Columbia Business School
AT THE VERY CENTER OF BUSINESS™
Sanford C. Bernstein & Co. Center
for Leadership and Ethics

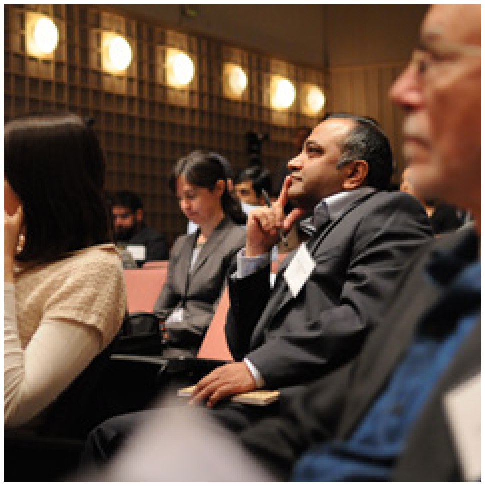

Innovation and The Value of Privacy

A joint project of the Data Science Institute and the Sanford C. Bernstein & Co. Center for Leadership and Ethics

Foreword

In 1999, Scott McNealy, Sun Microsystems' then-CEO, reportedly said, "You have zero privacy anyway. Get over it." Customers may claim that they value their privacy, but, in practice, they often give or trade it away at low value. Companies are then able to monetize this information by aggregating and analyzing it for their use or by selling or trading it with third parties.

Since the uproar over McNealy's frank words, the privacy-invading capabilities of businesses, governments, and motivated third parties have greatly expanded through the continuing progress made in data sciences and in analyzing big data. The ubiquity of points of access to once-personal data is stunning, moving far beyond websites visited to surveillance cameras images, voice pattern recognition, text analysis, and deep learning models. The situation has moved from protecting what individuals don't want to share to preventing intelligent systems from learning personal preferences and behaviors that individuals don't even know about themselves.

On February 5th, 2016, the Data Science Institute and the Sanford C. Bernstein & Co. Center for Leadership and Ethics at Columbia University brought together business leaders and academics to discuss how businesses are incorporating privacy considerations into their business strategies. Speakers included industry leaders, government regulators, and entrepreneurs from frontier firms, including Microsoft, IBM, Deloitte, RedMorph and Dstillery. They shared experiences regarding how they are accessing or protecting the customer data while setting new industry practices around privacy. Meanwhile the academics untangled often complex concepts of privacy and privacy harms or offered new algorithmic tools by which to counter intrusions on privacy. Above all, this was a conference that looked to solutions to the challenges of preserving privacy.

By viewing the video of the conference on the Bernstein Center website and by reading this report, prepared and written by Nancy J. Brandwein, you will find a lot of robust disagreement. We decided to highlight these disagreements to showcase the current debates on privacy. This debate is not only between the expected parties, such as US and European regulators, or hackers and companies, but also between and among firms. Echoing an earlier generation of software firms that identified viruses, startups today are being created to improve privacy protection as a way to differentiate their products.

For example, while firms such as Apple or Blackberry are competing on their reputations for keeping personal data private, others are actively monetizing data from consumers who are not vigilant or simply don't care about protecting their data. The economic explanation is that markets sort people by their heterogeneous (i.e., different) preferences and those outcomes are efficient. The human rights approach views privacy as inalienable: privacy cannot be sold and bought. This fundamental divide proved to be a touchstone for debate during the conference.

Keynote Presenter and Moderator:
Kate Crawford, Principal Researcher at
MSR-NYC (Microsoft) and Visiting Professor,
MIT Center for Civic Media

INNOVATION AND THE VALUE OF PRIVACY[1]

We are very thankful to the speakers, many of whom came a long way, and to our moderators Geoff Heal and Alondra Nelson. Kate Crawford was a superb keynote speaker, and she also advised us on invitations to other speakers. A number of speakers have connections to Columbia Engineering and Data Sciences or to the Business School. We thank our staff, Bianca Bellino and Sandra Navalli at the Bernstein Center, and Leslie Barna, Brianne Cortese and Jonathan Stark at the Data Science Institute. We also thank our Deans, Glenn Hubbard and Kathy Phillips from the Business School, Merit Janow from the School of International and Public Affairs, Alondra Nelson from Social Sciences and Mary Boyce from the Columbia University's Fu Foundation School of Engineering and Applied Science.

It was a lively and fast-paced workshop and both DSI and the Bernstein Center are pleased to present this account of their joint efforts.

Bruce Kogut
Sanford C. Bernstein & Co. Professor/Director
Columbia Business School
Columbia University

Kathleen McKeown
Henry and Gertrude Rothschild Professor of Computer Science
Director of the Date Science Institute
Columbia University

[1] Prepared and written by Nancy J. Brandwein.

Participants

In order of symposium agenda

Kathy McKeown
Director of the Data Science Institute

Kate Crawford
Principal Researcher at
MSR-NYC (Microsoft) and Visiting Professor,
MIT Center for Civic Media

Roxana Geambasu
Assistant Professor of Computer Science at
Columbia University

Deirdre Mulligan
Professor of Information at
UC Berkeley School of Information

Arvind Narayanan
Associate Professor at
Princeton University

Geoff Heal
Donald C. Waite III Professor of
Social Enterprise and Bernstein Faculty Leader
at Columbia Business School (moderator)

Brenda L Dietrich
IBM Fellow, Vice President,
Data Science, IBM

Abhay Edlabadkar '07
Founder of RedMorph

Claudia Perlich
Chief Data Scientist at Dstillery

Bruce Kogut
Director Sanford C. Bernstein & Co. Center for
Leadership and Ethics (moderator)

Irfan Saif
US Advisory Leader, Technology at Deloitte

Ashkan Soltani
Former Senior Advisor at the White House
Office of Science and Technology Policy and
Independent Researcher

Alondra Nelson
Professor and
Dean of Social Science at
Columbia University (moderator)

From the 10,000 recommended steps our Fitbits measure to the e-books we read, the text and e-mails we send and receive, the photos we share and the news videos we stream online, more and more of our daily activities are mediated through digital devices. And as we consume digital information and complete transactions, we provide personal information, often involuntarily, to industry, institutions and government. With the advancements of data science, that information becomes data points, sometimes transformed into inferences about everything from our buying behavior to our creditworthiness, from our gender to our sexual orientation. For business, government, and science, the more information the better—to target marketing, track terrorists, and personalize medicine. But if the innovations of big data have value, what about the value of the privacy lost in the process—or in the processing? Is privacy a human right or a societal value? In just a half-day long conference, there was no shortage of debate on what innovations in data science and predictive modelling imply for privacy and ethics. Kate Crawford, Principal Researcher at Microsoft Research, New York City, got straight to the heart of the controversy in her keynote address.

