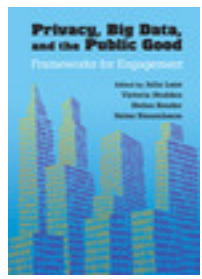

Privacy, Big Data, and the Public Good

Frameworks for Engagement

Edited by Julia Lane, Victoria Stodden, Stefan Bender, Helen Nissenbaum

Book DOI: http://dx.doi.org/10.1017/CBO9781107590205

Online ISBN: 9781107590205

Hardback ISBN: 9781107067356

Paperback ISBN: 9781107637689

Chapter

2 - Big Data's End Run around Anonymity and Consent pp. 44-75

Chapter DOI: http://dx.doi.org/10.1017/CBO9781107590205.004

Cambridge University Press

2 Big Data's End Run around Anonymity and Consent

Solon Barocas and Helen Nissenbaum

Introduction

Big data promises to deliver analytic insights that will add to the stock of scientific and social scientific knowledge, significantly improve decision making in both the public and private sector, and greatly enhance individual self-knowledge and understanding. They have already led to entirely new classes of goods and services, many of which have been embraced enthusiastically by institutions and individuals alike. And yet, where these data commit to record details about human behavior, they have been perceived as a threat to fundamental values, including everything from autonomy, to fairness, justice, due process, property, solidarity, and, perhaps most of all, privacy.[1] Given this apparent conflict, some have taken to calling for outright prohibitions on various big data practices, while others have found good reason to finally throw caution (and privacy) to the wind in the belief that big data will more than compensate for its potential costs. Still others, of course, are searching for a principled stance on privacy that offers the flexibility necessary for these promises to be realized while respecting the important values that privacy promotes.

This is a familiar situation because it rehearses many of the long-standing tensions that have characterized each successive wave of technological innovation over the past half-century and their inevitable disruption of constraints on information flows through which privacy had been assured. It should come as no surprise that attempts to deal with new threats draw from the toolbox assembled to address earlier upheavals. Ready-to-hand, anonymity and informed consent remain the most popular tools for relieving these tensions – tensions that we accept, from the outset, as genuine and, in many cases, acute. Taking as a given that big data implicates important ethical and political values,[2] we direct our focus instead on attempts to avoid or mitigate the conflicts that may arise. We do so because the familiar pair of anonymity and informed consent continues to strike

many as the best and perhaps only way to escape the need to actually resolve these conflicts one way or the other.

Anonymity and informed consent emerged as panaceas because they presented ways to 'have it all'; they would open the data floodgates while ensuring that no one was unexpectedly swept up or away by the deluge. Now, as then, conscientious industry practitioners, policymakers, advocates, and researchers across the disciplines look to anonymity and informed consent as counters to the worrisome aspects of emerging applications of big data. We can see why anonymity and consent are attractive: anonymization seems to take data outside the scope of privacy, as it no longer maps onto identifiable subjects, while allowing information subjects to give or withhold consent maps onto the dominant conception of privacy as control over information about oneself. In practice, however, anonymity and consent have proven elusive, as time and again critics have revealed fundamental problems in implementing both.[3]

The argument that we develop in this chapter goes further. Those committed to anonymity and consent do not deny the practical challenges; their solution is to try harder, to be more creative, to utilize more sophisticated mathematical and statistical techniques, and to become astute to the cognitive and motivational contours of users. Although we accept that improvements can result and have resulted from these efforts (e.g. more digestible privacy policies, more robust guarantees of anonymity, more usable choice architectures, and more supple policy), the transition to big data has turned definitional and practical fault lines that have worried policymakers, pundits, practitioners, and scholars into impassable chasms. After tracing progressive difficulties for anonymity and informed consent, respectively, we reveal virtually intractable challenges to both. In the case of anonymity, where important work has already shown it to be rather elusive, we argue that, even where strong guarantees of anonymity can be achieved, common applications of big data undermine the values that anonymity traditionally had protected. Even when individuals are not 'identifiable', they may still be 'reachable', may still be comprehensibly represented in records that detail their attributes and activities, and may be subject to consequential inferences and predictions taken on that basis. In the case of consent, too, commonly perceived operational challenges have distracted from the ultimate inefficacy of consent as a matter of individual choice and the absurdity of believing that notice and consent can fully specify the terms of interaction between data collector and data subject. Both, we argue, lead to the inescapable conclusion that procedural

approaches cannot replace policies based on substantive moral and political principles that serve specific contextual goals and values.

Definitions and Background Theory

Many of the terms in this chapter have ambiguous and often contested meanings. To avoid disagreements originating in terminological differences, we specify the interpretations of two key terms – big data and privacy – assumed throughout the rest of this chapter. We have reason to believe that these interpretations contribute positively to the substantive clarity, but, for the most part, we set these out as starting assumptions.

Big Data

Taking into consideration wide-ranging uses of 'big data' in public discussions, specialized applications,[4] government initiatives,[5] research agendas,[6] and diverse scientific,[7] critical,[8] and popular publications, we find that the term better reflects a paradigm than a particular technology, method, or practice. There are, of course, characteristic techniques and tools associated with it,[9] but, more than the sum of these parts, big data, the paradigm, is a way of thinking about knowledge through data and a framework for supporting decision making, rationalizing action, and guiding practice.[10] For better or worse, it is challenging entrenched epistemic and decision-making traditions across various domains, from climate science to medicine, from finance to marketing, from resource management to urban planning, and from security to governance.[11] Statistics, computer science, and information technology are crucial enablers and supporters of this paradigm,[12] but the ascent of big data involves, fundamentally, a belief in the power of finely observed patterns, structures, and models drawn inductively from massive datasets.[13]

Privacy as Contextual Integrity

There is some disagreement over how important privacy is among the various ethical and political issues raised by big data.[14] Downplaying privacy, the argument is that *real* problems include how we use the data, whether it is fair to treat people as part of a group, whether data is representative, whether we diminish the range of choices we make about their own lives and fates, whether data about us and the data that we generate belong to us, invoking thereby justice, fairness, autonomy, and property rights.

Revealing these wide-ranging ethical dimensions of big data is important, but an impoverished working conception of privacy can result in the failure to appreciate the crucial ways that these other values and privacy interact.

The conception we adopt here gives privacy a wider berth. To begin, we take privacy to be the requirement that information about people ('personal information') flows appropriately, where appropriateness means in accordance with informational norms. According to the theory of contextual integrity, from which this conception is drawn, informational norms prescribe information flows according to key actors, types of information, and constraints under which flow occurs ('transmission principles'). Key actors include recipients, information subjects, and senders, where the last two are often one and the same. Social contexts form the backdrop for this approach to privacy, accounting for the range over which the parameters of actors, information types, and transmission principles vary. Put more concretely, informational norms for a health care context would govern flow between and about people in their context-specific capacities, such as physicians, patients, nurses, insurance companies, pharmacists, and so forth. Types of information would range over relevant fields, including, say, symptoms, diagnoses, prescriptions, as well as biographical information. And notable among transmission principles, confidentiality is likely to be a prominent constraint on the terms under which information types flow from, say, patients to physicians. In drawing comparisons between contextual integrity and other theories of privacy, one key difference is that control over information about oneself is merely one in an indefinitely large class of transmission principles, not presumed unless the other parameters – (context specific) actors and information types – warrant it.[15]

Contextual informational norms, like other social norms, generally, are not fixed and static, but may shift, fade, evolve, and even reverse at varying rates, slowly or suddenly, sometimes due to deliberate cultural, legal, and societal alterations and other times in response to contingencies beyond human or societal control. Science and technology is a significant agent of change; in particular, computing and information technologies have been radically disruptive, enabling information practices that frequently diverge from entrenched informational norms. To explain why such disruptions are morally problematic – or rather to distinguish between those that are and are not – a norm-based account of privacy, such as contextual integrity, must offer a basis for drawing such distinctions. This enables a systematic critical perspective on informational norms in flux. For the theory of contextual integrity, the touchstones of moral legitimacy include interests

and general moral and political values (and associated rights), commonly cited in accounts of privacy. Beyond these, however, a further distinctive set of considerations are context-specific ends, purposes, and values. Although this is not the place to elaborate in detail, consider as a quick illustration the rules limiting access to results of an HIV test. Generally, we might consider embarrassment, job security, danger to sexual partners, autonomy, various freedoms, and so on. Beyond these, however, contextual integrity further considers how the shape of access rules may affect whether people choose to undergo testing at all. As such, access rules could influence how effectively the purposes and values of the health care context are achieved. Ideal norms, therefore, are those that promote relevant ends, purposes, and values. And since the world is a messy place, rife with conflict and uncertainty, it is usually on the basis of partial knowledge only that we seek to optimize on these factors. In concrete circumstances where science and technology enable disruptions of entrenched norms, a heuristic supported by contextual integrity sets entrenched norms as default but allows that if novel practices are more effective in promoting interests, general moral and political values, and context-specific ends, purposes, and values, they should be favored over the status quo.

