

Is privacy *privacy*?*

Kobbi Nissim[†]

Alexandra Wood[‡]

September 21, 2018

Abstract

This position paper observes how different technical and normative conceptions of privacy have evolved in parallel and describes the practical challenges that these divergent approaches pose. Notably, past technologies relied on intuitive, heuristic understandings of privacy that have since been shown not to satisfy expectations for privacy protection. With computations ubiquitously integrated in almost every aspect of our lives, it is increasingly important to ensure that privacy technologies provide protection that is in line with relevant social norms and normative expectations. Similarly, it is also important to examine social norms and normative expectations with respect to the evolving scientific study of privacy.

To this end, we argue for a rigorous analysis of the mapping from normative to technical concepts of privacy and vice versa. We review the landscape of normative and technical definitions of privacy and discuss specific examples of gaps between definitions that are relevant in the context of privacy in statistical computation. We then identify opportunities for overcoming their differences in the design of new approaches to protecting privacy in accordance with both technical and normative standards.

Keywords: Informational Privacy, Privacy Law, Formal Privacy Models, Differential Privacy

1 Introduction

Privacy concerns regarding the collection, storage, and use of personal information are a recurring topic of public discourse. Data analysis is now deeply and irrevocably embedded in our systems: from social networks, to electronic commerce, medical care, national security, research, and more. Analyses are being fed with increasingly detailed information—including data about our traits, our relationships, our behaviors, our interests, and our preferences. Such analyses, individually or in combination with other analyses, can reveal sensitive attributes of individuals, including information about their health, finances, political leanings, and social behaviors. The growing ubiquity of data analyses therefore implicates social norms about privacy. Misuse or disclosure of such data can adversely affect an individual’s relationships, reputation, employability, insurability, or financial status, or even lead to civil liability, criminal penalties, or bodily harm. It can also have

*For the journal version, see *Philosophical transactions of the Royal Society A*, Volume 376, Issue 2128, Spring 2018. DOI 10.1098/rsta.2017.0358

[†]Department of Computer Science, Georgetown University. kobbi.nissim@georgetown.edu.

[‡]Berkman Klein Center for Internet & Society, Harvard University. awood@cyber.law.harvard.edu.

negative consequences for essential rights and social values such as freedom of expression, freedom of association, and respect for private and family life, more broadly [1, 2].

In light of the attendant risks to individuals and to social values, organizations that manage personal data implement various measures to protect the privacy of those whose personal information is used. Yet the data privacy landscape is rapidly evolving, as advances in technology change how personal information is collected, stored, shared, analyzed, and disseminated. Organizations increasingly rely on technical safeguards to address growing data privacy risks. A common premise is that these technical safeguards protect individual privacy in accordance with legal and social norms. Privacy technologies are viewed through this lens because privacy is inherently a normative concept, with foundations in philosophical, legal, sociological, political, and economic traditions. Various privacy regulations and policies attempt to capture and codify these norms and values as enforceable constraints on behavior. Examples can be found in laws protecting educational records [3], medical records [4], information collected by the US Census Bureau [5], information maintained by US federal agencies [6], and personal data about individuals in the European Union [7].

Expectations and understandings of privacy are dramatically shifting, in response to a wide spectrum of privacy-invasive technologies underlying applications in everyday use by government agencies and businesses alike. A range of technical measures have been deployed to mitigate informational privacy risks. Common approaches include suppression, data swapping, noise addition, synthetic data, aggregation via statistical and machine learning tools, query logging and auditing, and more. With this expanded reliance on technological approaches, privacy is increasingly becoming a technical concept, in addition to being a normative concept. Further, the understanding of privacy reflected in these technologies is evolving over time—partly in response to the discovery of new vulnerabilities, a process that has significantly accelerated in the last two decades. Key examples illustrating the evolution of privacy practices can be found in the policies of federal statistical agencies, which have sought to strengthen protections in response to newly identified threats to privacy. For instance, the US Census Bureau has adopted different approaches to confidentiality over time, beginning with suppression and aggregation, then introducing data swapping methods, and, most recently, exploring uses of formal privacy models like differential privacy [8, 9].

However, significant gaps between technical and normative conceptions of privacy pose challenges for the development and deployment of privacy technologies. The technical language is often mathematical, precise, and somewhat rigid whereas the language describing social norms is flexible and less precise. There are substantive differences in scope, in what is considered personal information, in what is considered a privacy threat, in the protection required, and in the data formats the normative and technical conceptions contemplate. With such divergent foundations, it is difficult to reason about how these conceptions of privacy interact, or, more critically, whether they are in agreement at all.

The picture is even further complicated by the existence of gaps between normative conceptions of privacy in different cultures, and between disciplines. Gaps exist between these conceptions and their implementation in laws and regulations, and between the latter and technical conceptions of privacy. Describing and reconciling these gaps is a task of extremely large scope. For this reason, in this paper we restrict our attention to the interactions between technical and normative aspects of privacy, focusing specifically on the setting of data analysis and on conceptions of informational privacy coming from the legal and computer science literatures. Developing an understanding of the mechanics of the interactions between technical and normative aspects of privacy will be essential towards ensuring these various aspects of privacy are in harmony. For example, it will be critical

that legal conceptions of privacy be cognizant of what scientific knowledge deems deliverable. At the same time, technical definitions of privacy should provide a level of protection that meets societal needs and expectations.

2 The data privacy problem

The subject of informational privacy is wide, encompassing a vast multitude of concepts. For the purposes of providing a concrete illustration, this discussion focuses on a specific subset of the privacy problem, namely, an analysis of the differences between technical approaches for protecting privacy in statistical computation and the notions of de-identification and anonymization underlying many privacy regulations. This analysis yields a number of opportunities for overcoming the challenges created by the gaps between these varied concepts.