What does a Syrian refugee camp have to do with data science, innovation and privacy? Everything, as it turns out. Crawford said that Europe's burgeoning Syrian refugee camps are the site of a "fascinating machine-learning experiment." While she clicked on a slide of a mother and child walking between a row of grey tents, Crawford explained that an IBM initiative is turning **unstructured data** on the millions of Syrian refugees entering Europe into actionable intelligence.

On hearing concerns that healthy fighting-age males were among the many emaciated and sick refugees coming off the boats, IBM's team in Europe developed their i2 Enterprise Insight Analysis program to separate the sheep from the wolves.[2] Using real data from the **dark web**, such as where the refugees were buying passports, and combining it with a set of "fake data" to which a border guard would have access, IBM claims its computer model can create what amounts to a jihadist "credit score." Crawford said analysts were quick to say that such a score was not an absolute indicator of guilt or innocence, yet the model concocted incredibly detailed portraits of refugees—their names, religious beliefs, prior addresses and **social graphs**. Much of the unstructured data would seem to be innocuous, including municipal parking tickets and "likes" for certain DJs on Facebook and Twitter. Of the many reasons to be concerned about the refugee initiative, Crawford cited the potential for finding unwarranted correlations and false positives, the problem of verifying suspicious scores, and the unreliability of data drawn from the dark web.

Yet, the privacy harm Crawford finds most worrisome, and one raised throughout the conference, is the potential for discrimination:

> **Machine learning** models are creating a kind of intensification of surveillance of the least privileged, being tested out on the populations who are least able to object. We have a new and relatively untested system for predicting terrorists while the mechanism that's being used is, itself, in a black box. So if you are, for instance, one of the accused, it's very difficult to figure out what that data is about you or why that conclusion has been drawn, so that highlights for me the kernel of what we are here to discuss today: privacy, as it is currently configured, has no means to deal with a situation like this.

Crawford went on to outline three themes of the conference woven throughout two sessions, one on academic research and one on industry, and in the final keynote speakers' presentations as well:

Privacy's Threat Models: Privacy has always been built on mental models of the technologies of the time and ideas of human behavior. These threat models, Crawford explained, have evolved from Peeping Toms at the window and print newspapers revealing unwelcome tidbits to the development and use of the camera in the late 1880s, and then on to business and government databases of the 1960s through 1980s in which computer systems simply collected what you gave them. It was during this

[2] http://www.defenseone.com/technology/2016/01/refugee-or-terrorist-ibm-thinks-its-software-has-answer/125484/

Speakers: Arvind Narayanan, Associate Professor at Princeton University, Deirdre Mulligan, Professor of Information at UC Berkeley School of Information, and Roxana Geambasu, Assistant Professor of Computer Science at Columbia University

time period that privacy principles of notice and consent for access to your data and security and enforcement emerged. Yet Crawford, and almost every conference speaker, acknowledged that this transactional model of privacy, which still exists today, is inadequate to protect us from privacy threats of **metadata**, or data that provide information about other data, and **machine learning**. In the 2010s to 2020s, online trackers, mobile devices, wearable sensors and more are capturing enormous streams of behavioral data users never intended to share. "A predictive system can't give you notice," said Crawford, and, "it doesn't know how it's going to use your data, so it can't ask you for consent."

The Rise of Machine Learning: Machine learning has given society significant benefits, from Amazon's sophisticated recommendation systems to IBM's Watson, which draws on a corpus of public information about cancer protocols and treatments to offer patients personalized care. But there are enormous risks in how your data is stored and used, Crawford stressed. "Because predictions have material outcomes, some people are going to get benefits and opportunities and some are going to face negative consequences," she warned. Carnegie Mellon researchers found Google's online advertising system showed ads for high-income jobs to men more frequently than it showed the same ads to women. Ads for arrest records have been shown to pop up more in online searches for distinctly black names or historic black fraternities, and the Federal Trade Commission (FTC) has found advertisers target people in poor neighborhoods with ads for high-interest loans.[3]

Pointing to these discriminatory outcomes, Crawford questioned the autonomy of machine learning systems.

> "The greater autonomy these systems have," she said, "the more questions we need to ask about what kind of values or models of the world we use to create them. And what's interesting here is that these models get better only if they ingest more and more training data. That's how they get more intelligent... How can we ever know the value of data now or in the future? If we can't understand the systems we are giving our data to, how can we have equal power? How do you protect information from someone predicting it? What does it mean to have privacy rights to something personal, e.g. symptoms of a life-threatening disease or the desire to leave a relationship, that you haven't consciously recognized yet yourself? What if you're being judged on something you haven't yet done?"

The Turn to Ethics: Crawford said that privacy alone is insufficient to deal with the kinds of challenges created by the rise of machine learning, but "privacy and ethics get us somewhere. And ethics is becoming a core part of data science curriculum." She said there was promising work being done around **privacy by design**, an approach to systems engineering

[3] http://www.nytimes.com/2015/07/10/upshot/when-algorithms-discriminate.html?_r=0

Privacy Cheat Sheet

The multi-dimensional concept of privacy was characterized by some of the following touch points for presenters throughout the conference. Many of these dimensions are overlapping, proving why jurists, legal scholars, philosophers and now data scientists have yet to reach a satisfying definition of privacy. Lawyer and philosopher Tom Gerety once observed that privacy has "a protean capacity to be all things to all lawyers." And Yale Law School Dean Robert Post lamented that "[p]rivacy is a value so complex, so engorged with various and distinct meanings that I sometimes despair whether it can be usefully addressed at all."[4] Yet at the conference, Dierdre Mulligan valiantly introduced these dimensions of privacy in a dizzying ten minutes or less, placing them in two buckets (even these classifications overlap): those arising from philosophy and law and those developed from research in computer science. Which ones are salient rely heavily on context.

Philosophy and Law

"The right to be let alone" is considered the foundation of privacy law in the US. Formulated by Samuel Warren in an 1890 ruling, it began as a defense of physical privacy invasion but has been used against privacy breaches such as wiretapping without a warrant.

Psychological privacy is the right to keep secrets, and it can be violated by the public disclosure of previously concealed information.

Autonomous decision-making is the right to independence from the state in making the kinds of decisions that are central to defining a person, such as those relating to marriage, procreation, contraception, and child-rearing.