Now we are ready to weave together the disparate threads thus far spun. Big data involves practices that have radically disrupted entrenched information flows. From modes of acquiring to aggregation, analysis, and application, these disruptions affect actors, information types, and transmission principles. Accordingly, privacy, understood as contextual integrity, is fundamentally part of the big data story for it immediately alerts us to the ways any practice conflicts with the expectations we may have based on entrenched information-flow norms. But that is merely the beginning. Evaluating disruptive practices means judging whether they move us closer or farther from ideal informational flows, that is, whether they are more or less effective in promoting interests, general moral and political values, and context-specific ends, purposes, and values. In other words, we proceed from observing disruptive flows to assessing their comparative impacts on ethical and political values, such as fairness, justice, freedom, autonomy, welfare, and others more specific to the context in question. Take, for example, an applicant who is denied admission to college based on predictive analytics performed on a dataset aggregated from diverse sources, including many that have not traditionally featured into admissions decisions. Imagine further that these additional sources allowed the college to discriminate – perhaps unwittingly – against applicants on the basis of criteria that happen to correlate with socioeconomic status and thus with

the likely need for financial aid.[16] While the outcome of such decisions may be judged unfair for many reasons worth discussing, it is the role of privacy – the role of disruptive informational flow – that we wish to note in this case.

Why, one may ask, insist on the centrality of privacy? First, doing so deepens our understanding of privacy and its instrumental value and at the same time highlights the distinctive ways that other ethical values are impinged and sustained, specifically, by the ways information does and does not flow. Privacy is important, in part, because it implicates these other values. Second, doing so also allows us to better formulate interventions, regulations, or remediation for the sake of these values. By keeping in view connections with specific information flows, certain options become salient that might otherwise not have been. Parsing cases in which big data gives rise to discrimination in terms of contextual integrity forces us to be much more specific about the source of that unfairness because it compels us to account for the disruption that made such discrimination possible.[17] And it likewise allows us to ask if anonymity and informed consent limit or mitigate the potential consequences of such disruptions – that is, whether they actually protect the values at stake when novel applications of big data (threaten to) violate contextual integrity.

Anonymity

Anonymity obliterates the link between data and a specific person not so much to protect privacy but, in a sense, to bypass it entirely.[18] Anonymity is an attractive solution to challenges big data poses to privacy when identities associated with information in a dataset are not necessary for the analysis to proceed. For those in search of group-level regularities, anonymity may allow for relatively unfettered access to databases. The greatest consensus around the utility of anonymization seems to have emerged in the sciences, including medicine, public and population health, urban planning, and education, to name a few, with exciting prospects for advancing knowledge, diminishing risk, and improving decision making.[19] But incumbents in many other sectors have begun to stake out this moral high ground by claiming that their analytics apply only to anonymized datasets, particularly those in marketing and other commercial sectors.[20]

As we well know, however, anonymity is not unassailable. One of the earliest public demonstrations of its limits came with AOL's release of a large set of anonymized search queries with the stated purpose of facilitating academic research. This well-intended act backfired when a pair of

enterprising news reporters identified a number of individuals based on the content of searches.[21] Following these revelations, efforts to anonymize search query data, which were not particularly persuasive,[22] have more or less fizzled out. The promise of anonymization was further chipped away by rigorous demonstrations by Sweeney, joint work by Narayanan and Shmatikov, and ongoing efforts by Dwork,[23] with implications further drawn by Ohm and others in areas of law and policy, where debates rage on.[24]

It is impossible, within the scope of this article, to render anything close to a thorough account of the contemporary debate around anonymity; we merely mention key positions on threats to anonymity and attempts to defend it that are relevant to the general argument that we wish to develop. According to the literature, the promise of anonymity is impossible to fulfill if individual records happen to contain information – information that falls outside the scope of the commonly defined set of personally identifiable information – that nevertheless uniquely distinguishes a person enough to associate those records to a specific individual. So-called 'vanity searches' are an obvious example of this problem,[25] as AOL discovered,[26] but so, too, are records that contain extremely rich (e.g. location) data that necessarily map onto specific individuals.[27] The literature has also demonstrated many less obvious ways in which anonymity cannot be guaranteed due to the threat of so-called re-identification attacks.[28] These attacks depend on a variety of methods: overlaying an anonymized dataset with a separate dataset that includes identifying information, looking for areas of overlap (commonly described as a linkage attack)[29] or performing a sequence of queries on an anonymized dataset that allow the attacker to deduce that a specific person must be in the dataset because only one person has *all* of the queried attributes (differencing attack).[30] Responding to these challenges, computer scientists have developed a number of approaches to limit, if not eliminate, the chances of deducing identity, such as k-anonymity[31] and differential privacy,[32] which work in certain settings by abstracting or perturbing data to a level or degree set by data controllers. At the time of writing, this area of research is burgeoning, even though few real-world applications have been successfully implemented.

Let us review the main threads of this argument: anonymity is an attractive solution to challenges big data poses to privacy when identities associated with information in a dataset are not necessary for the analysis to proceed. Scientific and policy debates have swirled around whether robust anonymization is possible and whether the impact of intractable challenges is a fringe phenomenon of little practical importance (and thus merely of

academic interest) or fatal to the entire enterprise. *The concerns we have are neither about whether anonymization is possible nor about how serious a problem it poses for practical purposes; they are whether, in the first place, anonymization addresses privacy and related ethical issues of big data.* In so saying, we wish to shift the locus of attention away from the usual debates – conceding, at the same time, that they are extremely important and significant – to a different set of questions, where, for the sake of argument, we assume that the problem of anonymization, classically speaking, has been solved.

In order to see why anonymity does not solve ethical problems relating to privacy in a big data age, we should ask why we believe it does. And to do that, we need to ask not only whether in this age we are able to preserve the present-day equivalent of a traditional understanding of anonymity as namelessness, but whether this equivalent preserves what is at stake in protecting anonymity. In short, we need to ask whether it is worthwhile to protect whatever is being protected when, today, we turn to anonymity to avoid the ethical concerns raised by the big data paradigm.

Scholarship, judicial opinions, and legislative arguments have articulated the importance of anonymity in preserving and promoting liberal democratic values. We summarized these in earlier work, where we wrote that anonymity

> offers a safe way for people to act, transact, and participate without accountability, without others 'getting at' them, tracking them down, or even punishing them. [As such, it] may encourage freedom of thought and expression by promising a possibility to express opinions, and develop arguments, about positions that for fear of reprisal or ridicule they would not or dare not do otherwise. Anonymity may enable people to reach out for help, especially for socially stigmatized problems like domestic violence, fear of HIV or other sexually transmitted infection, emotional problems, suicidal thoughts. It offers the possibility of a protective cloak for children, enabling them to engage in internet communication without fear of social predation or – perhaps less ominous but nevertheless unwanted – overtures from commercial marketers. Anonymity may also provide respite to adults from commercial and other solicitations. It supports socially valuable institutions like peer review, whistle-blowing and voting.[33]

In this work, we argued that the value of anonymity inheres not in namelessness, and not even in the extension of the previous value of namelessness to all uniquely identifying information, but instead to something we called 'reachability', the possibility of knocking on your door, hauling you out of bed, calling your phone number, threatening you with sanction, holding you accountable – with or without access to identifying information.[34]

These are problematic because they may curtail basic ethical and political rights and liberties. But also at stake are contextual ends and values such as intellectual exploration, wholehearted engagement in social and economic life, social trust, and the like. The big data paradigm raises the stakes even further (to a point anonymity simply cannot extend and the concept of reachability did not locate) for a number of related reasons.

'Anonymous Identifiers'

First and perhaps foremost, many of anonymity's proponents have different meanings in mind, few of which describe practices that achieve unreachability. For example, when commercial actors claim that they only maintain anonymous records, they do not mean that they have no way to distinguish a specific person – or his browser, computer, network equipment, or phone – from others. Nor do they mean that they have no way to recognize him as the same person with whom they have interacted previously. They simply mean that they rely on unique persistent identifiers that differ from those in common and everyday use (i.e. a name and other so-called personally identifiable information (PII)). Hence the seemingly oxymoronic notion of an 'anonymous identifier', the description offered by, among others, Google for its forthcoming AdID,[35] an alternative to the cookie-based tracking essential for targeted advertising.[36] If its very purpose is to enable Google to identify (i.e. recognize) the same person on an ongoing basis, to associate observed behaviors with the record assigned to that person, and to tailor its content and services accordingly, AdID is anonymous only insofar as it does not depend on traditional categories of identity (i.e. names and other PII). As such, the identifier on offer does nothing to alleviate worries individuals might have in the universe of applications that rely on it. This understanding of anonymity instead assumes that the real – and only – issue at stake is how easily the records legitimately amassed by one institution can be associated with those held by *other* institutions, namely an association that would reveal the person's legal or real-world identity.[37]

The reasons for adopting this peculiar perspective on anonymity becomes clear when we explore why names, in particular, tend to generate such anxiety. As a persistent and common identifier, names have long seemed uniquely worrisome because they hold the potential to act as an obvious basis for seeking out *additional* information that refers to the same person by allowing institutions to match records keyed to the same name. Indeed, this is the very business of commercial data brokers: "Acxiom and other database marketing companies sell services that let retailers simply

type in a customer's name and zip code and append all the additional profile information that retailers might want".[38] But this is highly misleading because, as scholars have long argued, a given name and address is just one of many possible ways to recognize and associate data with a specific person.[39] Indeed, *any* unique identifier or sufficiently unique pattern can serve as the basis for recognizing the same person in and across multiple databases.[40]

The history of the Social Security Number is highly instructive here: as a unique number assigned to all citizens, the number served as a convenient identifier that *other* institutions could adopt for their own administrative purposes. Indeed, large institutions were often attracted to the Social Security Number because it was necessarily more unique than given names, the more common of which (e.g. John Smith) could easily recur multiple times in the same database. The fact that people had existing reasons to commit this number to memory also explains why other institutions would seize upon it. In so doing, however, these institutions turned the Social Security Number, issued by the government for administering its own welfare programs, into a *common* unique identifier that applied across multiple silos of information. A Social Security Number is now perceived as sensitive, not because of any quality inherent to the number itself, but rather because it serves as one of the few common unique identifiers that enable the straightforward matching of the disparate and detailed records held by many important institutions.