2.1 Privacy as a normative concept

Privacy is a normative concept deeply rooted in philosophical, legal, sociological, political, and economic traditions. Early principled discussion of privacy goes back to Aristotle’s distinction between public and private spheres of life [10]. An understanding of a vast range of privacy harms has since been developed by the literature and addressed by legal frameworks [11]. This piece focuses on the normative concepts regarding informational privacy embedded within various regulations and policies governing information privacy and data protection. Of relevance are regulations and policies restricting the release of statistical information about individuals or groups of individuals, whether released as raw data, de-identified data, or statistical summaries. These include regulatory requirements for de-identifying or anonymizing information prior to disclosure found in laws and related guidance protecting education [3] and health [4] records in the United States, data provided by respondents to the US Census Bureau [5], personally identifiable information held by US federal agencies [6], and personal data about individuals in the European Union [7], among many others around the world.

Privacy torts: An early precursor to information privacy law was an 1890 essay by Warren and Brandeis titled “The Right to Privacy” [12]. In this essay, Warren and Brandeis voiced concerns about the confluence of instantaneous photography and widespread newspaper circulation increasingly enabling journalists to intrude on private affairs. Characterizing privacy as the “right to be let alone” and as an essential component of “the right to one’s personality,” they invoked European philosophical and legal doctrine in articulating the right of an individual to develop his or her personality free from unwanted publicity [12]. They advocated a common law recognition of a right to privacy, an idea that has been highly influential, shaping the evolution of both common law and statutory law across the United States.

In 1960, Prosser published a summary of the subsequent jurisprudence on the right to privacy, finding that a majority of American courts recognized the right. Furthermore, he found that privacy law had evolved to comprise four distinct privacy torts, including intrusion upon a person’s seclusion or solitude, or into his or her private affairs; public disclosure of embarrassing private facts about an individual; publicity placing one in a false light in the public eye; and appropriation of one’s likeness for the advantage of another [13].

Modern courts continue to frame informational privacy interests in terms of the four privacy torts, carrying forward Prosser’s emphasis on redressing specific, tangible harms resulting from

a rather narrow subset of privacy-invasive activities. This narrow view of privacy-related harms, requiring a showing of actual or imminent physical, financial, or property injury, has led courts in the US to dismiss many data breach cases for lack of harm. Even cases alleging that an individual's sensitive information, such as his or her name and Social Security number, were breached as a result of a company's negligence have been dismissed for lack of injury [14]. These cases suggest that the privacy torts are ill-suited to address many types of modern data privacy risks, particularly those associated with leakages of information from releases of statistics for which the possible injuries are unlikely to be actual or imminent. Moreover, tort law's focus on the harms that result from unauthorized access to information about an individual arguably fails to capture other categories of privacy harms, namely the accumulated leakage of personal information from data to which access was properly granted.

Fair information practice principles: In 1973, the US Department of Health, Education and Welfare (HEW) published a report containing fair information practice principles for protecting personal data in record-keeping systems, in response to concerns about the growing use of automated data systems [15]. This report explicitly focuses on the protection of identifiable information. The principles set forth in this report govern the collection, use, and storage of "personal data that can be associated with identifiable information" either by "specific identification, such as name or Social Security number" or "because they include personal characteristics that make it possible to identify an individual with reasonable certainty" [15]. Based on a "concept of mutuality in record keeping," the report calls for safeguards enabling an individual "to find out what information about him is in a record and how it is used," "to prevent information about him obtained for one purpose from being used or made available for other purposes without his consent," and "to correct or amend a record of identifiable information about himself" [15]. The principles also require an organization maintaining identifiable personal data to "assure the reliability of the data for their intended use and [] take reasonable precautions to prevent misuse of the data" and hold that there must be no personal-data record-keeping systems "whose very existence is secret" [15].

In addition, the HEW report recognizes and provides a detailed discussion of the informational risks to individuals associated with statistical data publications. In order to protect research respondents, the report calls for new federal legislation to safeguard data "identifiable with, or traceable to, specific individuals" maintained by any statistical-reporting and research system. It notes that, when releasing data in statistical form, an organization should address the risk of "statistical disclosure," meaning "the risk that arises when a population is so narrowly defined that tabulations are apt to produce cells small enough to permit the identification of individual data subjects, or when a person using a statistical file has access to information which, if added to data in the statistical file, makes it possible to identify individual data subjects" [15].

Congress based provisions of the Privacy Act of 1974 [16] in large part on the findings of this report, thereby applying the fair information practice principles to all federal agencies in the United States. Similar statutory requirements have been enacted at the state level to govern the practices of state agencies (e.g., [17, 18]). International privacy guidelines based on the fair information practice principles have also been adopted by governments around the world, such as the European Union [19]. The most widely-followed guidelines are the privacy principles developed by the Organisation for Economic Co-operation and Development (OECD), which encompass the broadly-defined principles of collection limitation, data quality, purpose specification, use limitation, security safeguards, openness, individual participation, and accountability [20]. These principles apply only to "information relating to identified or identifiable individuals" and explicitly exclude

“anonymous data,” such as “collections of statistical data in anonymous form” [20].

Although strongly influenced by the 1973 HEW report, which provides an extended discussion of privacy risks and guidance on how to balance risks against uses of data, the laws and international guidelines that followed are significantly less detailed. For instance, because the “precise dividing line” between identifiable and anonymous data “may be difficult to draw,” the OECD principles leave further guidance on this issue to individual regulations at the national level [20]. More generally, critics have noted that the fair information practice principles are not self-implementing or self-enforcing, with actual implementation taking place at the statutory, regulatory, or organizational level [21]. Implementation of the principles is highly sector-specific and context-dependent and, further, “in any context is often more a matter of art and judgment rather than a science or mechanical translation of principles” [21].

