non-profit sectors and then communicated clearly and consistently to parents and their digitally savvy children.

In Chapter 15, "Stakeholders and High Stakes: Divergent Standards for Do Not Track," Aleecia M. McDonald provides an in-depth look at the history of Do Not Track, informed by McDonald's personal experience as an original cochair of the World Wide Web Committee standards group. In the United States, the Do Not Call list is considered one of the big successes in consumer privacy. In contrast, Do Not Track was dubbed "worse than a miserable failure" before it even got out of the standards committee trying to define it. At this time, Do Not Track is a soon-to-be published standard from the World Wide Web Committee (W3C), where standards emerge for web technologies such as HTML, which is the language of web pages. Meanwhile, the Electronic Frontier Foundation (EFF), an online rights group, has devised its own privacy-enhanced version of Do Not Track, with multiple companies pledging to use it. Several ad blockers will permit ads from companies that honor EFF's Do Not Track, providing a carrot and stick approach to user privacy and control. In yet a third approach, Do Not Track was suggested as a way to signal compliance with European Union privacy laws, both in a recent international Privacy Bridges project, as well as in publications by this author and leading European privacy

scholars. The best thing about standards, as the saying goes, is that there are so many to choose from. Yet from a user's perspective, how can the multiplicity of Do Not Track approaches be anything but confusion?

In Chapter 16, "Applying Ethics When Using Data Beyond Individuals' Understanding," Martin E. Abrams and Lynn A. Goldstein contend that with the expanding use of observational data for advanced analytics, organizations are increasingly looking to move beyond technical compliance with the law to the ethical use of data. Organizations need to understand the fair processing risks and benefits they create for individuals, whether they are ethically appropriate, and how they might be demonstrated to others. Their chapter explores the evolution of data-driven research and analytics, discusses how ethics might be applied in an assessment process, and sets forth one process for assessing whether big data projects are appropriate.

International Perspectives

In Chapter 17, "Profiling and the Essence of the Right to Data Protection," Bilyana Petkova and Franziska Boehm begin by reviewing the legislative history of the provision on automated decision-making in the 1995 EU Data Protection Directive (the 1995 Directive), as it was amended in the process of adopting a new EU General Data Protection Regulation that would enter into force in 2018. Next, they discuss profiling in the context of the case law of the Court of Justice of the European Union (CJEU) in the *Google Spain*, *Digital Rights Ireland*, and *Schrems* cases. Petkova and Boehm argue that the CJEU might be making a subtle move in its interpretation of the EU Charter of Fundamental Rights toward protecting against undesirable profiling measures instead of merely protecting against the identification of an individual. Finally, from the employment context, they discuss a few hypotheticals of algorithmic decision-making that illustrate how the relevant legislative framework might be applied.

In Chapter 18, "Privacy, Freedom of Expression, and the Right to be Forgotten in Europe," Stefan Kulk and Frederik Zuiderveen Borgesius discuss the relation between privacy and freedom of expression in Europe. In principle, the two rights have equal weight in Europe – which right prevails depends on the circumstances of a case. To illustrate the difficulties when balancing privacy and freedom of expression, Kulk and Borgesius discuss the *Google Spain* judgment of the Court of Justice of the European Union, sometimes called the "right to be forgotten" judgment. The court decided in *Google Spain* that, under certain conditions, people have the right to have search results for their name delisted. The authors discuss how Google and Data Protection Authorities deal with such delisting requests in practice. Delisting requests illustrate that balancing the interests of privacy and freedom of expression will always remain difficult.

In Chapter 19, "Understanding the Balancing Act Behind the Legitimate Interest of the Controller Ground: A Pragmatic Approach," Irene Kamara and Paul De Hert analyse the provision of the legitimate interest ground in the new EU data protection framework, the General Data Protection Regulation. The authors explain that the rationale of the legitimate interest ground is that under certain conditions, controllers' or third parties' interests might be justified to prevail over the interests, rights, and freedoms of the data subject. *When* and *how* the prevailing may take place under the GDPR provisions is not a one-dimensional assessment. De Hert and Kamara suggest a formalisation of the legitimate interest ground steps toward the decision of the controller on whether to base his or her processing on the legitimate interest ground. They argue that the legitimate interest ground should not be seen in isolation, but through the lens of the data protection principles of Article 5 GDPR and Article 8 Charter

Fundamental Rights EU. The authors further analyse the relevant case law of the Court of Justice EU, as well as the cases of Network and Information Security and Big Data and Profiling. Kamara and De Hert conclude that the legitimate interest of the controller is not a loophole in the data protection legislation, as has often been alleged, but an equivalent basis for lawful processing, which can distinguish controllers in bad faith from controllers processing data in good faith.