The history of the Social Security Number makes clear that any random string that acts as a unique persistent identifier should be understood as a pseudonym rather than an 'anonymous identifier',[41] that pseudonyms place no inherent restrictions on the matching of records, and that the protective value of pseudonyms decreases as they are adopted by or shared with additional institutions.[42] This is evident in the more recent and rather elaborate process that Facebook has adopted to facilitate the matching of its records with those maintained by outside advertisers while ensuring the putative anonymity of the people to whom those records refer:

> A website uses a formula to turn its users' email addresses into jumbled strings of numbers and letters. An advertiser does the same with its customer email lists. Both then send their jumbled lists to a third company that looks for matches. When two match, the website can show an ad targeted to a specific person, but no real email addresses changed hands.[43]

While there might be some merit to the argument, advanced by a representative of the Interactive Advertising Bureau, that such methods demonstrate that online marketers are not in the business of trying "to get people's

names and hound them",[44] they certainly fall short of any common understanding of the value of anonymity. They place no inherent limits on an institution's ability to recognize the same person in subsequent encounters, to associate, amass, and aggregate facts on that basis, and to draw on these facts in choosing if and how to act on that person. The question is whether, in the big data era, this still constitutes a meaningful form of unreachability.

Comprehensiveness

A further worry is that the comprehensiveness of the records maintained by especially large institutions – records that contain no identifying information – may become so rich that they subvert the very meaning of anonymity.[45] Turow, for instance, has asked, "[i]f a company knows 100 data points about me in the digital environment, and that affects how that company treats me in the digital world, what's the difference if they know my name or not?"[46] The answer from industry is that it seems to matter very little indeed: "The beauty of what we do is we don't know who you are [. . .] We don't want to know anybody's name. We don't want to know anything recognizable about them. All we want to do is [. . .] have these attributes associated with them."[47] This better accounts for the common refrain that companies have no particular interest in who someone is because their ability to tailor their offerings and services to individuals is in no way limited by the absence of such information. And it helps to explain the otherwise bizarre statement by Facebook's Chief Privacy Officer that they "serve ads to you based on your identity [. . .] but that doesn't mean you're identifiable."[48] On this account, your legal or real-world identity is of no significance. What matters are the properties and behaviors that your identity comprises – the kinds of details that can be associated with a pseudonym assigned to you without revealing your actual identity. Where these details are sufficiently extensive, as is the case with platforms that deal in big data, and where all of these details can be brought to bear in deciding how to treat people, the protections offered by 'anonymity' or 'pseudonymity' may amount to very little.[49] They may enable holders of large datasets to act on individuals, under the cover of anonymity, in precisely the ways anonymity has long promised to defend against. And to the extent that results in differential treatment that limits available choices and interferes with identity construction, it threatens individual autonomy and social justice. For these reasons, Serge Gutwirth and Paul Hert have warned that if it is "possible to control and steer individuals without the need to identify them, the time has probably come to explore the possibility of

a shift from personal data protection to data protection tout court."[50] In other words, we can no longer turn to anonymity (or, more accurately, pseudonymity) to pull datasets outside the remit of privacy regulations and debate.

Inference

But even this fails to appreciate the novel ways in which big data may subvert the promise of such protections: inference. However troubling the various demonstrations by computer scientists about the challenge of ensuring anonymity, there is perhaps more to fear in the expanding range of facts that institutions can infer and upon which they have become increasingly willing to act. As Brian Dalessandro has explained, "a lot can be predicted about a person's actions without knowing anything personal about them."[51] This is a subtle but crucially important point: insights drawn from big data can furnish additional facts about an individual (in excess of those that reside in the database) without any knowledge of their specific identity or any identifying information. Data mining breaks the basic intuition that identity is the greatest source of potential harm because it substitutes inference for using identifying information as a bridge to get at additional facts. Rather than matching records keyed to the same name (or other PII) in different datasets, data mining derives insights that simply allow firms to guess at these qualities instead. In fact, data mining opens people up to entirely new kinds of assessments because it can extend the range of inferable qualities far beyond whatever information happens to reside in records elsewhere. And as Dalessandro again explains, firms that adopt these tactics may submit to few, if any, constraints, because "PII isn't really that useful for a lot of predictive modeling tasks."[52] This explains a recent anecdote relayed by Hardy: "Some years ago an engineer at Google told me why Google wasn't collecting information linked to people's names. 'We don't want the name. The name is noise.' There was enough information in Google's large database of search queries, location, and online behavior, he said, that you could tell a lot about somebody through indirect means."[53] These indirect means may allow data collectors to draw inferences about precisely those qualities that have long seemed unknowable in the absence of identifying information. Rather than attempt to de-anonymize medical records, for instance, an attacker (or commercial actor) might instead infer a rule that relates a string of more easily observable or accessible indicators to a specific medical condition,[54] rendering large populations vulnerable to such inferences even in the absence of PII. Ironically, this is often the very thing about big data that generates the most excitement: the capacity to

detect subtle correlations and draw actionable inferences. But it is this very same feature that renders the traditional protections afforded by anonymity (again, more accurately, pseudonymity) much less effective.

Research Underwritten by Anonymity

The very robustness of the new guarantees of anonymity promised by emerging scholarship may have perverse effects if findings from the research that they underwrite provide institutions with new paths by which to infer precisely those attributes that were previously impossible to associate with specific individuals in the absence of identifying information. Ironically, this is the very purpose of differential privacy, which attempts to permit useful analysis of datasets while providing research subjects with certain guarantees of anonymity.[55] *However much these protect volunteers, such techniques may license research studies that result in findings that non-volunteers perceive as menacing because they make certain facts newly inferable that anonymity once promised to keep beyond reach.*

A recent study demonstrating that students suffering from depression could be identified by their Internet traffic patterns alone was met with such a reaction.[56] Much of this seemed to stem from one of the applications that the researchers envisioned for their results: "[p]roactively discovering depressive symptoms from passive and unobtrusive Internet usage monitoring."[57] The study is noteworthy for our purposes for having taken a number of steps to ensure the anonymity and privacy of its research subjects while simultaneously – if unintentionally – demonstrating the limits of those very same protections for anyone who might be subject to the resulting model. The point is not to pick on these or other academic researchers; rather, it is to show that anonymity is not an escape from the ethical debates that researchers should be having about their obligations not only to their data subjects, but also to others who might be affected by their studies for precisely the reasons they have chosen to anonymize their data subjects.

Informed Consent

Informed consent is believed to be an effective means of respecting individuals as autonomous decision makers with rights of self-determination, including rights to make choices, take or avoid risks, express preferences, and, perhaps most importantly, resist exploitation. Of course, the act of consenting, by itself, does not protect and support autonomy; individuals

must first understand how their assent plays out in terms of specific commitments, beliefs, needs, goals, and desires. Thus, where anonymity is unachievable or simply does not make sense, informed consent often is the mechanism sought out by conscientious collectors and users of personal information.

Understood as a crucial mechanism for ensuring privacy, informed consent is a natural corollary of the idea that privacy means control over information about oneself. For some, these are the roots of privacy that must be respected in all environments and against all threats. Its central place in the regulation of privacy, however, was solidified with the articulation and spread of the Fair Information Practice Principles (FIPPs) in the domains of privacy law and countless data protection and privacy regulation schemes around the world. These principles, in broad brushstrokes, demand that data subjects be given notice, that is to say, informed who is collecting, what is being collected, how information is being used and shared, and whether information collection is voluntary or required.[58]

The Internet challenged the 'level playing field' embodied in FIPPS.[59] It opened unprecedented modalities for collecting, disseminating, and using personal information, serving and inspiring a diverse array of interests. Mobile devices, location-based services, the Internet of things, and ubiquitous sensors have expanded the scope even more. For many, the need to protect privacy meant and continues to mean finding a way to support notice and choice without bringing this vibrant ecology to a grinding halt. This need has long been answered by online privacy policies offered to individuals as unilateral terms-of-service contracts (often dubbed 'transparency and choice' or 'notice and consent'). In so doing, privacy questions have been turned into practical matters of implementation. As in the arena of human subjects research, the practical challenge has been how to design protocols for embedding informed consent into interactions of data subjects and research subjects with online actors and researchers, respectively. In both cases, the challenge is to come up with protocols that appropriately model both notice and consent. What has emerged online are privacy policies similar to those already practiced in hard copy by actors in the financial sector, following the Gramm-Leach-Bliley privacy rules.[60]

Over the course of roughly a decade and a half, privacy policies have remained the linchpin of privacy protection online, despite overwhelming evidence that most of us neither read nor understand them.[61] Sensitive to this reality, regulatory agencies, such as the Federal Trade Commission, have demanded improvements focusing attention on (1) ways privacy policies are expressed and communicated so that they furnish more effective

notice and (2) mechanisms that more meaningfully model consent, reviving the never-ending stalemate over opt-in versus opt-out.[62] While the idea that informed consent *itself* may no longer be a match for challenges posed by big data has been floated by scholars, practitioners, advocates, and even some regulators,[63] such thinking has not entered the mainstream. As before, the challenge continues to be perceived as purely operational, as a more urgent need for new and inventive approaches to informing and consenting that truly map onto the states of understanding and assenting that give moral legitimacy to the practices in question.