De-identification: The fair information practice principles—and the statutes, regulations, and policies implementing them—center on concepts such as identifiable data and anonymous data. These concepts are grounded, in large part, in the longstanding practice of using an identifier to retrieve an individual’s record from an administrative record-keeping system (see [15]). Organizations have an obligation to protect this information, as it can potentially be used to access sensitive information about an individual from a database. For instance, federal statistical agencies have established a wide range of safeguards to ensure their statistical data releases do not contain identifiers that could be used to locate an individual’s record in an administrative system (see [15]).

Since then, notions of identifying information have continued to play a central role in information privacy laws enacted in jurisdictions around the world. Although they vary significantly with respect to the types of information protected and the safeguards required, they often turn on a specific definition of personally identifiable information, personal information, or personal data [22]. If information falls within a particular law’s definition of personal information, it is generally protected from disclosure. Definitions of personal information differ considerably across sectors, jurisdictions, and contexts. Some regulations offer narrow definitions, such as the Massachusetts data security regulation, which defines personal information as a Massachusetts resident’s name in combination with his or her Social Security number, driver’s license number, or financial account number [23]. Other regulations are considerably broader, such as the EU General Data Protection Regulation, which defines personal data as “any information relating to a data subject” [7].

Some information privacy laws expressly exclude information classified as de-identified or anonymous data from the definition of personal information. Information that has been transformed in accordance with regulatory requirements for de-identification can generally be shared more widely, or, in some cases, can even be disclosed publicly without further restriction on use or redisclosure. For example, the HIPAA Privacy Rule provides a safe harbor that permits the disclosure of health information that has been de-identified through the removal of information from a list of eighteen identifiers, such as names, Social Security numbers, and dates of birth [24]. Other regulations may require case-by-case determinations to be made when de-identifying information. For example, regulators have declined to specify whether a particular set of methods for de-identifying information is sufficient to meet the requirements of FERPA and, instead, instruct educational agencies and institutions to make a determination based on the data set itself, other data sources that may be available, and other context-specific factors (see [25]). The open-ended, contextual nature of de-identification standards poses challenges for data holders and other practitioners who often may not know with certainty whether they have adequately satisfied the de-identification requirements of the applicable laws.

Moreover, definitions of personal information are evolving over time in response to expanded notions of what information could be used to identify an individual via a privacy attack on a database. For instance, in 2017, the US Office of Management and Budget updated its guidance to federal agencies on preparing for and responding to a breach of personally identifiable information. The new definition of personally identifiable information advises that “information that is not PII can become PII whenever additional information becomes available—in any medium or from any source—that would make it possible to identify an individual” [6]. Differences between regulatory definitions, uncertainty regarding their scope, and changes in how they are defined and interpreted in light of the broader data privacy landscape are widely cited as weaknesses of the regulatory framework for privacy protection [22].

Contextual integrity: There is growing recognition that traditional conceptions of privacy fail to accurately capture normative expectations of privacy. A highly influential alternative framework is *contextual integrity*, introduced by Nissenbaum in 2004 [26] and incorporated in legislative proposals such as the Consumer Privacy Bill of Rights Act of 2015 [27]. This framework rejects notions of privacy as control over information about oneself or as a strict dichotomy between public vs. private information or sensitive vs. non-sensitive information. Privacy is, instead, best understood in terms of normative expectations about the appropriate flow of information.

Contextual integrity assesses how closely the flow of personal information conforms to context-relative informational norms. More precisely, in a context, the flow of information of a certain type about a subject from a sender to a recipient is governed by a particular transmission principle. Contextual integrity is violated when the norms in the relevant context are breached. Intuitively, it recognizes that certain parties may obtain certain types of information about other parties under the right terms and for the right reasons. For example, a patient (subject and sender) may communicate concerns about his or her medical condition (information type) to her physician (recipient) during an annual medical exam, with the understanding that the patient-physician relationship demands confidentiality (transmission principle). A physician sharing details about the patient’s medical condition with the patient’s insurer would not constitute a violation of contextual integrity, whereas sharing the same information with the patient’s employer would.

As an ethical justificatory framework for socio-technical systems, contextual integrity seeks to take into consideration the individual interests and preferences of affected parties, the ethical and political principles and values of societal distribution, and societal contextual functions, purposes, and values. It recognizes that contexts evolve over time in cultures and societies, subject to historical, cultural, geographic factors, and aims to capture how normative expectations depend on such contextual factors. An active area of research has sought to use this framework to disentangle various justifications for normative theories and legal conceptions of privacy using the contextual integrity framework [26]. Extensions of this research have also explored technical applications, such as the design of access control and privacy policies based on a formalization of privacy law requirements [28].

2.2 Privacy as a technical concept

In addition to being a normative concept, privacy is a technical concept underlying the design of privacy-preserving technologies. Such technologies draw from a range of definitions of privacy presented in the statistics and information security literature. As discussed below, these technical concepts often disagree in fundamental ways about what it means to protect individual privacy.

Anonymization and de-identification: Many privacy technologies are designed with the goal of de-identifying personal information. This approach equates privacy protection with making personal information anonymous or de-identified, i.e., preventing an individual’s information from being linked with him or her. The premise is that it is impossible (or, at least, very difficult) to infer personal information pertaining to an individual from a de-identified dataset or use it to violate an individual’s privacy in other ways.

The process of de-identifying data typically involves a combination of data redaction and coarsening techniques. Examples include suppression of directly identifying attributes (e.g., names, addresses, and identification numbers); suppression of indirectly identifying attributes, i.e., those that can be linked to external data sources that contain identifying information (e.g., the combination of ZIP code, sex, and date of birth); generalization or coarsening of data (e.g., by grouping ages into ranges or grouping geographic locations with small populations); and aggregating and perturbing data (see, e.g., [29, 30]).