Personhood, or the protection of the integrity of one's personality and of what is irreducible in a human being, is perhaps the broadest dimension of privacy and more akin to a human right than a societal value. Such privacy should protect what Warren and Brandeis called one's "inviolate personality." It's a hop, skip, and a jump from there to **anti-totalitarianism** in which privacy helps mitigate against the state imposing a defining identity upon an individual.

Control over personal information simply states that one should be able to control the information that others have about oneself; however its interpretation depends upon what is meant by "control," what is considered "personal" and what is considered "information."[5]

Computer Science Research

Anonymity: the ability to interact without revealing personally identifying information.

Confidentiality: information shared between two parties that is not made public, an example of which is the privacy agreement signed between patient and doctor.

Boundary regulation: control of information shared between individuals and institutions, including third party dissemination.

Differential privacy: the cryptographic protection of data to permit accurate queries from statistical databases while minimizing the chances of identifying records and the individual.

[4] Solove, Daniel J., "Conceptualizing Privacy," California Law Review, Vol.90. Iss.4, July 2002, pp.1087-1054.

[5] Ibid.

Co-Sponsor of the Event: Kathy McKeown, Director of the Data Science Institute

that takes privacy into account throughout the whole engineering process, and the precise, mathematical **differential privacy**, which seeks to protect against the de-anonymization threats mentioned earlier. And she pointed to exciting research coming out of the privacy and ethics space on transparency and fairness,[6] human autonomy vs. machine autonomy,[7] power asymmetries, and due process—the ways you can have legal recourse to intervene on a system that has made a decision about you.[8]

Tackling any of these themes is daunting, but Crawford said, "I am thrilled that academic researchers, industry practitioners and government regulators are coming together, because it's only by putting all our heads together that we will resolve some of the very real concerns… about data privacy today."

What We Talk About When We Talk About Privacy

To understand privacy and how to value it in the big data era, we need to know what it is. In the research session, Deirdre Mulligan, Associate Professor of the University of California Berkeley's School of Information, took on this challenge of untangling what she called an unwieldy ball of yarn. Building on Crawford's outline of threat models, Mulligan said that legislators, engineers and academics all need to understand "what it is people are trying to protect against what kinds of threats."

"Privacy is like democracy and art," stated Mulligan, "We are never going to nail it down." The best way to think about privacy, she said, is as a touch point. "Part of the term 'privacy's' work in the world," Mulligan told attendees, "is to give us a vehicle to have conversations about privacy's meaning, in practice and in everyday life"—and in conferences like this one. She recommended viewing privacy as a living concept, more as an unfinished "tapestry" that's fringy at the edges and that invites us to add to it and, as she said, "expand how we operationalize privacy to protect against today's threats." We can't build the right solutions unless we are sure we are focused on the "right" privacy. Yet with so many definitions of privacy (see "Privacy Cheat Sheet"), and in a world constantly being reinvented by technology, we need a heuristic rule to forge this malleable and multi-faceted concept of privacy that can meet the requirements of actual situations.

As an example of what happens when the "wrong" privacy is addressed, Mulligan cited the Transportation Safety Authority's (TSA's) full body scanners or "naked machines." It turns out passengers' concerns on implementing the scanners were not what the government thought they would be—that their **Personally**

[6] http://www.stanfordlawreview.org/online/privacy-and-big-data/its-not-privacy-and-its-not-fair

[7] http://sth.sagepub.com/content/early/2016/05/31/0162243916650491.full.pdf+html

[8] http://lawdigitalcommons.bc.edu/cgi/viewcontent.cgi?article=3351&context=bclr

Keynote Speaker: Irfan Saif, US Advisory Leader, Technology at Deloitte

Research Speaker: Roxana Geambasu, Assistant Professor of Computer Science at Columbia University speaking to Bruce Kogut, Director of The Sanford C. Bernstein & Co. Center for Leadership and Ethics

Identifiable Information (PII) was being stored in a massive cloud-based park of servers. No, their concerns are the ones that Mulligan, as the mother of a teenage girl, felt keenly: guys in a booth ogling bodies." The privacy threat here is access to the physical self, exposure of one's naked body, and dignity interests. And design solutions that protect against privacy solutions need to be geared toward what is supposed to protected.

Is It Privacy Or Is It Fairness? Power Asymmetries Between Users And Businesses

Though not explicitly discussed, definitions of privacy also differ by cultures and countries. European countries are much more eager than the US to view privacy in the context of an inviolable human right. This right includes a derivative "right to be forgotten," resulting from a 2014 European Court Ruling, which grants individuals the right, under certain conditions, to ask search engines to remove personal information about them.[9] The US tendency is to treat privacy more as property, where users give consent to provide data to the service-provider as a condition for the service. This consent may arguably be voluntary, but there is also a coercive element to it. If everyone is on Facebook, the decision not to give consent is effectively to be socially ostracized. And once an individual agrees to share data, the data may be kept in perpetuity. These issues, if not explicit in the conference panels, came up in Q&As and nevertheless set the stage for one of the more interesting debates regarding the relative powerlessness of consumers in many situations.

Are we focused on the wrong privacy threats? In her keynote address, Kate Crawford stated, "…while these models are ultimately creating very detailed pictures of us that have privacy implications, their outcomes have implications that go far beyond privacy, that touch on things like discrimination, how we think about fairness… Privacy discourse has failed to account for growing power asymmetries between data industries and the rest of us, getting to the point where I fear those asymmetries."

Data scientist and consultants working in business often have a different view of these asymmetries. Industry speaker Brenda Dietrich, IBM Fellow and Vice President in the IBM Research Divisions, is one of them. "If you think back to pre-Internet days, consumers had relatively little information about prices, quality, availability," she said. Dietrich contrasted that with the process of buying a car today, in which the consumer is armed with vastly more information and there is much more transparency for the consumer:

> Today you can know what the last 100 or so models of that car sold for. You have price

[9] http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf

Question and answer sessions with members from the greater community

> information at your fingertips. You have the upper hand. Retailers, consumer packaged goods companies, etc. were very concerned about this asymmetry. 'If my customers know everything there is about my product and competitor product, I want to know everything there is to know about my customer so I can have an equal amount of information that levels the playing field again.' We can all argue, 'Well, the businesses had the advantages for decades, so it's our turn as consumers.' The reality is that they [businesses] don't think that way.