In this chapter, we take a different path. We accept that informed consent is a useful privacy measure in certain circumstances and against certain threats and that existing mechanisms can and should be improved, but, against the challenges of big data, consent, by itself, has little traction. After briefly reviewing some of the better-known challenges to existing models of informed consent, we explore those we consider insurmountable.

The Transparency Paradox

There is little value in a protocol for informed consent that does not meaningfully model choice and, in turn, autonomy. The ideal offers data or human subjects true freedom of choice based on a sound and sufficient understanding of what the choice entails. Community best practices provide standards that best approximate the ideal, which, because only an approximation, remains a subject of philosophical and practical debate.[64] Online tracking has been one such highly contentious debate[65] – one in which corporate actors have glommed onto the idea of plain language, simple-to-understand privacy policies, and plain-to-see boxes where people can indicate their assent or consent. A number of scholars continue to hold out hopes for this approach,[66] as do regulators, such as the FTC, who continues to issue guiding principles that reflect such commitments.[67] But situations involving complex data flows and diverse institutional structures representing disparate interests are likely to confront a challenge we have called 'the transparency paradox',[68] meaning that simplicity and clarity unavoidably results in losses of fidelity. Typical of the big data age is the business of targeted advertising, with its complex ecology of back-end ad networks and their many and diverse adjuncts. For individuals to make considered decisions about privacy in this environment, they need to be informed about the types of information being collected, with whom it is shared, under what constraints, and for what purposes. Anything less

than this requires a leap of faith. Simplified, plain-language notices cannot provide information that people need to make such decisions. The detail that would allow for this would overwhelm even savvy users because the practices themselves are volatile and indeterminate as new parties come on board and new practices, squeezing out more value from other sources of information (e.g. social graphs), are constantly augmenting existing flows. Empirical evidence is incontrovertible: the very few people who read privacy policies do not understand them.[69] But the paradox identified above suggests that even when people understand the text of plain-language notices, they still will not – indeed cannot – be informed in ways relevant to their decisions whether to consent.

Indeterminate, Unending, Unpredictable

What we have said, thus far, emerges from a discussion of notice and choice applied to online behavioral advertising, but with clear parallels for the big data paradigm generally. Consider typical points of contact for data gathering: signing up for a smart utility meter, joining an online social network, joining a frequent flier program, buying goods and services, enrolling in a MOOC, enrolling in a health self-tracking program, traveling, participating in a medical trial, signing up for a supermarket loyalty card, clicking on an online ad, commenting on a book, a movie, or a product, applying for insurance, a job, a rental apartment, or a credit card. Because these mundane activities may yield raw material for subsequent analysis, they offer a potential juncture for obtaining consent, raising the natural question of how to describe information practices in ways that are relevant to privacy so that individuals meaningfully grant or withhold consent. The machinations of big data make this difficult because data moves from place to place and recipient to recipient in unpredictable ways. Further, because its value is not always recognized at collection time, it is difficult to predict how much it will travel, how much it will be in demand, and whether and how much it may be worth. In the language of contextual integrity, unless recipients and transmission principles are specified, the requirements of big data are for a blank check.

While questions of information type and use might, at first, seem straightforward, they are extremely difficult when considered in detail: it may be reasonably easy for a utility company to explain to customers that, with smart meters, it can monitor usage at a fine grain, can derive aggregate patterns within and across customers, and can use these as a basis for important decisions about allocation of resources and for

targeted advisement about individual customers' energy usage. It may clearly explain who will be receiving what information and to what end. With notice such as this, consent is meaningful. However, big data analytics typically do not stop here; an enterprising company may attempt to figure out how many people are associated with a given account, what appliances they own, their routines (work, bedtime, and vacations). It may fold other information associated with the account into the analysis and other information beyond the account – personal or environmental, such as weather. The company may extract further value from the information by collaborating with third parties to introduce further data fields. Not anomalous, practices such as these are the life blood of the big data enterprise for massive corporate data brokers and federal, state, and local government actors. How can they be represented to data subjects as the basis for meaningful consent?

Let us consider the challenges. The chain of senders and recipients is mazelike and potentially indefinite, incorporating institutions whose roles and responsibilities are not circumscribed or well understood. The constraints under which handoffs take place are equally obscure, including payments, reciprocity, obligation, and more. What can it mean to an ordinary person that the information will be shared with Axciom or Choicepoint, let alone the NSA? Characterizing the type of information is even tougher. Is it sufficient for the utility company to inform customers that it is collecting smart meter readings? The case is strong for arguing that notice should cover not only this information but, further, information that can be directly derived from it and even information that more sophisticated analysis might yield, including that which follows from aggregations of smart meter readings with information about other matters, personal or contextual. Intuitions on this matter are challenging, almost by definition, because the value of big data lies in the unexpectedness of the insights that it can reveal.

Even if we knew what it meant to provide adequate notice to ensure meaningful consent, we would still not have confronted the deepest challenges. One is the possibility of detecting surprising regularities across an entire dataset that reveal actionable correlations defying intuition and even understanding. With the best of intentions, holders of large datasets willing to submit them to analyses unguided by explicit hypotheses may discover correlations that they had not sought in advance or anticipated. A lot hangs on what informed consent means in such cases.[70] Does the data controller's obligation end with informing subjects about data that is explicitly recorded, or must the data controller adopt a more encompassing approach,

explaining what further information the institution may be able to glean?[71] If the more encompassing approach is taken, how does the data controller explain that it is impossible to know in advance what further information might be discoverable? These factors diminish the value of informed consent because they seem to require notice that does not delimit future uses of data and the possible consequences of such uses. As many have now argued, consent under those conditions is not meaningful.[72]

The Tyranny of the Minority

But big data troubles the long-standing focus on individual choice in a slightly more roundabout way because, as discussed earlier, the willingness of a few individuals to disclose certain information implicates everyone else who happens to share the more easily observable traits that correlate with the revealed trait. This is the tyranny of the minority: the volunteered information of the few can unlock the same information about the many. This differs markedly from the suggestion that individuals are ill equipped to make choices that serve their actual interests; rather, even if we accept that individuals can make informed, rational decisions concerning their own privacy, these decisions nonetheless affect what institutions (to whom these individuals have disclosed information) can now know (i.e. infer) about others.[73]

Such inferences can be drawn in a number of ways. In registering some kind of connection to another person through the formal process of 'friending' on a social networking site, we signal that this is a person with whom we share certain interests, affinities, and history. In associating with this person, we open ourselves up to inferences that peg us as people who share certain qualities with this other person. This is the familiar trope about 'the company I keep': what my friends say and do – or rather, what they are willing to say and do on social networking sites – will affect what others think of me. Hence danah boyd's point that "[i]t's no longer about what you do that will go down on your permanent record. Everything that everyone else does that concerns you, implicates you, or might influence you will go down on your permanent record."[74]

Computer scientists have turned this into a formal problem, asking whether techniques drawing from social network analysis and data mining can be used to infer undisclosed attributes of a user based on the disclosed attributes of the user's friends on social networking sites. And indeed a recent study has demonstrated that, where a certain portion of their friends disclose such facts, social networking sites may be able to infer

users' undisclosed major, graduation year, and dorm.[75] Other – more widely reported – research has also shown that homosexuality can be inferred with some reliability from the fact that a user holds a number of relationships and interacts with an otherwise disproportionate number of 'out' users.[76] Yet another study, building on this earlier work, has even shown that it is possible to make inferences about people who are not even a part of an online social network (i.e. to learn things about obviously absent *non*members).[77]

These demonstrations have tended to focus on cases of explicit association and the drawing of inferences based on confirmed relations, but, when we move away from discussions of online social networking, we find that no such explicit associations are necessary to engage in this same kind of guesswork. More significantly, similar inferences can be made about an entire population even if only a small fraction of people who share no ties are willing to disclose. This describes the dynamics of the Target pregnancy prediction score.[78] In this case, Target did not infer the likelihood of a woman giving birth by looking at her group of friends; rather, the company looked over the records from its baby shower registry to find women who had actively disclosed the fact that they had given birth and then went about trying to figure out if these women's shopping habits, leading up to the baby shower, seemed to differ from other customers' habits such that Target could then recognize the telltale signs in the future shopping habits of *other* women.[79] Which is to say that Target was able to infer a rule about the relationship between purchases and pregnancy from what must have been a tiny proportion of all its customers who actually decided to tell the company that they recently had a baby. Not only is this the tyranny of the minority, it is a choice forced upon the majority by a minority with whom they have no meaningful or recognized relations.[80]

Computer science researchers are tackling this question head-on: what proportion of people need to disclose that they possess a certain attribute for an adversary to then be able to identify all the other members in the population who also have this attribute? The findings from Mislove et al.'s study are rather startling: "multiple attributes can be inferred globally when as few as 20% of the users reveal their attribute information."[81] Of course, reaching this minimum threshold is really just a matter of arriving at a sufficiently representative sample whose analysis generates findings that are generalizable to an entire population. As such, the value of any particular individual's withheld consent diminishes incrementally the closer the dataset of those who granted consent approaches representativeness – a

point beyond which companies may have no further reason to pass. So long as a data collector can overcome sampling bias with a relatively small proportion of the consenting population,[82] this minority will determine the range of what can be inferred for the majority and it will discourage firms from investing their resources in procedures that help garner the willing consent of more than the bare minimum number of people. In other words, once a critical threshold has been reached, data collectors can rely on more easily observable information to situate all individuals according to these patterns, rendering irrelevant whether or not those individuals have consented to allowing access to the critical information in question. Withholding consent will make no difference to how they are treated!