Concerns about the efficacy of de-identification techniques have been raised in the technical literature on privacy since the late 1990s. Researchers have repeatedly demonstrated that it is possible to re-identify data believed to be de-identified. In many cases, re-identification has been accomplished via successful linkage attacks, whereby a de-identified dataset is joined with a publicly available dataset containing identifying information. Any record in the de-identified dataset which is uniquely linked with a record in the publicly available dataset (using the attributes common to both) is re-identified. For example, Sweeney demonstrated that de-identified hospital records could be joined with voter registration records, as both datasets contained birth date, sex, and ZIP code information for each individual. Sweeney used these three attributes to uniquely identify Massachusetts Governor William Weld’s health records and observed that the process she used resulted in many other uniquely identified records [31]. More generally, she showed that a large portion of the US population could be uniquely re-identified given just these three pieces of their information [32]. Subsequent attacks on de-identified data (e.g., [33, 34]) suggest that as few as three or four data points can be sufficient to re-identify an individual in a de-identified dataset. In fact, to be effective, de-identification must strip the data of most of its informational content, rendering it almost void of value of analytic purposes (see discussion in [35, 36]). There have been many attempts to define the concept of de-identification heuristically (e.g., k -anonymity [37] and its variants l -diversity [38] and t -closeness [39]). However, as researchers continually discover new privacy vulnerabilities, new heuristic definitions must be developed to address them. In practice, heuristic approaches can be used to protect personal information against a small number of specific attacks, but they do not provide comprehensive protection against all feasible attacks.

To date, there is no formal mathematical definition for de-identification, or for related concepts such as personally identifiable information and linkage. Because these concepts are heuristic rather than formal, they must be periodically updated in response to newly discovered weaknesses. Privacy standards based on these heuristic concepts become moving targets, creating uncertainty for practitioners (i.e., whether the techniques they use provide sufficient privacy protection) and for the individuals in the data (i.e., whether and when data publications expose them to risks). Note that although the heuristic approaches mentioned above have been defined using mathematical language, they are syntactic in nature (i.e., specifying properties of how an anonymized dataset should look) rather than semantic (i.e., specifying restrictions on what an attacker may infer about the personal information that is the input for the anonymization process by observing its outcome). Consequently, these notions do not provide a precise understanding of the privacy protection pro-

vided, such as how the potential harm to an individual may be affected by being included in a k -anonymized, l -diverse, or t -close dataset. The lack of formal understanding limits the scope of scientific discussion of de-identification to, roughly, producing a collection of techniques, the efficacy of which are difficult or even impossible to measure and compare.

Semantic security: The foregoing discussion refers to the advantages of semantic over syntactic definitions of privacy. A key example of a semantic definition is *semantic security*, a definition that was introduced by Goldwasser and Micali [40]. Semantic security is a standard privacy requirement of encryption schemes. To understand what it requires, consider a scenario in which Alice uses a public-key encryption scheme to communicate a confidential message m to Bob. She encrypts the message using Bob’s encryption key and sends him the resulting ciphertext c . Using his (secret) decryption key, Bob can then recover the message m from the ciphertext c . The definition of semantic security compares what an attacker (without access to Bob’s decryption key) can predict about the message m given the ciphertext c with what the attacker can predict about the message m without being given the ciphertext c . The advantage that access to the ciphertext gives to any attacker is quantified. Encryption schemes are designed to make this advantage so negligible that access to the ciphertext does not give the attacker any practical advantage over not getting any information about the message m at all.

An early influential work on data privacy by the statistician Tore Dalenius defined disclosure as follows. “If the release of the statistics S makes it possible to determine the value D_k more accurately than is possible without access to S , a disclosure has taken place” (here D_k refers to the personal information of subject k) [41]. This view of disclosure is interpreted as having a goal similar to semantic security: “access to a statistical database should not enable one to learn anything about an individual that could not be learned without access” [42].¹

Note that, while not fully formalized, this desiderata is essentially equivalent to semantic security. There are, however, significant differences between applying this concept to encryption and to data privacy. In particular, the setting of encryption schemes clearly distinguishes between the party who should be able to learn the message m (i.e., Bob) and an eavesdropping attacker who should not gain any information about m . In contrast, when statistics are computed over a collection of personal data, the analyst (understood broadly as any party with access to the published outcome) is both the proper consumer of the statistics and a potential privacy attacker.

Semantic security has proved to be a fundamental concept for encryption. In fact, encryption schemes that are in common use today are semantically secure (under some mathematical assumptions). However, Dwork and Naor [42] demonstrated that the concept cannot be applied in the context of private data analysis as it would imply (reusing Dalenius’ words in [41]) “elimination of statistics.” To understand what this means, consider an example, taken from [42]. Suppose the attacker has the auxiliary information “Terry Gross is two inches shorter than the average Lithuanian woman.” Without access to statistical information about Lithuanian women, the attacker has little information about Terry Gross’ (secret) height, but the situation changes dramatically once statistics about the average heights of the Lithuanian population are published, as the attacker can now combine the published average height with his auxiliary information to learn Gross’ height. The release of the average height of Lithuanian women (whether exact or an approximation)

¹Dalenius recognizes that “[a] reasonable starting point is to discard the notion of elimination of disclosure” as “it may be argued that elimination of disclosure is only possible by elimination of statistics” [41]. The argument by Dwork and Naor still holds even if their interpretation is relaxed to allowing an attacker to obtain a small advantage in predicting private personal information when given access to the statistical database.

does not satisfy semantic security as it allows some attackers (those who possess relevant auxiliary knowledge) to improve their predictions of Gross' height. This illustrates how the semantic security desiderata fails to address access to auxiliary information, i.e., information available to an attacker from outside the system.

A second example illustrates how the semantic security desiderata fails to differentiate between information about a person and information that is specific to a person (i.e., information that cannot be inferred unless his or her information were used in the analysis). Suppose a research study explores the effects of wine drinking in 60-year-old women and finds a strong positive correlation between wine drinking and kidney disease that was not known to exist before the study was conducted. If one knows that Gertrude is a 60-year-old woman and that she is a regular wine drinker, the results of the study can be used to better estimate whether Gertrude suffers from kidney disease than before the study was conducted. The release of the study's results therefore does not satisfy semantic security. Further research extends these examples and generalizes them in a formal mathematical argument that making semantic security a requirement in data analysis would essentially prohibit the discovery of any new knowledge through research using personal data [42].