No they don't. Let's take Dietrich's new car example. While we as customers have more access to information on dealerships and car models than ever before, when we go online to search for the abundance of information on that new car, websites, browsers, companies, and third party data brokers are collecting our information to show us targeted ads for cars, and, privacy advocates argue, that they then have the upper hand. Speaking in the research session, Arvind Narayan, Assistant Professor of Computer Science at Princeton, said:

> In classical economics we say that users know something about themselves that the seller or information provider doesn't know. So there's an information asymmetry. What's happening here is the exact opposite. There's a lot of tracking going on. The people doing it know what information they're sending and who they're sending it to. You don't have that information yourself, which hurts your ability to make an informed decision in that transaction.

This is a bold observation. Indeed, no less than a Nobel Prize was given to Berkeley Professor George Akerlof for his analysis of the failure of markets to find prices for used cars when the seller knows whether or not the car is a 'lemon'. In the new age of data science, the buyer or aggregator will often know more about the quality of the good being sold than the seller.

Transparency and Measurement: Antidotes to Information Asymmetry?

Is transparency the answer to the new economics of the information asymmetry problem? The research session moderated by Geoffrey Heal, Donald Waite III Professor of Social Enterprise at the Columbia Business School, suggested yes. Roxana Geambasu, Assistant Professor of Computer Science at Columbia University, and Arvind Narayan both presented their own research solutions that use transparency and measurement to counter tracking mechanisms.

Geambasu introduced "Sunlight," the first system that can diagnose online ad targeting in fine detail, at a wide scale, and with solid statistical justification. This second-generation tool builds on its predecessor "X-ray." Geambasu and her fellow researchers set up a Gmail account and populated it with simple single-topic e-mails. The researchers then set up 119 Gmail accounts and sent 300 e-mail messages all with sensitive keywords both in the subject line and bodies of the e-mail. Of the ads shown, about 15 percent appeared to be targeted, and more worrying, many violated Google's own policy against targeting ads based on "race, religion, sexual orientation, health, or sensitive financial categories." For instance, an ad with the keyword "unemployed" netted an "easy auto financing" ad. And of course, Geambasu reminded us, Gmail is not alone in contradicting its own ad targeting policy. "The web is a complex and opaque ecosystem," she said, "driven by massive collection and monetization of personal data. "

Interestingly, during the last day that the Sunlight project collected its data, Google shut down its Gmail ads program and began something called organic ads.[10] The shutdown of Gmail ads and its transmutation into something else are likely not coincidental. In the process of creating greater transparency, Arvind Narayan's research shows that the very act of tracking the trackers can stop them dead in their tracks. He and his fellow researchers measured the presence of three persistent tracking techniques and published the results in their paper, "The Web Never Forgets: Persistent Tracking Mechanisms in the Wild." Of all the techniques measured, the **canvas fingerprinting** method received the most media attention—and elicited the most public outrage.

The response was immediate. "Canvas fingerprinting was only a known technique for two years," said Narayan, "but in just days after our work was published, the largest users of canvas fingerprinting [AddThis and Ligatus] stopped cold. Why?" Narayan's diagnosis of the sudden halt in the surreptitious fingerprinting technique is that "you have some kind of weird Heisenberg principle analogy for the Web, that measurement automatically changes the behaviour [of the trackers]." In a blog post on his "Freedom to Tinker," one of Narayan's co-authors, Steve Engelhart, posited that the visual representation of canvas fingerprinting in particular was especially helpful in removing the asymmetry of information because users could intuitively grasp that a unique "picture" of their online ID was being created click by click. In addition, the paper included measurements of the trackers responsible for fingerprinting as well as the sites on which the fingerprinting was taking place, thus allowing users to voice their frustrations or visit alternate sites.

Inevitably, though, it was only a matter of time before canvas fingerprinting started up again. When the research team ran the measurements again in January of 2016, the number of distinct trackers utilizing canvas fingerprinting had more than doubled.[11] The lesson: in order to maintain transparency for users, measurements must be collected on a continuous basis. Narayan has initiated what he calls Open Web Privacy Measurement (WPM), which runs monthly taking one million site measurements. He hopes using it will increase transparency, reduce the asymmetry of information in favor of trackers, and answer questions ranging from "Are new web technologies being used and abused for tracking?" to "How is the NSA able to piggyback on ad tracking?" and "How effective are the various tracking protecting agencies?"

However, the tools that Geambasu and Narayan and their teams are building are likely of most use to regulators, privacy advocates and privacy tool developers. As for the general public, while a transparency study generates a burst of interest and outrage, Narayan admits, in "The Web Never Forgets," that the level of user effort and sophistication required for mitigation make it doubtful that individuals can effectively adopt and maintain the necessary privacy tools.

Peering Into The Black Box

Rather than looking towards individuals to mitigate against privacy harms, we can look towards markets and competition to drive the innovations needed to protect privacy. In the industry session, moderated by Professor Bruce Kogut, IBM's Brenda Dietrich, Dstillery's chief data scientist Claudia Perlich, and privacy advocate and RedMorph founder Abhay Edlabadkar each spoke about this approach.

Born under Communism in East Germany, Clau-

[10] http://datascience.columbia.edu/new-tool-expands-tracking-personal-data-web

[11] https://freedom-to-tinker.com/blog/englehardt/retrospective-look-at-canvas-fingerprinting/

Privacy as Promotional Plus

Privacy has increasingly been touted as a selling point by major technology companies. The Fall of 2015 saw the launch of Blackberry PRIV, for instance, which ditched the company's own mobile software for Google's android platform—with a difference—it boasts privacy-enhancing features . Compared to smartphones powered by Apple's iOS, android devices have long been known for their lack of data security. With Blackberry's "Crackberry" days long gone, the company's launch of the PRIV—and its assurance to make its Android system secure—is a last-ditch effort to regain market share.[12] That same fall, both Microsoft and Apple updated and touted their privacy policies. Microsoft re-explained the privacy policy for its Windows 10, assuring users, "We don't use what you say in e-mail, chat, video calls or voice mail, or your documents, photos or other personal files to choose which ads to show you." Apple couldn't resist a barely veiled swipe at Google in promoting its privacy policy, "Some companies mine your cloud data or e-mail for personal information to serve you targeted ads. We don't."[13]

Apple has long used its privacy model to market its phones to consumers. Unlike Facebook, Google, and Amazon, which monetize the data users feed them, Apple does not sell customer data and makes a point of telling consumers that almost all personal information is kept on the device, not sent to Apple's servers. And when it does upload users' data, that data is anonymous, untethered to the users' Apple IDs. For instance, when Apple's iPhone location services do send information on your whereabouts to its own servers, Apple explains in its privacy policy that it uses anonymized rotating identifiers to keep searches and locations from being traced to you personally.[14]