Conclusion

Those swept up in the great excitement that has placed big data at the forefront of research investment and the national scientific policy agenda may take courage. For them, these findings, particularly those concerning consent, prove once and for all that privacy is an unsustainable constraint if we are to benefit, truly, from big data. Privacy and big data are simply incompatible and the time has come to reconfigure choices that we made decades ago to enforce certain constraints. The arguments presented here give further reason to dislodge privacy from its pedestal and allow the glorious potential of big data to be fulfilled.[83] We think these people are wrong in part because they adhere to a mistaken conception of privacy, often as control or as secrecy. Because they see privacy at odds with any distribution and use of data instead of focusing only on the inappropriate, they set up a false conflict from the start. They also may wrongly be conflating the *operationalization* of informed consent with informed consent *itself*.

Others say that we should remain concerned about ethical issues raised by big data, that, while privacy may be a lost cause, the real problems arise with use.[84] Those deserving urgent attention include unfair discrimination, being limited in one's life choices, being trapped inside stereotypes, being unable to delineate personal boundaries, being wrongly judged, embarrassed, or harassed.[85] Pursuing privacy as a way to address these issues is not only retrograde but a fool's errand, a conclusion reinforced by the arguments in our paper. Better, they would say, to route around privacy and pursue directly its ends. We agree that individual interests and ethical and, we would add, context-specific values are vitally important, but we think that it is reckless to sever, prematurely, the conceptual and practical

ties between privacy and these moral and political ends. To fathom the ways that big data may threaten interests and values, we must distinguish among the origins and nature of threats to individual and social integrity, between, say, unfair discrimination originating in inappropriate information flows and unfair discrimination originating from other causes. For one thing, different sources may indicate different solutions.

We are not yet ready to give up on privacy, nor completely on anonymity and consent. The paradigm shift of big data calls for a paradigm shift in our responses and, though it may seem that the arguments of this chapter leave no place for anonymity and consent and, for some, therefore, no place for privacy, we reach different conclusions.

Let us begin with informed consent and imagine it foregrounded against a social landscape. In academic and regulatory circles, attention has focused on the foreground, suggesting ways to shape, tweak, and augment informed consent so that it covers everything important about the relationship between a data controller and a data subject. FIPPS and its innumerable descendants are a case in point. These efforts ensure that, in principle, nothing should go unremarked, unrevealed, unnoticed; in practice, informed consent has groaned under the weight of this burden with results – such as the transparency paradox – that have been noted here and elsewhere.

Informed consent also has a great legacy in the domain of human subjects research, where it remains the subject of ongoing deliberation, and has generated a mature philosophical literature. In *Rethinking Informed Consent in Bioethics*, philosophers Neil Manson and Onora O'Neill address a concern, analogous to the one confronted by privacy researchers and regulators, over how to communicate with human subjects to ensure that consent is meaningful. They observe that the transaction of informed consent in medical treatment and biomedical research can only be understood against a rich social backdrop, which integrates medical practice and research into the background fabric of social and political life. When individuals – human subjects – enter into a study or treatment regime, they engage not as tabula rasa in a vacuum expecting that the protocol of informed consent will specify fully what will happen and respective rights, obligations, and responsibilities. It does not and cannot constitute the complete relationship between the medical researcher or practitioner and the subject. Instead, the protocol is set against a rich background of social and professional roles, ethical standards, and legal and other obligations, which shape a subject's reasonable expectations. Notice generally only covers notable departures from these expectations and consent is a

limited and selective waiver of rights that subjects normally would expect to be respected. In other words, individuals understand that

> obligations and expectations are presupposed by informed consent practices. When they are waived by giving consent, they are not discarded or marginalized: they are merely waived in limited ways, for a limited time, for a limited purpose. In consenting to an appendectomy I do not consent to other irrelevant incisions, or to incisions by persons other than the relevant surgeon. In consenting to take part in a clinical trial I do not consent to swallow other novel medicines, let alone medicines that are irrelevant to my condition. Informed consent matters because it offers a standard and controllable way of setting aside obligations and prohibitions for limited and specific purposes.[86]

According to O'Neill and Manson, consent is not required for acceptable, expected behaviors, but only for those that depart from it. The burden on notice, therefore, is to describe clearly the violations of norms, standards, and expectations for which a waiver is being asked and not to describe everything that will be done and not done in the course of treatment or research, which both the researcher and the subjects can safely presume. Manson and O'Neill decline to produce a general or universal list of legal and ethical claims that applies to all treatment and research scenarios because, while all would surely include a common set of obvious prohibitions on, say, killing, stealing, injury, torture, fraud, deception, manipulations, and so forth, each would further include prohibitions and prescriptions relevant to the particular treatment or study in which subjects are engaged. For example, subjects may reasonably expect physicians, researchers, and others to perform in accordance with the training and professional commitments required in their respective fields, for example, to prescribe only the treatment and medication they believe to be the best and necessary for a patient's condition.

It is not sufficient for researchers to provide assurances that subjects are given a choice to waive or not to waive; they must be able to justify "actions that otherwise violate important norms, standards or expectations."[87] According to O'Neill and Manson, "[a]ny justification of informed consent has therefore to start from a recognition of the underlying legal and ethical claims and legitimate expectations that are selectively waived by consent transactions, and the reasons individuals may have for waiving them in particular cases."[88] In other words, selective waivers may not be requested for just anything but are acceptable under two conditions, either concerning actions for which individuals are presumed to have reasons to waive rights and obligations, or concerning actions that promise

significant benefits to others and to society at large. In other words, consent cannot exist as an excuse for anything, a limitation further emphasized by the second and third key principles of scientific integrity in the treatment of human subjects, namely, justice and beneficence (or non-maleficence.) Scientists requesting a limited waiver must ensure that subjects are well informed of departures from expected behaviors and they should ensure that the waiver they are requesting is consistent with the reasons their subjects have for waiving these rights. But informed consent is constrained in one further, crucial way – namely, by the requirements of beneficence, non-maleficence, and justice. These constrain what a subject can be asked to consent to.

When we understand informed consent as a limited waiver of rights and obligations, certain aspects of existing practices applied to privacy come to light. To begin, since FIPPs have served as a guide to law and policy, the focus has been on specifying the characteristics of notice and consent and very little on rights and obligations. Drawing on Manson and O'Neill, it is quite clear why this has not worked; it is impossible, even absurd to believe that notice and consent can fully specify the terms of interaction between data collector and data subject. The arguments in our paper attest to this. For too long, we have focused on the foreground, working at it from every angle. In good faith, we have crammed into the notice and consent protocol all our moral and political anxieties, believing that this is the way to achieve the level playing field,[89] to promote the autonomy of data subjects, to energize a competitive marketplace for good data practices, and more. In our view, this became a futile effort at some point along the way for reasons we and others have repeatedly offered. It is time to contextualize consent by bringing the landscape into focus. It is time for the background of rights, obligations, and legitimate expectations to be explored and enriched so that notice and consent can do the work for which it is best suited.[90]

Until now, the greatest obligation of data gatherers was either to anonymize data and pull it outside various privacy requirements or to inform and obtain consent. After charting the increasing difficulty of fulfilling these obligations in the face of big data, we presented the ultimate challenge: not of practical difficulty but of irrelevance. Where, for example, anonymizing data, adopting pseudonyms, or granting or withholding consent makes no difference to outcomes for an individual, we had better be sure that the outcomes in question can be defended as morally and politically legitimate. When anonymity and consent do make a difference, we learn from the domain of scientific integrity that simply because someone is anonymous or pseudonymous or has consented does not by itself

legitimate the action in question. A burden is upon the collector and user of data to explain why a subject has good reason to consent, even if consenting to data practices that lie outside the norm. That, or there should be excellent reasons why social and contextual ends are served by these practices.

We have argued that background and context-driven rights and obligations have been neglected in favor of anonymity and consent to the detriment of individuals and social integrity. Although our chapter will be deeply vexing to those who have placed anonymization and consent at the foundation of privacy protection, we welcome the shift in focus to the purposes to which data practices are being put and how these comport with individual interests as well as ethical, political, and context-driven values.

Acknowledgements The authors gratefully acknowledge research support from Intel Science and Technology Center for Social Computing, DHHS Strategic Healthcare Information Technology Advanced Research Projects on Security (SHARPS), NSF Cyber-Trust Collaborative Research (CNS-0831124), and Lady Davis Trust, The Hebrew University of Jerusalem.