The discussion above should not be understood to undermine the importance of cryptographic notions such as semantic security to privacy. In fact, modern cryptographic methods such as secure multiparty computation [43] and homomorphic encryption [44] extend the notion of semantic security to settings in which several parties (each holding a private input) compute a function $f()$ of their joint data without having to reveal any information beyond what is revealed at the conclusion of the following "ideal" (imaginary) process: a fully trusted party collects the data from all the parties, computes $f()$, informs each party of her designated part of $f()$, and erases the entire process from its memory. Such cryptographic approaches are now beginning to be deployed in real-world settings, such as personalized medicine [45].

Formal privacy models and differential privacy: Failures of traditional privacy-preserving approaches to control disclosure risks in statistical publications have motivated computer scientists to develop a strong, formal approach to privacy. The formal study of privacy has grown successfully out of the study of cryptography, beginning with the seminal work of Diffie and Hellman [46]. This work helped to establish rigorous mathematical foundations for the field of cryptography. One of the benefits of this approach was an ability to avoid the "penetrate-and-patch" dynamics of older cryptographic schemes whose design had to be repeatedly modified to counter newly found attacks. New cryptographic schemes could be designed to be provably resilient to any feasible adversarial attack. For this reason, it is now extremely rare that a security system is vulnerable because of a weakness in a cryptographic algorithm (though vulnerabilities may nevertheless exist in the implementation of the algorithm or in other components of the system).

The formal approach to privacy includes both a study of what cannot be computed with any reasonable notion of privacy and what can be computed with strong notions of privacy, with the goal of closing the gaps between the two. At the center of this approach stands a collection of formal concepts defined with mathematical precision. Furthermore, statements about these concepts are proved mathematically (rather than, say, empirically). For example, research seeks to develop a formal notion of what should be considered non-private under any reasonable notion of privacy. This very minimal view of privacy, called a *reconstruction attack*, corresponds to a setting in which a privacy attacker can use the database's purportedly privacy-preserving mechanism to accurately reconstruct the database. Dinur and Nissim [47] used reconstruction attacks to prove a lower bound on the magnitude of noise that is essential for preserving privacy in the context of a

database answering multiple statistical queries. In this setting, each of the statistical queries seems to leak very little personal information; however, the accumulated leakage amounts to a massive loss of privacy. This phenomenon was shown to hold in other settings, indicating that if *too many* statistics are released *with too high accuracy* then privacy is lost.

Central to the current study of formal privacy models is the notion of *differential privacy*, introduced by Dwork, McSherry, Nissim, and Smith in 2006 [48]. Differential privacy is a formal mathematical standard for quantifying and managing privacy risk. The definition requires the output distribution of a privacy preserving analysis to remain “stable” under any possible change to a single individual’s information. Differential privacy has a compelling intuitive interpretation as it guarantees to every individual that the consequences to her would be similar regardless of whether her information were used in the analysis or not. We refer the reader to [51] for a more detailed discussion of the protection provided by differential privacy and to [48] for a more technical presentation of the mathematical definition and its properties.

It is worth emphasizing that differential privacy is not a specific tool or technique for privacy protection but a *definition* or *standard* for quantifying and managing privacy risks. A range of technological tools can be devised to satisfy the differential privacy standard. It differs from all notions of anonymization mentioned above in that it restricts the informational relationships between the personal information that is used as input to a privacy-preserving analysis and the outcome of the privacy-preserving analysis. The excess risk to an individual resulting from her participation in a differentially private analysis is bounded (or strictly controlled), whereas such a guarantee is impossible for the anonymity concepts. Another key property of differential privacy is that it self-composes. Any mechanism resulting from the composition of two or more differentially private mechanisms is also differentially private (albeit, with worse parameters). Composition is a property that prior attempts to define privacy lacked. Currently, differential privacy is also the only framework giving meaningful privacy guarantees in the face of adversaries having access to arbitrary external information. Furthermore, analyses satisfying differential privacy provide provable privacy protection against any feasible adversarial attack, whereas the anonymity concepts only counter a limited set of specific attacks.

Comparing differential privacy with semantic security, the crucial difference is that semantic security compares what an attacker can infer about an individual with and without access to the statistics disclosed, whereas differential privacy compares what an attacker that has access to the statistics can infer about an individual whether her information is used in computing the statistics or not. Recall that the analysis of Dwork and Naor [42] showed that semantic security entails no utility in data analysis. In sharp contrast, there is a continually growing list of tasks that has been shown, in principle, to be computable with differential privacy, including descriptive and inferential statistics, machine learning algorithms, and production of synthetic data. Existing real-world applications of differentially private analyses include implementations by federal agencies such as the US Census Bureau and companies such as Google, Apple, and Uber.

In settings in which one seeks to analyze data that are distributed among several agencies but cannot be shared among themselves or with a data curator, differential privacy may be implemented with cryptographic techniques, such as secure multiparty computation, for computing over the distributed data. In some cases, integrating differential privacy with cryptographic techniques can yield greater accuracy at the same level of privacy.

2.3 Gaps between normative and technical concepts

There are substantial gaps between technical and normative conceptions of privacy, and these gaps are magnified by the emergence of formal privacy models. Chiefly, formal privacy models such as differential privacy are defined with mathematical precision, and statements about whether their requirements are satisfied can be proved rigorously. Uncertainty with respect to the correctness of such statements is eliminated once their proofs are verified. In contrast, normative approaches to privacy are inherently flexible and subject to interpretation. Assessing privacy with respect to normative expectations can lead to varying—and even contradictory—conclusions. Practical guidance on interpreting normative concepts often, by design, leaves the boundaries between what is private and what is non-private not fully determined. On one hand, this allows interpretation of the standard to evolve, and, on the other hand, creates uncertainty. The following discussion illustrates a number of the gaps created by these differences, through a comparison of formal privacy models like differential privacy to the normative approaches reflected in information privacy laws.