Three weeks after the Innovation and the Value of Privacy Conference, the world saw just how far Apple would go to protect customer data. Apple was besieged by lawsuits from the United States Department of Justice (DOJ), first to create a backdoor key to unlock the iPhone 5C of one of the San Bernardino, California, shooters, Syed Farook, who killed 14 people and injured 22 more on December 22, 2015, and then to unlock the phone of a Brooklyn drug kingpin. The cases never went to court because third-party hackers came forward to unlock the phones. But the case pushed data security and privacy concerns to the forefront for the public. Directly after the DOJ suits, Facebook-owned messaging app, Whatsapp, gained positive press because they announced end-to-end encryption of all chat messages. Yet in all the laudatory press about the encryption, Facebook's Whatsapp still collects all the metadata—your e-mail address, cell phone number and, if not what you say in your chat, with whom you are chatting. Users' comfort with the level of privacy touted by a tech business often comes down to their trust in that particular business.

[12] http://www.cnet.com/news/inside-blackberrys-last-ditch-plan-to-win-you-back-with-android/

[13] http://www.computerworld.com/article/2987249/data-privacy/apple-microsoft-wield-privacy-as-marketing-tool.html

[14] https://www.wired.com/2015/06/apples-latest-selling-point-little-knows/
and [Apple location services] http://www.apple.com/privacy/approach-to-privacy/#personalization

Moderator: Alondra Nelson, Professor and Dean of Social Science at Columbia University

dia Perlich could not recall seeing an ad until she was 16. Now, as chief scientist at marketing intelligence company Dstillery, Perlich fine-tunes consumer marketing by capturing data from consumers' digital footprints. "Predictive modelling has come up because of the downside," Perlich said, referencing Kate Crawford's keynote. However, Perlich argued, the very technology behind the machine learning models provides "a level of security around sharing data and limiting data breeches" and also a way to share data without reaching unintended recipients.

Perlich explained how her models decide who are potential consumers of client products and, hence, who can be shown an ad and when. The machines crunching the number are "agnostic," in that "they don't need to know or understand anything about the meaning of the data… Instead of using things like income, we're directly using the history of URLs to do exactly the same thing, because what you do online is a hell of a lot more predictive than what somebody else thinks you are…. Your actions speak louder than your demographics."

And not only does the machine not "know" anything about your data, but Perlich said she doesn't either. "With all due respect to people arguing for transparency, " she admitted, "I have no clue what these things do. I build ten thousand a week. I have no idea what any of these models predict and why." In an e-mail discussion after the conference, Perlich expanded on this notion:

> While our brain is not capable of calculating the optimal route to the airport (as your GPS can), we can perfectly understand the resulting instructions. This is no longer true when we are talking about algorithms of the ilk used to predict human behavior. The short answer is yes, I know exactly how the prediction was calculated. Think of it as a very complicated recipe.[15] But our human mind cannot possibly understand the meaning behind it…The second problem is that I do not even know what information the model is using because I am trying to respect the privacy of users.[16]

Both Dietrich and Perlich told attendees privacy by design was a crucial part of the models built at IBM and Dstillery respectively. During her talk, Dietrich spoke of her work with IBM Cloud Insight Services and its recent acquisition of The Weather Company's digital assets. According to a press release announcing the deal, IBM's "sophisticated models can now analyze data from three billion weather forecast reference points, more than 40 million smartphones and 50,000 airplane flights per day…[enabling] IBM to offer a broad range of data-driven products and services to more than 5000 clients in the media, aviation, energy, insurance and government industries."[17] Dietrich remarked, "I don't think people have privacy concerns about

[15] https://www.linkedin.com/pulse/way-algorithm-claudia-perlich?trk=mp-reader-card
[16] E-mail from Claudia Perlich to Nancy J. Brandwein dated April 28, 2016

Keynote Speaker: Ashkan Soltani, Former Senior Advisor at the White House Office of Science and Technology Policy and Independent Researcher

knowing what weather conditions are at any point in time and space, but we also have the forecasted weather, and so the question is how do people respond to weather forecasts? How does their behavior change? This is the question that gets into privacy issues."

On the privacy protections in place at IBM, Dietrich said, "We adhere to privacy by design as an objective. Any time we [at IBM] are taking data in or giving data out we get a reading from the chief privacy officer on whether what we're doing is compliant with both the law and the intent and aspirations of the company." In correspondence after the conference Dietrich added, "as Claudia also emphasized, the data is calculated by a computer and almost never looked at by human eyes."

Anonymizing Data and Accounting for Bias

Toward the end of her presentation, Claudia Perlich started to talk about the ways data is anonymized so even if human eyes look at it, the humans looking can't recognize what they are seeing. Expanding upon these techniques after the conference, Perlich elaborated on her point that her company keeps the users' digital footprint from turning into a full body portrait. She explained:

> Our customers are primarily marketers and their agencies are trying to connect to potential customers. They expect that we NOT show ads to bots, that ads do NOT appear on questionable content, that we do NOT fake the evaluation metric (knowingly or unknowingly). On the user side, they expect us NOT to use device fingerprinting to establish a unique identity, to NOT reconnect cookies once they are deleted by the user, to NOT store PII of any kind, to NOT label people's behaviour into segments, to NOT sell user data, and to NOT keep activity logs in humanly readable form. [18]

One of the ways data is parsed, for instance, is through a technique called **feature hashing**. When users with a Dstillery **cookie** visit URLs, the model does not store the actual string of URLs visited but translates them into a "hashed" version, a random string. The URL for www.cnn.com may look like '1he83', but there are maybe 15 other websites other than CNN that get recorded as 1he83, so engineers and scientists like Perlich are unable to trace back the exact user browsing history. In this way, Perlich and Dstillery are attempting to satisfy the privacy dimensions of computer science introduced by Mulligan earlier (see "Privacy Cheat Sheet"), such as anonymity, confidentiality and differential privacy.

Yet because machine learning algorithms learn and evolve based on what users do online, can they really be "agnostic," as Perlich stated? "The amoral status of an algorithm does not negate its effects on society," said Carnegie Mellon researcher Anupan Datta, one of the authors of the Google ads study mentioned earlier.[19] One

[17] http://www.ibmbigdatahub.com/blog/introducing-insight-cloud-services-ibm-insight-2015
[18] Op. Cit.

can't look at the inputs that go into that black box but one can note that the outputs reflect user bias and can potentially lead to discrimination in real life (as in the case with the Google ads study).