New Approaches to Improve the Status Quo

In Chapter 20, "The Intersection of Privacy and Consumer Protection," Julie Brill explores the intersection between privacy and consumer protection in the United States. She surveys the consumer protection laws that simultaneously address privacy harms, and also examines how the Federal Trade Commission's consumer protection mission has allowed the Commission to become a lead privacy regulator. Along the way, Brill delves into the challenges posed by data brokers, lead generators, and alternative credit scoring – as well as potential avenues for the United States to strengthen privacy protections.

In Chapter 21, "A Design Space for Effective Privacy Notices," Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor argue that notifying users about a system's data practices is supposed to enable users to make informed privacy decisions. Yet, current notice and choice mechanisms, such as privacy policies, are often ineffective because they are neither usable nor useful, and are therefore ignored by users. Constrained interfaces on mobile devices, wearables, and smart home devices connected in an Internet of Things exacerbate the issue. Much research has studied the usability issues of privacy notices and many proposals for more usable privacy notices exist. Yet, there is little guidance for designers and developers on the design aspects that can impact the effectiveness of privacy notices. In this chapter, Schaub, Balebako, Durity, and Cranor make multiple contributions to remedy this issue. They survey the existing literature on privacy notices and identify challenges, requirements, and best practices for privacy notice design. Further, they map out the design space for privacy notices by identifying relevant dimensions. This provides a taxonomy and consistent terminology of notice approaches to foster understanding and reasoning about notice options available in the context of specific systems. Our systemization of knowledge and the developed design space can help designers, developers, and researchers identify notice and choice requirements and develop a comprehensive notice concept for their system that addresses the needs of different audiences and considers the system's limitations and opportunities for providing notice.

In Chapter 22, "Enter the Professionals: Organizational Privacy in a Digital Age," J. Trevor Hughes and Cobun Keegan observe that contemporary privacy professionals apply legal, technological, and management knowledge to balance the important concerns of citizens and consumers with the interests of companies and governments worldwide. They further note that the field of information privacy has rapidly matured into an organized, interdisciplinary profession with international reach. Their chapter compares the burgeoning privacy profession with other modern professions, describing its history and similar growth curve while highlighting the unique characteristics of a profession that combines law, policy, technology, business, and ethics against a rapidly shifting technological landscape. As it has grown into a profession, Hughes and Keegan argue that privacy has developed a broad body of knowledge with multiple specialties, gained recognition as a vital component of organizational management, and become formally organized through professional associations and credentialing programs. Government recognition and enforcement actions have legitimized the role of privacy professionals even as these

professionals work collectively to synthesize comprehensive and lasting ethical norms. In an era increasingly fueled and defined by data, significant changes in the shape of our economy and professional workforce are inevitable. By guiding the governance and dissemination of personal information, Hughes and Keegan argue that the privacy profession is well situated to grow and mature in these rapidly changing times.

In Chapter 23, "Privacy Statements: Purposes, Requirements, and Best Practices," Mike Hintze addresses common criticisms of privacy statements and argues that many criticisms misunderstand the most important purposes of privacy statements, while others can be addressed through careful and informed drafting. Hintze suggests that while drafting a privacy statement may be considered by some to be one of the most basic tasks of a privacy professional, doing it well is no simple matter. One must understand and reconcile a host of statutory and self-regulatory obligations. One must consider different audiences who may read the statement from different perspectives. One must balance pressures to make the statement simple and readable against pressures to make it comprehensive and detailed. A mistake can form the basis for an FTC deception claim. And individual pieces can be taken out of context and spun into public relations debacles. Hintze's chapter explores the art of crafting a privacy statement. It explains the multiple purposes of a privacy statement. It lists and discusses the many elements included in a privacy statement – some required by law and others based on an organization's objectives. Finally, it describes different approaches to drafting privacy statements and suggests best practices based on a more complete understanding of a privacy statement's purposes and audiences.