Generality of protection afforded: Regulatory requirements for privacy protection vary according to industry sector, jurisdiction, institution, types of information involved, or other contextual factors [22]. For example, the Family Educational Rights and Privacy Act protects only certain types of information contained in education records maintained by schools, universities, and educational agencies. However, in practice, privacy risks are not limited solely to the information categories and contexts contemplated in the law. Furthermore, interpreting and applying regulatory standards is challenging in cases in which an analyst seeks to combine data from multiple sources. In contrast to regulatory requirements, a formal privacy model like differential privacy offers general protection and can, in principle, be applied wherever statistical or machine learning analysis is performed on collections of personal information, regardless of the contextual factors at play. In recognition of the fact that a broader notion of personal information must be protected in order to preserve privacy and that information that currently seems innocuous may prove sensitive in the future, formal privacy models protect all information specific to an individual, not just information traditionally considered to be identifiable or able to be linked to an individual.

Scope of attacks contemplated: Privacy regulations and related guidance contemplate a limited set of specific attacks and privacy failure modes. As one example, many regulations make an implicit assumption that re-identification via record linkage—i.e., the re-identification of one or more records in a de-identified dataset by uniquely linking these records with identified records in a publicly available dataset—is the primary or sole privacy failure mode. Other central concepts appearing in privacy regulations, including personally identifiable information, (de-)identification, linkage, and inference, are often defined from this point of view. For example, many privacy regulations require data providers to protect information that can be linked to an individual in order to safeguard against record linkage. As a result, these requirements are often interpreted as requiring the protection of information one can foresee being used in a record linkage attack.

However, in the last two decades, researchers have identified new attacks and privacy failure modes. In some cases, learning whether an individual participated in a research study could be considered a privacy violation, even if the individual’s exact information cannot be identified. For instance, if an employer learns that a job candidate previously participated in a research study investigating the efficacy of interventions for substance abuse, he or she may infer that the candidate has a history of substance abuse, even without identifying the candidate’s record in the database. Privacy risks can take other forms as well, such as singling out an individual (even if not

fully identified), or inferring information that is specific to an individual with less than absolute certainty. Privacy regulations that focus on re-identification via record leakage can fail to address this broader understanding of the potential modes of privacy failure.

In contrast with existing privacy regulations, formal privacy models provide protection against a wide collection of privacy attacks, even those that are not currently known. More specifically, formal models focus on provable limitations to excess harm due to participation in a computation, regardless of the failure mode that occurred.

Expectations vs. the scientific understanding: Regulatory standards that rely on the concept of de-identification to protect privacy are often not in agreement with the current scientific understanding of privacy. For example, the HIPAA Privacy Rule allows the publication of personal health records, as long as certain pieces of information deemed to be identifying have been removed [24]. This is now understood to be a weak standard, as research has demonstrated that redaction of identifiers can fail to protect privacy, especially when applied to information that is very detailed, such as that found in medical records. In fact, any information about individuals, including information not traditionally considered to be identifying, has the potential to leak information specific to individuals [36]. Moreover, this issue is not limited to HIPAA, as many legal standards of privacy rely on the concepts of de-identification and personally identifiable information.

In some cases, the law may be interpreted to require something that is not technically feasible, such as absolute privacy protection when sharing personal data. For example, Title 13 of the US Code protects the confidentiality of respondent information collected by the US Census Bureau by prohibiting “any publication whereby the data furnished by . . . [an] individual . . . can be identified” [5]. Whether an individual can be identified in a publication is not precisely defined. If this concept were interpreted very conservatively, Title 13 would disallow any leakage of information about individuals. This, in turn, would prohibit the Census Bureau from publishing any statistics based on data furnished by individuals, as every release of statistics unavoidably implies some level of privacy loss for the individuals whose information has been analyzed.

The binary view of privacy found in Title 13—whereby information is either identifiable or not—is common to many regulations. Such an approach is problematic (1) because information can never be made completely non-identifiable and (2) because it fails to recognize that privacy loss accumulates with successive releases of information about the same individuals. Not only is some privacy loss inevitable in every release of statistics, but these leakages accumulate and can eventually amount to a significant disclosure of personal information. Formal privacy models such as differential privacy bound the privacy leakage of each release, and furthermore, are equipped with a host of tools (called “composition theorems”) that bound the total privacy leakage across multiple releases. However, composition, a key feature of formal privacy models, is generally ignored by regulatory standards for privacy protection.

(In)-stability over time: Notions of privacy embedded within regulatory standards are continually evolving in response to new discoveries of vulnerabilities. Practitioners seeking to implement privacy safeguards in accordance with these regulations therefore face a moving target. As one example, the US Office of Management and Budget’s guidance on protecting personally identifiable information has been updated over time to address an evolving understanding of the ways in which de-identified data may be vulnerable to potential attacks. The latest update to the guidance advises government agencies that they must consider that non-personally identifiable information may become personally identifiable information in the future [6]. This approach to defining the scope of information to be protected requires regulatory and policy standards to be updated as new attacks

are identified, much like the “penetrate-and-patch” approach to patching software incrementally as new bugs are discovered.

However, when a regulatory definition of privacy is periodically amended over time, it is an indication that the definition is not a strong, general definition of privacy. Rather, the definition reflects a much narrower presumption that certain approaches are likely sufficient to provide adequate privacy protection. Regulations such as the HIPAA Privacy Rule [24] that require the removal of certain pieces of identifying information have hardwired a specific technique—redaction of identifiers—into their standards. Over time, as techniques such as redaction are shown to be inadequate to protect privacy, practitioners are left with uncertainty regarding what is required to satisfy regulatory standards that are out of step with best practice.