Mulligan pressed this point with Perlich, leading to a lively exchange after Mulligan posed this conundrum:

> If I can't look at [inside] the black box, we can at least avoid bad outcomes. We can say 'If these people are alike on all these different parameters we shouldn't treat them differently, just because their race or gender is different.' But that presumes you actually have data about race, about gender. All the [demographic] data you're doing away with in your model actually limits your ability to avoid bad outcomes because you don't actually have data.

To which Perlich responded, "If I give my model race, then I can use it, but that's what you specifically don't want my model to use. Yet, I take it out, and you tell me I can't control it." The model will still effectively identify race by using correlates.

"But can't you set a rule, 'Don't use race as a determining factor'?" asked Mulligan.

"But that's a totally useless rule!" Perlich exclaimed.

"I completely agree," said Mulligan, and then got to the heart of her concern over bias and the black box. "But the thing that you want to be able to say is 'Do these things the model seems to be using correlate with race?', and if you've done away with capturing race you can never do that." Several members of the audience murmured their assent and there was even a chorus of "Yes" that echoed around the room.

A nonplussed Perlich responded, "I am already half out of business, because I refuse showing ads to bots, " she said, referencing Dstillery's privacy policies. "The problem becomes, I'm in an economic incentive system, where I'm paid on getting the highest click-through rate, and if that click-through rate happens to be subject to racial preference… I can quit doing advertising or I will be off the plan and someone will be hired who has much less concerns about going where the signal leads." Perlich clearly stood by Dstillery as a company in which there is a strong attempt to "do the right thing," but as she said earlier, it comes at a cost. "There certainly have been tradeoffs that have not been beneficial to us in terms of market valuation. We have competitors getting a lot further in their economic development by disregarding basic rules around what we should and should not be doing."

Do Businesses Value Privacy? Do Individuals?

"But is privacy something that is always a negative externality, "bad" for business?" Professor Kogut asked all three industry practitioners. "Or can privacy become a benefit for everyone? It seems that deploying privacy solutions would diminish ad revenue, or do you think that privacy can be a positive?" (see "Privacy as Promotional Plus")

In response to Kogut's questions, Abhay Edlabadkar entered a spirited discussion with his fellow panelists. With prior experience in Telecom & Semiconductors including a stint at Bell Labs and a Columbia Business School alumni, Edlabadkar started his company, RedMorph, in 2013 when he became concerned about the amount of data his teens were sharing online. His firm markets a privacy tool that promises to visualize and block trackers, cookies, and third party content, to encrypt your data, and to filter or block inappropriate content. "Advertising is not 'wrong,' affirmed Edlabadkar, but "the way it is happening is wrong with the indiscriminate way data is collected, the asymmetry that has been created."

"My very short answer, " replied Perlich bluntly, "is this: 'You're living in a capitalist society where there is always misalignment of incentives.' "

[19] http://www.nytimes.com/2015/07/10/upshot/when-algorithms-discriminate.html?_r=0; see also https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

But IBM's Dietrich pushed back against the notion of asymmetry, "I don't think there is a disequilibrium here. Some of us don't mind giving information to the businesses that we do business with because it enables them to serve us better. On the other hand, we have some extremely paranoid people who don't want to let any information leak, so they don't reap the necessary benefits."

Edlabadkar leapt from his seat to disagree. "That's not right. I believe one should be able to control what's public and what you want to keep private. That control does not exist right now," he exclaimed, to which Dietrich confirmed calmly, "You're right. However, I'm not convinced that if it *did* exist that most people would be able to control it to their benefit."

The Tradeoff Between Privacy and Utility: Does it Exist?

Dietrich and Perlich assume there is a tradeoff between privacy and utility.

Irfan Saif, the penultimate conference speaker, agreed with this assumption. As a principal in Deloitte & Touche LLP's Cyber Risk Services practice, Saif helps technology companies design and build platforms for the marketplace, with a primary focus on combatting cyber risks.

"I don't believe that users are necessarily victims, " he said, "There's lots of people who intentionally give up certain privacy [rights] and certain information to get something in return…and it's driving a massive economy."

Citing Facebook as one of the forces driving this economy and offering utility to its users, Saif said that in January 2016, "Facebook's user rate was up to 1.6 billion individual users and growing at 100 million users a month…People want to share. They want to share a lot. But what they have, though, is an expectation that the people they are sharing with are going to have a limited and constrained set of use cases for that data. There is a tradeoff that is the whole heart of data as currency…If I can get a game or service or some functionality or something that benefits me from a user experience standpoint, I might be willing to give up some of my privacy."

But what about the question Dietrich raised: if some way to control data actually existed, would people avail themselves of it? The answer is paradoxical. Citing a Pew Research poll on Americans' attitudes on privacy, Edlabadkar reeled off these stats:

- Over 90% of people have a strong desire to have control of their online content and privacy
- 85% of online consumers now oppose tracking and are unwilling to trade their personal data, even anonymously, for the sake of being served relevant ads.

Yet—and this is the paradox—Edlabadkar also pointed out that Americans are not doing much to protect themselves by utilizing tools like RedMorph's or ad blocker Disconnect, despite the fact that such tools create speedier Internet surfing, which is the positive externality Professor Kogut was looking for. Citing his experience trying Disconnect, Irfan Saif said that the tool found twenty trackers on a Superbowl 50 website page—for one two-paragraph article. Saif acknowledged, "The user experience for me is far greater with Disconnect. Not because I'm worried that I won't get targeted ads… I'm a security guy, so I don't click on ads or links, but it speeds up the performance of the web page by more than double."

So if we can protect our privacy and improve our Internet experience to boot, why aren't we doing so? Crawford had mentioned the work of Professor Joseph Turow, at NYU's Annenberg School for Communication. Turow reported the results of his study on the tradeoff between privacy and utility in *The Tradeoff Fallacy: How Marketers are Misrepresenting American Consumers and Opening them Up to Exploitation*. In a study on Americans' attitudes to exchanging online data for supermarket discounts, Turow and his fellow researchers found:

> A majority of Americans are resigned to giving up their data—and that is why many appear to be engaging in tradeoffs. Resig-

Industry Insights Presenters: (from left) Brenda Dietrich, IBM Fellow, Vice President, Data Science, IBM, Abhay Edlabadkar '07, Founder of RedMorph and Claudia Perlich, Chief Data Scientist at Dstillery

> nation occurs when a person believes an undesirable outcome is inevitable and feels powerless to stop it. Rather than feeling able to make choices, Americans believe it is futile to manage what companies can learn about them. Our study reveals that more than half do not want to lose control over their information but also believe this loss has already happened.[20]

For those users who are not resigned, can they opt out? One audience member, an info-tech employee who works on highly secure systems, posed this dilemma to Kate Crawford. "I'm writing a novel set in the war on terrorism. I'm learning Arabic, I'm researching GITMO online. I'm sure I'm in an NSA database. While I don't fit their [terrorism] profile in some ways, I'm sure if I bought a large load of different types of fertilizer, someone would show up at my door very quickly! Are there valid ways of opting out that don't raise a different red flag?" Crawford told the techie/novelist:

> Each piece of information is a signal that can be read a certain way, and the system reads cumulatively… I worry that the 1990s and 2000s model of opting out is gone. We've lost out. There are data infrastructures built into the cities we live in, the airports we fly through. Without the ability to opt out, what can we do to strengthen people's agency, capacity and privacy? We have less than five years to come up with a really strong framework.