In Chapter 24, "Privacy Versus Research in Big Data," Jane R. Bambauer analyzes how traditional notions of privacy threaten the unprecedented opportunity to study humans in the Big Data era. After briefly describing the set of laws currently constraining research, the chapter identifies puzzles and potential flaws in three popular forms of privacy protection. First, data protection laws typically forbid companies from repurposing data that was collected for a different, unrelated use. Second, there is a growing appreciation that anonymized data can be reidentified, so regulators are increasingly skeptical about using anonymization to facilitate the sharing of research data. And third, research law generally prohibits researchers from performing secret interventions on human subjects. Together, these restrictions will interfere with a great amount of Big Data research potential, and society may not get much in return for the opportunity costs.

In Chapter 25, "A Marketplace for Privacy: Incentives for Privacy Engineering and Innovation," Courtney Bowman and John Grant inquire into what drives businesses to offer technologies and policies designed to protect consumer privacy. The authors argue that in capitalist systems, the primary levers would be market demand supplemented by government regulation where the market fails. But when it comes to privacy, consumers' demand can appear inconsistent with their expressed preferences, as they ignore high-profile data breaches and gleefully download trivial smartphone apps in exchange for mountains of their own personal data. Yet, even in places where government regulation is light (such as the United States), many companies increasingly appear to be pursuing high profile – and sometimes costly – positions, practices, and offerings in the name of protecting privacy. Ultimately, Bowman and Grant suggest that in order to understand the true market for privacy, beyond consumer-driven demand, it is necessary also to consider the ethos of the highly skilled engineers who build these technologies and their level of influence over the high-tech companies that have created the data economy.

In Chapter 26, "The Missing Role of Economics in FTC Privacy Policy," James Cooper and Joshua Wright note that the FTC has been in the privacy game for almost twenty years. In that

time span, the digital economy has exploded. As a consequence, the importance to the economy of privacy regulation has grown as well. Unfortunately, Cooper and Wright insist, its sophistication has yet to keep pace with its stature. As they see it, privacy stands today where antitrust stood in the 1970s. Antitrust's embrace then of economics helped transform it into a coherent body of law that – despite some quibbles at the margin – almost all agree has been a boon for consumers. Cooper and Wright thus argue that privacy at the FTC is ripe for a similar revolution. The chapter examines the history of FTC privacy enforcement and policy making, with special attention paid to the lack of economic analysis. It shows the unique ability of economic analysis to ferret out conduct that is likely to threaten consumer welfare, and provide a framework for FTC privacy analysis going forward. Specifically, Cooper and Wright argue that the FTC needs to be more precise in identifying privacy harms and to develop an empirical footing for both its enforcement posture and such concepts as "privacy by design" and "data minimization." The sooner that the FTC begins to incorporate serious economic analysis and rigorous empirical evidence into its privacy policy, the authors maintain, the sooner consumers will begin to reap the rewards.

In Chapter 27, "Big Data by Design: Establishing Privacy Governance by Analytics," Dale Skivington, Lisa Zolidis, and Brian P. O'Connor argue that a significant challenge for corporate big data analytics programs is deciding how to build an effective structure for addressing privacy risks. They further contend that privacy protections, including thoughtful Privacy Impact Assessments, add essential value to the design of such programs in the modern marketplace where customers demand adequate protection of personal data. The chapter thus provides a practical approach to help corporations weigh risks and benefits for data analytics projects as they are developed to make the best choices for the products and services they offer.

In Chapter 28, "The Future of Self-Regulation is Co-Regulation," Ira Rubinstein contends that privacy self-regulation – and especially voluntary codes of conduct – suffers from an overall lack of transparency, weak or incomplete realization of the Fair Information Practice Principles, inadequate incentives to ensure wide-scale industry participation, and ineffective compliance and enforcement mechanisms. He argues that the US experiment with voluntary codes has gone on long enough and that it is time to try a new, more co-regulatory approach. In co-regulation, firms still enjoy considerable flexibility in shaping self-regulatory guidelines, but consumer advocacy groups have a seat at the table, and the government retains general oversight authority to approve and enforce statutory requirements. Rubenstein examines three recent co-regulatory efforts: (1) privacy management programs designed by multinational firms to demonstrate accountability under both European and US privacy laws; (2) the NTIA multistakeholder process, under which industry and privacy advocates have sought to develop voluntary but enforceable privacy codes without any explicit legal mandate; and (3) Dutch codes of conduct under national data protection law, which allows industry sectors to draw up privacy codes specifying how statutory requirements apply to their specific sector. He concludes by identifying lessons learned and offering specific policy recommendations that might help shape any future consumer privacy legislation in the United States or abroad.