NOTES

1. For a wide-ranging set of opinions on these matters, see David Bollier, *The Promise and Peril of Big Data* (Washington, DC: The Aspen Institute, 2010) and Janna Anderson and Lee Rainie, *The Future of Big Data* (Washington, DC: Pew Research Center, July 20, 2012).
2. For a broad overview, see Solon Barocas, "Data Mining: An Annotated Bibliography," *Cyber-Surveillance in Everyday Life: An International Workshop* (Toronto, Canada: University of Toronto, 12–15 May 2011), http://www.digitallymediatedsurveillance.ca/wp-content/uploads/2011/04/Barocas_Data_Mining_Annotated_Bibliography.pdf.
3. See e.g. Latanya Sweeney, "K-Anonymity: A Model for Protecting Privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, no. 5 (October 2002): 557–570, doi:10.1142/S0218488502001648; Arvind Narayanan and Vitaly Shmatikov, "Robust De-Anonymization of Large Sparse Datasets" (presented at the 2008 IEEE Symposium on Security and Privacy, IEEE, 2008), 111–125, doi:10.1109/SP.2008.33; Paul Ohm, "Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization," *UCLA Law Review* 57, no. 6 (August 2010): 1701–1777; Solon Barocas and Helen Nissenbaum, "On Notice: The Trouble with Notice and Consent" (presented at the Engaging Data: First International Forum on the Application and Management of Personal Electronic Information, Cambridge, MA, 2009); Lorrie Faith Cranor, "Necessary but

Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice," *Journal on Telecommunications and High Technology Law* 10, no. 2 (Summer 2012): 273–445; Alessandro Acquisti and Jens Grossklags, "Privacy and Rationality in Individual Decision Making," *IEEE Security and Privacy Magazine* 3, no. 1 (January 2005): 26–33, doi:10.1109/MSP.2005.22; Daniel J. Solove, "Privacy Self-Management and the Consent Dilemma," *Harvard Law Review* 126, no. 7 (May 2013): 1880–1880.

4. James Manyika et al., *Big Data: The Next Frontier for Innovation, Competition, and Productivity* (McKinsey Global Institute, 2011).
5. *Demystifying Big Data: A Practical Guide to Transforming the Business of Government* (Washington, DC: TechAmerica Foundation, 2012).
6. "NSF Advances National Efforts Enabling Data-Driven Discovery" (Washington, DC: National Science Foundation, November 12, 2013).
7. E.g. *Big Data* (http://www.liebertpub.com/big).
8. E.g. *Big Data and Society* (http://bigdatasoc.blogspot.com/p/big-data-and-society.html).
9. See Tony Hey, Stewart Tansley, and Kristin Tolle, eds., The Fourth Paradigm: Data-Intensive Scientific Discovery (Redmond, WA: Microsoft Research, 2009); *Frontiers in Massive Data Analysis* (Washington, DC: The National Academies Press, 2013); Pete Warden, *Big Data Glossary* (Sebastopol, CA: O'Reilly Media, 2011).
10. Mireille Hildebrandt, "Defining Profiling: A New Type of Knowledge?" in *Profiling the European Citizen: Cross-Disciplinary Perspectives*, ed. Mireille Hildebrandt and Serge Gutwirth (Dordrecht, Netherlands: Springer, 2008), 17–45, doi:10.1007/978-1-4020-6914-7_2; danah boyd and Kate Crawford, "Critical Questions for Big Data," *Information, Communication & Society* 15, no. 5 (June 2012): 662–679, doi:10.1080/1369118X.2012.678878; Christopher Steiner, *Automate This: How Algorithms Came to Rule Our World* (New York: Portfolio/Penguin, 2012); Viktor Mayer-Schönberger and Kenneth Cukier, *Big Data: A Revolution That Will Transform How We Live, Work, and Think* (New York: Houghton Mifflin Harcourt, 2013).
11. Manyika et al., *Big Data: The Next Frontier for Innovation, Competition, and Productivity*; Steiner, *Automate This: How Algorithms Came to Rule Our World*; Mayer-Schönberger and Cukier, *Big Data: A Revolution That Will Transform How We Live, Work, and Think*.
12. *Frontiers in Massive Data Analysis*.
13. See Usama Fayyad, "The Digital Physics of Data Mining," *Communications of the ACM* 44, no. 3 (March 1, 2001): 62–65, doi:10.1145/365181.365198; David Weinberger, "The Machine That Would Predict the Future," *Scientific American* 305, no. 6 (November 15, 2011): 52–57, doi:10.1038/scientificamerican1211-52; Foster Provost and Tom Fawcett, "Data Science and Its Relationship to Big Data and Data-Driven Decision Making," *Big Data* 1, no. 1 (March 2013): 51–59, doi:10.1089/big.2013.1508; Vasant Dhar, "Data Science and Prediction," *Communications of the ACM* 56, no. 12 (December 1, 2013): 64–73, doi:10.1145/2500499.
14. See e.g. Oscar H. Gandy Jr., "Consumer Protection in Cyberspace," *tripleC: Communication, Capitalism & Critique* 9, no. 2 (2011): 175–189; Cynthia Dwork and Deirdre K. Mulligan, "It's Not Privacy, and It's Not Fair," *Stanford Law*

Review Online 66 (September 3, 2013): 35–40; Omer Tene and Jules Polonetsky, "Judged by the Tin Man: Individual Rights in the Age of Big Data," *Journal on Telecommunications and High Technology Law*, August 15, 2013; Jonas Lerman, "Big Data and Its Exclusions," *Stanford Law Review Online* 66 (September 3, 2013): 55–63; Kate Crawford and Jason Schultz, "Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms," *Boston College Law Review* 55, no. 1 (2014).

15. For a more detailed account, see Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford, CA: Stanford University Press, 2010).
16. Applicants' ability to pay is already a controversial factor in the admissions decisions of many colleges; in locating less obvious correlates for the ability to pay, analytics may grant colleges the capacity to pursue similar ends without direct access to such information while also shielding such contentious practices from view.
17. Solon Barocas, "How Data Mining Discriminates," in *Data Mining: Episteme, Ethos, and Ethics*, PhD dissertation, New York University (Ann Arbor, MI: ProQuest Dissertations and Theses, 2014).
18. Such was the thinking in the so-called HEW report, where anonymized datasets were treated differently and separately under the heading of 'statistical databases'. Secretary's Advisory Committee on Automated Personal Data Systems, *Records, Computers and the Rights of Citizens* (U.S. Department of Health, Education, and Welfare, July 1973).
19. White House Office of Science and Technology Policy and the Networking and Information Technology R&D, *Data to Knowledge to Action*, Washington, DC, November 12, 2013, http://www.nitrd.gov/nitrdgroups/index.php?title=Data_to_Knowledge_to_Action.
20. Emily Steel and Julia Angwin, "On the Web's Cutting Edge, Anonymity in Name Only," *The Wall Street Journal*, August 4, 2010.
21. Michael Barbaro and Tom Zeller, "A Face Is Exposed for AOL Searcher No. 4417749," *The New York Times*, August 9, 2006.
22. Vincent Toubiana and Helen Nissenbaum, "An Analysis of Google Logs Retention Policies," *Journal of Privacy and Confidentiality* 3, no. 1 (2011): 2.
23. Sweeney, "K-Anonymity: A Model for Protecting Privacy;" Narayanan and Shmatikov, "Robust De-Anonymization of Large Sparse Datasets"; Cynthia Dwork, "Differential Privacy" (presented at the ICALP'06 Proceedings of the 33rd International Conference on Automata, Languages and Programming, Berlin: Springer, 2006), 1–12, doi:10.1007/11787006_1; Cynthia Dwork, "A Firm Foundation for Private Data Analysis," *Communications of the ACM* 54, no. 1 (January 1, 2011): 86, doi:10.1145/1866739.1866758.
24. Ohm, "Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization"; Jane Yakowitz, "Tragedy of the Data Commons," *Harvard Journal of Law & Technology* 25, no. 1 (Autumn 2012): 1–67; Felix T Wu, "Defining Privacy and Utility in Data Sets," *University of Colorado Law Review* 84, no. 4 (2013): 1117–1177; Jane Bambauer, Krishnamurty Muralidhar, and Rathindra Sarathy, "Fool's Gold: An Illustrated Critique of Differential Privacy," *Vanderbilt Journal of Entertainment & Technology Law* 16 (2014).