In contrast, a formal model like differential privacy is the subject of ongoing scientific research, *regardless of implementation*. This research provides a strong assurance that differential privacy provides a sufficient level of privacy in an extremely wide collection of settings. This assurance is supported by mathematical theory providing provable privacy guarantees for any combination of differentially private analyses, provided they are correctly implemented.²

Relationship to normative expectations: On first impression, the relationship between some technical and normative concepts of privacy may appear to be straightforward. For example, it may seem to follow intuitively that differential privacy satisfies the requirements of many regulatory requirements, as well as many of the normative expectations underlying such requirements. This is because, in many cases, differential privacy provides protection that is more robust than that provided by traditional statistical disclosure limitation techniques commonly used to satisfy regulatory requirements for privacy protection. Moreover, because they are founded on rigorous mathematical grounds, formal privacy models provide protection against a wide range of potential privacy attacks, including attacks that have succeeded against traditional techniques.

We observe, however, that there are significant conceptual gaps between formal privacy models and normative standards and expectations of privacy. We caution that an analysis comparing the protection afforded by a particular privacy technology to a normative standard must be done with care. For illustration, we provide a few examples demonstrating why making a sufficiency claim with respect to differential privacy and legal requirements for privacy is a non-trivial task.

First, formal privacy models are, by and large, “privacy-first” definitions. These definitions make the privacy desiderata primary, and they subject other desiderata, such as accuracy, computational costs, and sample complexity, to the limitations implied by the privacy guarantee. In many practical settings, uses of data may demand a compromise between protecting privacy and carrying out accurate computations on the data, and finding the right balance can be challenging. For instance, some existing real-world applications of differential privacy have been implemented with parameters that were selected to improve accuracy but, in principle, may allow for a significant leakage of personal information [50].

Second, some regulatory standards express requirements for privacy protection that can be interpreted to exceed the protection provided by various privacy technologies, including formal privacy models like differential privacy. For example, the US Census Bureau has an obligation to protect the privacy of both individual respondents and establishments. Haney et al. observed that, while differential privacy is a notion suitable for the protection of information pertaining to individuals or small groups of individuals, such as a family, it does not necessarily protect establishments [49].

²We note that variants of differential privacy are the subject of research. The goal of these variants is to provide improved accuracy while providing provable privacy guarantees similar to those of differential privacy.

Their research offers suggestions for modifying differential privacy with additional requirements to prevent precise inferences of establishment size and the composition of the establishment workforce [49]. This real-world case illustrates how normative requirements, such as those embodied in the law, can direct the development and implementation of formal privacy models.

Lastly, as we discussed above, differential privacy protects information that is specific to a data subject—i.e., information that can only be inferred about the subject if his or her information is used in the analysis. However, as discussed earlier in this section, some regulations can be interpreted to require the protection of data that is not specific to an individual. This poses a research direction for the formal mathematical study of privacy: to understand formally which part of this expectation can be met with a rigorous mathematical definition, and the implications of such a definition for analyses performed in a variety of tasks, including both research- and commercially-focused analyses.

3 A way forward: Hybrid concepts of privacy

Adopting an understanding of privacy that is consistent across its technical and normative dimensions will be critical to ensuring personal data are adequately safeguarded over the long term. However, significant conceptual gaps between existing technical and normative concepts create challenges for arriving at a universal notion of privacy. For instance, normative concepts embedded in existing regulatory requirements for privacy protection often rely on intuitive assumptions about how pieces of information interact, rather than (and often contradicting) scientific and mathematical principles. Framing privacy in this way can result in expressions of unrealistic privacy desiderata, leading practitioners to pursue an idealized privacy goal that is impossible to achieve. At the same time, purely technical approaches may adopt a narrow view of privacy that fails to capture the fundamental normative expectations of privacy.

Bridging these gaps will be necessary to ensure robust privacy protection in practice. A significant first step is to recognize that privacy concepts have a hybrid nature. They are neither purely legal nor purely technical, but rather a multi-dimensional combination of the two. Understanding the hybrid nature of these concepts and developing tools for implementing them is necessary to bridge the gaps between the current legal and technical understanding of privacy.

Contextual integrity: Nissenbaum’s seminal framework—contextual integrity [26]—is an approach that aims to capture the dual technical-normative nature of privacy. Contextual integrity is a justificatory framework for privacy. It states that privacy breaches can be tracked to violations of societal norms regarding appropriate information flows within a particular context. Where it is possible to encode norms formally and determine whether a particular information flow respects these norms, i.e., when the norms are unambiguous and furthermore effectively testable, then the framework can be used to precisely predict or flag violations of the societal norms. As an example, Barth et al. [28] provide a model for encoding norms appearing in legal standards such as HIPAA.³

Normative concepts are often not defined explicitly, and, when they are, they are not expressed in a formal language that enables a precise analysis. As a result, there is uncertainty with respect to which information flows are in agreement with normative concepts. In many cases, it may be difficult to determine with reasonable certainty whether an information flow is appropriate, whether it creates a risk of a privacy breach, or even whether a privacy breach has in fact occurred.

³Technically, this model is based on first-order logic (allowing predicates as well as quantifying over variables) extended with temporal operators that allow expressing conditions on the timing relationships between events.

In order to bridge normative and technical concepts, the framework of contextual integrity could in the future be equipped with formal mathematical definitions. If defined formally, its fundamental concepts, i.e., context, information flow, and norm, could provide an interface for integrating formal privacy models into the framework. As an example, the formalization of context, information flow, and norm would allow researchers and policymakers to reason about the appropriateness of using differentially private analyses in various contexts. This reasoning would involve making and proving quantifiable statements regarding the extent to which differential privacy respects social norms around privacy in different contexts.