How do other countries view the same privacy dilemmas?

Inevitably, questions came up about how the United States compares with other countries when it comes to expectations around privacy and whether there was any convergence in the international space on privacy that would allow business to proceed smoothly.

In the industry session Q&A session, the consensus was that European countries are much more privacy conscious than Americans—Edlabadkar said a higher percentage of Germans used ad blockers. (In the summer of 2016, the European Union adopted a due process requirement for data-driven decisions based "solely on automated processing" that are shown to "significantly affect" citizens.[21] (Writer's note, NJB) Yet while it's tempting to think that privacy rules and regulations necessarily make a populace more secure, Dierdre Mulligan added an interesting perspective from research conducted for her book with Kenneth A. Bamberger, *Privacy on the Ground: Driving Corporate Behavior in the United States and Europe*.[22] The authors found that there is a divergence between how privacy is treated on the ground—in daily and business

[20] https://www.asc.upenn.edu/sites/default/files/TradeoffFallacy_1.pdf

[21] http://www.nytimes.com/2016/08/01/opinion/make-algorithms-accountable.html?_r=1

[22] Bamberger, Kenneth A. and Mulligan, Dierdre K., Privacy on the Ground: Driving Corporate Behavior in the United States and Europe, Boston, MA (MIT Press: 2015).

Keynote Presenter and Moderator: Kate Crawford, Principal Researcher at MSR-NYC (Microsoft) and Visiting Professor, MIT Center for Civic Media

life—and on the books, i.e., the law. Mulligan summarized her findings:

> The countries that have the most robust interpretations of privacy... are the US and Germany. If you just focused on the laws on the books, you would think them the furthest apart. Privacy is a strategic issue, not just a compliance issue; privacy influences business and tech decision-making, market entry, lines of business, not just what does the legal agreement look like at the end of the day. Privacy in both countries is relatively ambiguous from a legal perspective, and the company itself is vested with figuring out what is required. We think of Germany as top down with lots of rules, but the simple fact is they vested responsibility for understanding and protecting privacy with a data protection officer who is in the firm, and as there has been more transparency around privacy failures, those people have become more influential.

These findings are at odds with Crawford's urgency in demanding a framework for privacy within five years and the statement she made, at the outset, in her keynote: If you were to ask me 'How do we end privacy?' The answer: Just keep doing what we're doing."

Pushing Back on Paranoia

On the whole, this conference and the intense debates it spawned on the nexus between data science and privacy—as well as other forums for the debate that have sprouted up—run counter to a pessimistic outlook. What "we're doing"—the academics, industry practitioners and government regulators—is what Crawford lauded as so important at the outset: "putting all our heads together."

From his own experience with enterprises, albeit on the cybersecurity end, Irfan Saif noted just how important this dialogue will be. Regarding the information that is aggregated behind the scenes and correlated in a way that generates information you didn't consciously give up, Saif said there is no consistent approach from company to company. "It's not necessarily a problem…that you want the engineers in the field to solve as they build a product," he said, implying that privacy by design cannot be left solely to data science engineers but "requires a lot of debate and dialogue [across disciplines], and I think that's a lot of what has gone on today."

In introducing the conference's closing speaker Alondra Nelson, Columbia's Dean of Social Science, echoed the need for interdisciplinary dialogue. Nelson said this conference was particularly important to the social science schools "because issues of data and computational social science, privacy, surveillance and discrimination are one of the strategic priorities for the division."

The conference's final keynote speaker, Ashkan Soltani, is an independent technologist and researcher; he was on the team of Pulitzer Prize winning Washington Post reporters who broke the story of Edward Snowden and NSA surveillance, a researcher on the *Wall Street Journal's* award-winning "What They Know" series, and chief technologist for the Federal Trade Commission, where he helped create its new office of technology research and investigation. Wouldn't he sound the most alarms for the potential for privacy breeches in the big data era?

Surprisingly, Soltani began his talk by pushing back on paranoia and attempting to ease privacy fears. "Much of this [discussion] has generated a feeling of being frightened or…of helplessness. That's one response," Soltani said, and he went on to remind attendees, "This field is emerging. We've had databases only less than a century, data brokerages for a few decades, ad targeting for maybe ten or twenty years…Prior to 2006 the big ad companies were not doing behavioral ad targeting. There are a lot of different business models we might end up at. It's not predetermined that all our privacy is gone, and we're walking around naked all day."

At the same time that he pushed back on paranoia, Soltani pushed back on the notion of hiding behind an outdated legal framework of notice and consent, or on the idea Claudia Perlich brought up that "we don't understand how these things [computer models] work." To that, Soltani said, "That's not good enough," and here he echoed Mulligan's concerns about bias: There are certain things the system should *not* do. Using the example of flight control systems safety improvements, he said even seasoned pilots might ask, "What is the system doing now?" in the same way that seasoned data scientists will ask, "What is the model doing now?" The Federal Aviation Administration says, "That may be the case but there are certain things the system should not do and these should be known." In the case of predictive modeling, Soltani said, "Computers or systems don't necessarily understand what data it's processing, but…we as practitioners, developers and designers must make sure that there are certain things baked in." He gave an example of casinos marketing to people who have disclosed gambling addictions online, and concluded, "People's self-disclosed vulnerabilities should not be an area we allow targeting to."