25. Christopher Soghoian, "The Problem of Anonymous Vanity Searches," *I/S: A Journal of Law and Policy for the Information Society* 3, no. 2 (2007).
26. Barbaro and Zeller, "A Face Is Exposed for AOL Searcher No. 4417749."
27. Yves-Alexandre de Montjoye et al., "Unique in the Crowd: The Privacy Bounds of Human Mobility," *Scientific Reports* 3 (2013): 1376, doi:10.1038/srep01376.
28. Khaled El Emam et al., "A Systematic Review of Re-Identification Attacks on Health Data," ed. Roberta W Scherer, *PLoS ONE* 6, no. 12 (December 2, 2011): e28071, doi:10.1371/journal.pone.0028071.s001.
29. Narayanan and Shmatikov, "Robust De-Anonymization of Large Sparse Datasets."
30. Dwork, "A Firm Foundation for Private Data Analysis."
31. Sweeney, "K-Anonymity: a Model for Protecting Privacy."
32. Dwork, "Differential Privacy."
33. Helen Nissenbaum, "The Meaning of Anonymity in an Information Age," *The Information Society* 15, no. 2 (May 1999): 142, doi:10.1080/019722499128592.
34. Of course, this is why anonymity, in certain contexts and under certain conditions, can be problematic.
35. Alistair Barr, "Google May Ditch 'Cookies' as Online Ad Tracker," *USA Today*, September 17, 2013.
36. Ashkan Soltani, "Questions on the Google AdID," *Ashkan Soltani*, September 19, 2013, http://ashkansoltani.org/2013/09/19/questions-on-the-google-adid/.
37. In practice, this has tended to refer to what we commonly conceive as 'contact information'.
38. Natasha Singer, "Acxiom, the Quiet Giant of Consumer Database Marketing," *The New York Times*, June 16, 2012.
39. Gary T. Marx, "What's in a Name? Some Reflections on the Sociology of Anonymity," *The Information Society* 15, no. 2 (May 1999): 99–112, doi:10.1080/019722499128565.
40. Arvind Narayanan and Vitaly Shmatikov, "Myths and Fallacies of 'Personally Identifiable Information'," *Communications of the ACM* 53, no. 6 (June 1, 2010): 24–26, doi:10.1145/1743558.
41. This explains the ongoing attempt, as part of European Data Protection reform, to broaden the definition of 'personal data' to cover any such data that allows for the "singling out" of individuals, whether or not they can be identified as traditionally understood. The Article 29 Working Party, for instance, has advised "that a natural person can be considered identifiable when, within a group of persons, (s)he can be distinguished from others and consequently be treated differently. This means that the notion of identifiability includes singling out." *Statement of the Working Party on Current Discussions Regarding the Data Protection Reform Package* (European Commission, February 27, 2013). For a more detailed discussion of the contours of this debate in the context of online behavioral advertising, see Frederik Zuiderveen Borgesius, "Behavioral Targeting: A European Legal Perspective," *IEEE Security and Privacy Magazine* 11, no. 1 (January 2013): 82–85, doi:10.1109/MSP.2013.5.
42. This is not to suggest that pseudonyms are valueless. Most immediately, pseudonyms limit the potential to infer gender, race, national origin, religion,

or class position from names that possess any such obvious associations. *One-off* pseudonyms (i.e., unique identifiers that are *not* common to multiple databases) also do not lend themselves to the kind of straightforward matching of records facilitated by traditional categories of identity. In principle, only the institution that assigns a one-off pseudonym to a specific person can recognize that person *according* to that pseudonym. And where this pseudonym has been abandoned or replaced (e.g. by expiring or deleting a cookie), even the institution that assigned it to a specific individual will no longer be able to recognize or associate prior observations with that person.

43. Jennifer Valentino-Devries and Jeremy Singer-Vine, "They Know What You're Shopping for," *The Wall Street Journal*, December 7, 2012.
44. "'Drinking from a Fire Hose': Has Consumer Data Mining Gone Too Far?," *Knowledge@Wharton*, November 22, 2011, http://knowledge.wharton.upenn.edu/article.cfm?articleid=2886.
45. Steel and Angwin, "On the Web's Cutting Edge, Anonymity in Name Only."
46. "'Drinking From a Fire Hose': Has Consumer Data Mining Gone Too Far?"
47. Cindy Waxer, "Big Data Blues: The Dangers of Data Mining," *Computerworld*, November 4, 2013, http://www.computerworld.com/s/article/print/9243719/Big_data_blues_The_dangers_of_data_mining.
48. Valentino-Devries and Singer-Vine, "They Know What You're Shopping For."
49. This, too, explains the dispute over an article in the current draft of the proposed revision of the European Data Protection laws that stipulates that profiling "based solely on the processing of pseudonymous data should be presumed not to significantly affect the interests, rights or freedoms of the data subject." For more details about this point of debate, see Monika Ermert, "EU Data Protection: Bumpy Piece of Road Ahead," *Internet Policy Review*, October 24, 2013, http://policyreview.info/articles/news/eu-data-protection-bumpy-piece-road-ahead/209.
50. Serge Gutwirth and Paul Hert, "Regulating Profiling in a Democratic Constitutional State," in *Profiling the European Citizen: Cross-Disciplinary Perspectives*, ed. Mireille Hildebrandt and Serge Gutwirth (Dordrecht, Netherlands: Springer, 2008), 289, doi:10.1007/978-1-4020-6914-7_14.
51. Brian Dalessandro, "The Science of Privacy," *Ad:Tech*, July 30, 2013, http://blog.ad-tech.com/the-science-of-privacy/.
52. Dalessandro, "The Science of Privacy."
53. Quentin Hardy, "Rethinking Privacy in an Era of Big Data," *The New York Times*, June 4, 2012, http://bits.blogs.nytimes.com/2012/06/04/rethinking-privacy-in-an-era-of-big-data/.
54. Solon Barocas, "Extending the Frontier of the Inferable: Proxies, Proximity, and Privacy," in *Data Mining: Episteme, Ethos, and Ethics.*
55. Dwork, "A Firm Foundation for Private Data Analysis."
56. The authors of the study, somewhat surprised by the fierce reaction to news of their findings, assembled and responded to various criticisms: Frances H. Montgomery et al., "Monitoring Student Internet Patterns: Big Brother or Promoting Mental Health?" *Journal of Technology in Human Services* 31, no. 1 (January 2013): 61–70, doi:10.1080/15228835.2012.756600.

57. Raghavendra Katikalapudi et al., "Associating Internet Usage with Depressive Behavior among College Students," *IEEE Technology and Society Magazine* 31, no. 4 (Winter 2012): 73–80, doi:10.1109/MTS.2012.2225462.
58. Among other things, FIPPs also require that adequate steps be taken to secure the information.
59. OECD, *The Evolving Privacy Landscape: 30 Years after the OECD Privacy Guidelines*, vol. 176, April 6, 2011, doi:10.1787/5kgf09z90c31-en.
60. Gramm-Leach-Bliley Act, 15 USC, § 6801-6809. See also Chapters 1 and 7 in this volume.
61. Yannis Bakos, Florencia Marotta-Wurgler, and David R. Trossen, "Does Anyone Read the Fine Print? Testing a Law and Economics Approach to Standard Form Contracts," *SSRN Electronic Journal* (2009), doi:10.2139/ssrn.1443256; Aleecia M. McDonald and Lorrie Faith Cranor, "The Cost of Reading Privacy Policies," *I/S: a Journal of Law and Policy for the Information Society* 4, no. 3 (2008): 540–565; Aleecia M. McDonald et al., "A Comparative Study of Online Privacy Policies and Formats" (presented at the PETS 2009, Springer, 2009), 37–55. Also discussed in Helen Nissenbaum, "A Contextual Approach to Privacy Online," *Daedalus*, 140, no. 4 (Fall 2011): 32–48.
62. Julie Brill, "Reclaim Your Name: Privacy in the Age of Big Data" (presented at the Sloan Cyber Security Lecture, Brooklyn, NY, 2013).
63. In 2012, Microsoft hosted a series of 'dialogs' in which scholars. executives, advocates, and regulators discussed the future of notice and consent in the wake of big data, many of whom expressed this sentiment. For a summary of the full range of opinions at these events, see Fred H. Cate and Viktor Mayer-Schönberger, "Notice and Consent in a World of Big Data," *International Data Privacy Law* 3, no. 2 (May 20, 2013): 67–73, doi:10.1093/idpl/ipt005. Similar arguments have been advanced in Mireille Hildebrandt, "Who Is Profiling Who? Invisible Visibility," in *Reinventing Data Protection?* ed. Serge Gutwirth et al. (Dordrecht, Netherlands: Springer, 2009), 239–252, doi:10.1007/978-1-4020-9498-9_14; Christopher Kuner et al., "The Challenge of 'Big Data' for Data Protection," *International Data Privacy Law* 2, no. 2 (April 23, 2012): 47–49, doi:10.1093/idpl/ips003; *Big Data and Analytics: Seeking Foundations for Effective Privacy Guidance* (Washington, DC: The Centre for Information Policy Leadership, February 28, 2013); Omer Tene and Jules Polonetsky, "Big Data for All: Privacy and User Control in the Age of Analytics," *Northwestern Journal of Technology and Intellectual Property* 11, no. 5 (April 2013): 239–272; Ira Rubinstein, "Big Data: The End of Privacy or a New Beginning?" *International Data Privacy Law* 3, no. 2 (May 20, 2013): 74–87, doi:10.1093/idpl/ips036.
64. Nir Eyal, "Informed Consent," in *The Stanford Encyclopedia of Philosophy* (Fall 2012 Edition), ed. Edward N. Zalta (ed.), http://plato.standford.edu/archives/fall2012/entries/informed-consent/.
65. Lorrie Faith Cranor, "Can Users Control Online Behavioral Advertising Effectively?" *IEEE Security and Privacy Magazine* 10, no. 2 (n.d.): 93–96, doi:10.1109/MSP.2012.32; Pedro Giovanni Leon et al., "What Do Online Behavioral Advertising Privacy Disclosures Communicate to Users?" (presented at the WPES '12 Proceedings of the 2012 ACM workshop on Privacy in the electronic society, New York, NY: ACM Press, 2012), 19–30,

doi:10.1145/2381966.2381970; Omer Tene and Jules Polonetsky, "To Track or 'Do Not Track': Advancing Transparency and Individual Control in Online Behavioral Advertising," *Minnesota Journal of Law, Science & Technology* 13, no. 1 (Winter 2012): 281–357; Frederik J. Zuiderveen Borgesius, "Consent to Behavioural Targeting in European Law – What Are the Policy Implications of Insights from Behavioural Economics?" *SSRN Electronic Journal* (2013), doi:10.2139/ssrn.2300969; Joseph Turow, "Self-Regulation and the Construction of Media Harms: Notes on the Battle over Digital 'Privacy'," in *Routledge Handbook of Media Law*, ed. Monroe E Price, Stefaan Verhulst, and Libby Morgan (New York, NY: Routledge, 2013).