In Chapter 29, "Privacy Notices: Limitations, Challenges, and Opportunities," Mary J. Culnan and Paula J. Bruening contend that openness is the first principle of fair information practices. While in practice "notice" has been used to create openness, notices have been widely criticized as being too complex, legalistic, lengthy, and opaque. Culnan and Bruening argue that to achieve openness, data protection should move from a "notice" model to a model that requires organizations to create an environment of "transparency." They assert that while often used interchangeably, the terms "notice" and "transparency" are not synonymous. In their

chapter, Culnan and Bruening review the history of notice in the United States, its traditional roles in data protection, the challenges and limitations of notice, the efforts to address them, and the lessons learned from these efforts. They examine the challenges emerging technologies pose for traditional notice and propose a move away from a reliance on notice to the creation of an environment of transparency that includes improved notices, attention to contextual norms, integrating notice design into system development, ongoing public education, and new technological solutions. Finally, Culnan and Bruening present arguments for business buy-in and regulatory guidance.

In Chapter 30, "It Takes Data to Protect Data," David A. Hoffman and Patricia A. Rimo note that we live in a world of constant data flow, and safeguarding data has never been more important. Be it medical records, financial information or simple online passwords, the amount of private data that needs to be protected continues to grow. Along with this growth in the need to secure data, Hoffman and Rimo insist, however, are the privacy concerns people have with their data. While some would pit security and privacy against each other, arguing that individuals must choose one over the other, the two actually can and should reinforce each other. It's this model that forms the basis of the chapter: Privacy and security should be pursued hand-in-hand as we move toward an increasingly connected, digital world. To fully realize the benefits of information technology, big data, and Internet of Things, Hoffman and Rimo argue individuals must be confident that their devices are designed in a way that protects their data and that any data being collected and processed from those devices is used responsibly. Using internationally recognized mechanisms such as the Fair Information Privacy Principles, public and private organizations can enable both the innovative and ethical use of data. The key is not avoiding data but using it mindfully. It takes data to protect data.

In Chapter 31, "Are Benefit-Cost Analysis and Privacy Protection Efforts Incompatible?" Adam Thierer argues that benefit-cost analysis (BCA) helps inform the regulatory process by estimating the benefits and costs associated with proposed rules. At least in the United States, BCA has become a more widely accepted part of regulatory policy-making process and is formally required before many rules can take effect. The BCA process becomes far more contentious, however, when the variables or values being considered are highly subjective in character. This is clearly the case as it pertains to debates over online data collection and digital privacy. The nature and extent of privacy rights and privacy harms remain open to widely different conceptions and interpretations. This makes BCA more challenging, some would say impossible. In reality, however, this same problem exists in many different fields and does not prevent BCA from remaining an important part of the rule-making process. Even when some variables are highly subjective, others are more easily quantifiable. Thierer thus contends that policymakers should conduct BCA for any proposed rules related to data collection and privacy protection to better understand the trade-offs associated with those regulatory proposals.

In Chapter 32, "Privacy After the Agile Turn," Seda Gürses and Joris van Hoboken explore how recent paradigmatic transformations in the production of everyday digital systems are changing the conditions for privacy governance. Both in popular media and in scholarly work, great attention is paid to the privacy concerns that surface once digital technologies reach consumers. As a result, the strategies proposed to mitigate these concerns, be it through technical, social, regulatory or economic interventions, are concentrated at the interface of technology consumption. The authors propose to look beyond technology consumption, inviting readers to explore the ways in which consumer software is produced today. By better understanding recent shifts in software production, they argue, it is possible to get a better grasp of how and why software has come to be so data intensive and algorithmically driven, raising a

plethora of privacy concerns. Specifically, the authors highlight three shifts: from waterfall to agile development methodologies; from shrink-wrap software to services; and, from software running on personal computers to functionality being carried out in the cloud. Their shorthand for the culmination of these shifts is the "agile turn." With the agile turn, the complexity, distribution, and infrastructure of software have changed. What are originally intended to be techniques to improve the production of software development, e.g., modularity and agility, also come to reconfigure the ways businesses in the sector are organized. In fact, the agile turn is so tectonic, it unravels the authors' original distinction: The production and consumption of software are collapsed. Services bind users into a long-term transaction with software companies, a relationship constantly monitored and improved through user analytics. Data flows, algorithms, and user profiling have become the bread and butter of software production, not only because of business models based on advertisements, but because of the centrality of these features to a successful disruptive software product. Understanding these shifts has great implications for any intervention that aims to address, and mitigate, consumer privacy concerns.