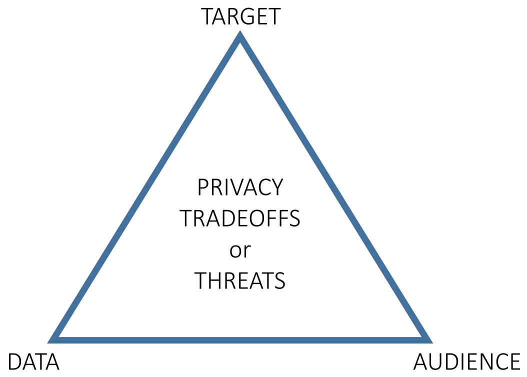

Figure 1. Soltan's Privacy Triangel

Adding to Deirdre Mulligan's take on privacy as a "morass" and an "unwieldy ball of yarn," Soltani agreed that one of the key difficulties of tackling privacy and the ethics surrounding it in the age of innovation in big data is its very amorphousness. Yet no matter what definition you use, Soltani said they all share three elements (see Figure 1):

1. The data: what are the different levels of sensitivity, from browsing history to high sensitivity data like credit card or healthcare information?
2. The target: who are the individuals or class of people from whom data are collected?
3. The audience: who sees/uses the information?

It is only by measuring each of these elements that you can get a sense of privacy tradeoffs or privacy threats, Soltani summarized. He also cautioned that this "privacy triangle" might not work as well with inferential privacy issues that arise when data analytics can take a piece of information—e.g., favorite ice cream flavor—and infer sexual orientation. Still, one might be able to use the triangle as a rubric to analyze a new technology and its effect on privacy.

Soltani urged the attendees to do just this, to approach innovation and solutions from an optimistic, inclusive viewpoint:

> Privacy is information regret management.

> We all want to make sure individuals or classes of society don't regret the disclosures and the changes in information structure to which we trust them. As a result, if we feel like a new change or new development or new policy is going to push us down the spectrum in terms of being more privacy sensitive, we have to make sure we have decided that that's the outcome we want, that we're comfortable with those nuances.

As an example of reducing that link to individuals or reducing the impact of sensitive data floating about in the information ecosystem, Soltani suggested reducing the storage of individualized behaviors once the aggregate inferences are made, thus reducing the potential privacy harm. There are "levers" to play with as a group, he explained, but it's not something that can be left to any one group—to individuals, researchers, businesses or policymakers.

Where Can We Find A Sense of Agency?

As the only person from the regulatory sphere in the conference, Soltani was cautious around the idea of implementing privacy laws. "What I want to emphasize as a policymaker," he warned, "is that we probably don't want technology-specific laws to deal with this [big data and discrimination] too quickly, because oftentimes the people who are least informed about these nuances are policymakers themselves—and that will have downsides for everyone."

What, then, is the "upside"?

At the closing of the conference, Kate Crawford asked Ashkan Soltani, "Where can people find a sense of agency?" If we can't trust policymakers or trust that businesses will behave as ethically as Dstillery or follow the word of their chief policy officers to the letter, what then? If we as consumers find it too complicated to use ad blockers or find that they don't work perfectly yet, how can we protect our information online? And even then, there are always the fumes of what is called "digital exhaust" emanating from our devices, from the applications, and from day-to-day online interactions we make without realizing that traces of our behaviour, or perhaps our very identity, will be left behind like contrails.

One way to find agency is in banding together as concerned citizens and companies. Red-Morpheus's Edlabadkar wrote that his firm is launching a non-profit consortium called the "Ethical Internet Initiative" (EII). "As part of this," he wrote, "websites and advertisers become members and agree not to identify and track users…They submit the URLs through which…context based ads would be pushed. These would then be allowed with user permission in our product or any other…tracker blockers, browsers, etc.…This is an ethical solution as the control is in users' hands." But Edlabadkar says the plan benefits websites as well, since they can regain their ad revenues while branding themselves as ethical and pro-user privacy.[23]

For his part, Soltani said, "The key is articulating the balance we want to maintain, in terms of agency, autonomy, rights, and ethics" that shares the space with a market that provides the commercial benefits and ads supporting the products and services that consumers want. One should think of where we are with privacy, he said, as analogous to where we were decades ago in terms of environmental policy.

It's likely that everyone at the conference was aware of the need to couple concerns for privacy with all the promising innovations in data science. However, the general public is still largely unaware of the issues raised at the conference, in the same way that they were once unaware of the dangers of fossil fuel consumption. Most consumers are now keenly aware of environmental harms and, hopefully, they will become aware of privacy harms in the future. This conversation between academic researchers, business practitioners, and regulators was a fruitful one and part of a larger, national conversation that is far from settled.

[23] E-mail correspondence between Nancy J. Brandwein and Abhay Edlabadkar on April 22, 2016.

Data Science Glossary

Bots: a software application that runs automated tasks (or scripts) over the Internet.

Canvas fingerprinting: a surreptitious online user tracking technique that relies on minute differences in text or images drawn on command by users' browsers.

Cookie: a small amount of data generated by a website and saved by your web browser; its purpose is to remember information about you.

Feature hashing: in machine learning, feature hashing is a fast and space-efficient way to convert arbitrary features into indices in a vector or matrix.

Unstructured data: refers to information that either does not have a pre-defined data model or is not organized in a pre-defined manner; for example, text data is usually unstructured.

Dark Web and Deep Web: the deep web refers to anything that a search engine can't find, and the dark web is a small portion of the deep web that has been intentionally hidden and is inaccessible through standard web browsers.

Differential Privacy: a privacy definition that aims to provide means to maximize the accuracy of queries from statistical databases while minimizing at a cost the chances of identifying personal records by adding a small amount of random noise to the data.

Machine learning: a type of artificial intelligence (AI) that enables machines to learn through inferential algorithms.

Metadata: a set of data that describes other data; metadata for a document will include author, date created, date modified, and file size, e.g.

Privacy by Design (PbD): an approach to protecting privacy by embedding it into the design specifications of technologies, business practices, and physical infrastructures.

Privacy harms: one researcher has defined subjective privacy harm as a feeling that your privacy has been invaded—something that "creeps you out" perhaps—and objective privacy harm as when your own personal information has been used against you. Privacy harms are not the same as privacy violations, which are legal violations, though the line between them is a fine one.

Social graph: a graphic or network representation of the interconnection of relationships in an online social network.

Threat models: refers to the procedure for optimizing network security by identifying and measuring vulnerabilities and then defining countermeasures to prevent, or mitigate the effects of, threats to the system.

Columbia Business School
AT THE VERY CENTER OF BUSINESS™

**The Sanford C. Bernstein & Co. Center
for Leadership and Ethics**
Uris Hall
3022 Broadway
New York, NY 10027
P. 212-854-1060
F. 212-854-0016
www.gsb.columbia.edu/leadership
leadershipethics@gsb.columbia.edu