66. "Carnegie Mellon Leads NSF Project to Help People Understand Web Privacy Policies," *Carnegie Mellon News* (Pittsburgh, PA: Carnegie Mellon University, August 20, 2013). See Usable Privacy Policy Project: http://www.usableprivacy.org/.
67. *Protecting Consumer Privacy in an Era of Rapid Change* (Washington, DC: Federal Trade Commission, March 2012).
68. Helen Nissenbaum, "A Contextual Approach to Privacy Online," *Daedalus* 140, no. 4 (October 2011): 32–48, doi:10.1162/DAED_a_00113.
69. For a summary of the relevant research, see Solove, "Privacy Self-Management and the Consent Dilemma."
70. Mireille Hildebrandt, "Profiling and the Rule of Law," *Identity in the Information Society* 1, no. 1 (December 19, 2008): 55–70, doi:10.1007/s12394-008-0003-1.
71. Data miners were concerned with the principle of consent for this very reason from early on in the field's history; see e.g. Daniel E. O'Leary, "Some Privacy Issues in Knowledge Discovery: The OECD Personal Privacy Guidelines," *IEEE Expert: Intelligent Systems and Their Applications* 10, no. 2 (1995): 48–59.
72. Herman T. Tavani, "KDD, Data Mining, and the Challenge for Normative Privacy," *Ethics and Information Technology* 1, no. 4 (1999): 265–273, doi:10.1023/A:1010051717305; Mireille Hildebrandt, "Who Is Profiling Who?"; Mireille Hildebrandt, "Profiling and AmI," in *The Future of Identity in the Information Society*, ed. Kai Rannenberg, Denis Royer, and André Deuker (Berlin: Springer, 2009), 273–310, doi:10.1007/978-3-642-01820-6_7; Serge Gutwirth and Mireille Hildebrandt, "Some Caveats on Profiling," in *Data Protection in a Profiled World*, ed. Serge Gutwirth, Yves Poullet, and Paul De Hert (Dordrecht, Netherlands: Springer, 2010), 31–41, doi:10.1007/978-90-481-8865-9_2; Cate and Mayer-Schönberger, "Notice and Consent in a World of Big Data"; *Big Data and Analytics: Seeking Foundations for Effective Privacy Guidance*; Tene and Polonetsky, "Big Data for All: Privacy and User Control in the Age of Analytics"; Rubinstein, "Big Data: The End of Privacy or a New Beginning?"
73. We should stress that this is not the same argument, advanced by a number of scholars, that the capacity to assess any particular individual depends on the willingness (or unwillingness, as the case may be) of other individuals to reveal data about themselves. The common example in these arguments is that *I* am a good customer because *you* are less profitable. Taking the example of car insurance, Tverdek explains that what constitutes a 'good' driver is a statistical artifact that can only be made in contrast to a statistically 'reckless' driver. But the phenomena that we are describing here, to stick with the same example, is the capacity for

insurers to predict that *I* am bad driver because I share certain qualities with the limited number of other bad drivers who chose to report their accidents. Edward Tverdek, "Data Mining and the Privatization of Accountability," *Public Affairs Quarterly* 20, no. 1 (2006): 67–94. See also Scott R. Peppet, "Unraveling Privacy: The Personal Prospectus and the Threat of a Full Disclosure Future," *Northwestern University Law Review* 105, no. 3 (2011): 1153–1203, and, more recently, Evgeny Morozov, "The Real Privacy Problem," *MIT Technology Review*, October 22, 2013.

74. danah boyd, "Networked Privacy" (presented at the Personal Democracy Forum 2011, New York, NY, 2011).
75. Alan Mislove et al., "You Are Who You Know: Inferring User Profiles in Online Social Networks" (presented at the WSDM '10 Proceedings of the Third ACM International Conference on Web Search and Data Mining, New York, NY: ACM Press, 2010), 251–260, doi:10.1145/1718487.1718519.
76. Carter Jernigan and Behram F. T. Mistree, "Gaydar: Facebook Friendships Expose Sexual Orientation," *First Monday* 14, no. 10 (September 25, 2009), doi:10.5210/fm.v14i10.2611.
77. Emöke-Ágnes Horvát et al., "One Plus One Makes Three (for Social Networks)," ed. Sergio Gómez, *PLoS ONE* 7, no. 4 (April 6, 2012): e34740, doi:10.1371/journal.pone.0034740.s011.
78. Charles Duhigg, "How Companies Learn Your Secrets," *The New York Times Magazine*, February 16, 2012.
79. Rachel Nolan, "Behind the Cover Story: How Much Does Target Know?" *The New York Times*, February 21, 2012, http://6thfloor.blogs.nytimes.com/2012/02/21/behind-the-cover-story-how-much-does-target-know/.
80. Scholars have described this as the problem of 'categorical privacy',[80] whereby an individual's apparent membership in a group reveals more about them than can be observed directly (i.e., inferring that they likely possess the same traits as other group members). But the focus of this line of thinking has been on the impossibility of individuals foreseeing these potential inferences, the problem of inaccurate stereotyping, and the absence of associational ties, rather than the limited amount of examples that would be necessary to induce the rule and then apply it to others. See, in particular, Anton Vedder, "KDD: the Challenge to Individualism," *Ethics and Information Technology* 1, no. 4 (1999): 275–281, doi:10.1023/A:1010016102284.
81. Mislove et al., "You Are Who You Know: Inferring User Profiles in Online Social Networks," 255. We should put this result into context to ensure that we do not overstate its significance: Mislove et al. were looking at the relatively innocuous details posted by college students on Facebook, specifically their major, year of expected graduation, and college (at this particular university, students are assigned to a residential college where they tend to remain for the duration of their college career). The stakes are not especially high in this case. That said, these inferences were only based on the very limited set of variables (in fact, they only looked at the same attributes that they hoped to infer), rather than the far richer data that social networking sites accrue about their users. The authors speculate that inferences about far more sensitive attributes would be possible if the analysis were to consider a larger set of criteria that might prove statistically relevant.

82. Scholars have already proposed such a method: Christina Aperjis and Bernardo A. Huberman, "A Market for Unbiased Private Data: Paying Individuals According to Their Privacy Attitudes," *First Monday* 17, no. 5 (May 4, 2012), doi:10.5210/fm.v17i5.4013.
83. Yakowitz, "Tragedy of the Data Commons."
84. Tal Z. Zarsky, "Desperately Seeking Solutions Using Implementation-Based Solutions for the Troubles of Information Privacy in the Age of Data Mining and the Internet Society," Maine Law Review 56, no. 1 (2004): 14–59; Cate and Mayer-Schönberger, "Notice and Consent in a World of Big Data"; *Big Data and Analytics: Seeking Foundations for Effective Privacy Guidance*; Tene and Polonetsky, "Big Data for All: Privacy and User Control in the Age of Analytics"; Rubinstein, "Big Data: the End of Privacy or a New Beginning?"
85. Oscar H. Gandy, *Coming to Terms with Chance: Engaging Rational Discrimination and Cumulative Disadvantage* (Burlington, VT: Ashgate, 2009); Dwork and Mulligan, "It's Not Privacy, and It's Not Fair"; Tene and Polonetsky, "Judged by the Tin Man"; Omer Tene and Jules Polonetsky, "A Theory of Creepy: Technology, Privacy and Shifting Social Norms," *Journal on Telecommunications and High Technology Law*, September 16, 2013; Gutwirth and Hildebrandt, "Some Caveats on Profiling"; Tal Z. Zarsky, "'Mine Your Own Business!': Making the Case for the Implications of the Data Mining of Personal Information in the Forum of Public Opinion," *Yale Journal of Law & Technology* 5 (2004): 1–57.
86. Neil C. Manson and Onora O'Neill, *Rethinking Informed Consent in Bioethics* (New York: Cambridge University Press, 2012), 73.
87. Ibid., 75.
88. Ibid., 73.
89. Secretary's Advisory Committee on Automated Personal Data Systems, *Records, Computers and the Rights of Citizens*.
90. This idea has been picked up by the Federal Trade Commission, which, as one of the key recommendations of its 2012 report, suggested that "companies do not need to provide choice before collecting and using consumers' data for practices that are consistent with the context of the transaction, consistent with the company's relationship with the consumer," *Protecting Consumer Privacy in an Era of Rapid Change*, vii.