Bridging the legal and technical perspectives: Motivated by the lack of translational work to bridge the legal and technical perspectives of privacy, we, together with our colleagues, have sought in prior work to develop an approach to demonstrating that a particular privacy technology can be used to satisfy a regulatory standard for privacy protection. As an example, we applied the proposed approach to demonstrate that differential privacy can be used to satisfy the de-identification requirements of the Family Educational Rights and Privacy Act (FERPA). This approach involved identifying a description of a potential privacy attacker and the attacker’s goals, which were embedded within FERPA’s definition of personally identifiable information. With a legal analysis of this regulatory language, we developed a detailed description of the potential attacker contemplated by FERPA. We then developed a mathematical model of FERPA’s privacy requirements and made conservative assumptions in the model with the goal of accounting for possible ambiguities in the regulatory standard. Finally, we analyzed differential privacy with respect to the mathematical model extracted from FERPA’s privacy requirements [52].

Efforts to extend this work to other laws, such as the law protecting the confidentiality of information the US Census Bureau collects from respondents, led to the observation that other laws do not describe a potential attacker or the attacker’s knowledge in the way that FERPA does. This makes it difficult to generalize the specific approach that was used to model FERPA. Despite this challenge, there is an alternative approach that may provide a solution to the modeling problem. This approach involves identifying fundamental concepts used in regulatory standards for privacy protection and the statistical disclosure limitation literature and performing a detailed legal analysis of these concepts. Based on this analysis, one can model the concept mathematically, and then check whether the modeling agrees with the legal analysis. If it does, then one can proceed by comparing the mathematical model of the concept to a technical definition of privacy in order to demonstrate whether the technical definition meets the definition of the privacy concept extracted from the regulation or literature. We suggest that such an approach can be used to design *hybrid concepts of privacy*, i.e., concepts that can be interpreted and used consistently technically and normatively.

Future regulation: Developing an understanding of the gaps between technical and normative approaches to privacy can also help identify ways in which to improve future privacy regulations. We argue that future regulations should aim to articulate clear goals for privacy protection, that are in line with the scientific understanding of privacy, rather than implicitly or explicitly endorsing heuristic de-identification techniques. As an example of a recent attempt to bring greater clarity to regulatory requirements by explaining the goals of privacy protection, consider the EU’s General Data Protection Regulation, which protects personal data but not anonymous data [7]. The regulation does not state a clear goal, making it difficult to interpret what exactly it is intended to protect. However, guidance on interpreting the regulation outlines goals that go beyond the traditional notion of de-identification, namely, protection from singling out, linking, or inferring

an individual’s personal data from a dataset [53]. Although these privacy concepts have not yet been defined precisely and formally from a mathematical perspective, they aim to describe what the goal of a privacy-preserving mechanism should be, rather than prescribing a specific family of techniques.

Bridging legal and other normative perspectives: It is also important to acknowledge that the normative concepts embedded within privacy regulations may not be a perfect mirror of more fundamental normative concepts of privacy. For instance, the choices made in the design of privacy regulations may reflect expedient political or practical compromises, rather than actual individual and societal expectations. It may in fact not be possible to close the gap between legal and other more fundamental normative concepts. However, by learning how to analyze whether a technology satisfies a legal definition, it may also be possible to develop analogous approaches for analyzing whether a technology satisfies other normative expectations of privacy, including those that serve as the motivation for regulatory standards for privacy protection.

Acknowledgment The authors are supported by the US Census Bureau under cooperative agreement no. CB16ADR0160001. The authors thank Simson Garfinkel and the editor and reviewers of *Philosophical Transactions of the Royal Society A* for their helpful comments.

References

- [1] J. M. Cohen. *What Privacy Is For*. Harvard Law Review 126(7): 1904-1933. 2013.
- [2] D. J. Solove. *‘I’ve Got Nothing to Hide’ and Other Misunderstandings of Privacy*. San Diego Law Review 44(4): 745-772. 2007.
- [3] Family Educational Rights and Privacy Act of 1974, 20 U.S.C. § 1232g; 34 C.F.R. Part 99 (2013).
- [4] Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule, 45 C.F.R. Part 160 and Subparts A and E of Part 164.
- [5] 13 U.S.C. § 9. Information as confidential; exception.
- [6] Office of Management and Budget, Memorandum M-17-12. *Preparing for and Responding to a Breach of Personally Identifiable Information*. January 3, 2017.
- [7] Regulation of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data, COM (2012) 11 final (Jan. 25, 2012).
- [8] G. Gatewood, *Census Confidentiality and Privacy: 1790–2002*. Monograph, 2001.
- [9] J. M. Abowd. *Why the Census Bureau Adopted Differential Privacy for the 2020 Census of Population*. Presentation at Harvard University, December 11, 2017.
- [10] J. A. Swanson. *The Public and Private in Aristotle’s Political Philosophy*. 1992.
- [11] D. J. Solove. *A Taxonomy of Privacy*. University of Pennsylvania Law Review 154(3): 477-560, 2006.

- [12] S. Warren and L. Brandeis. *The Right to Privacy*. Harvard Law Review 4:193 (1890).
- [13] W. L. Prosser. *Privacy*. California Law Review 48:383 (1960).
- [14] D. J. Solove and D. K. Citron. *Risk and Anxiety: A Theory of Data Breach Harms*. Texas Law Review 96 (2018): 737-786.
- [15] US Dept. of Health, Education, & Welfare, Records, Computers, and the Rights of Citizens: Report of the Secretary's Advisory Committee on Automated Personal Data Systems (1973).
- [16] Privacy Act of 1974, 5 U.S.C. § 552a.
- [17] Mass. Gen. Laws ch. 66A.
- [18] Minn. Stat. § 13.01 et seq.
- [19] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (Data Protection Directive).
- [20] Organisation for Economic Co-operation and Development. *The OECD Privacy Framework*. 2013.
- [21] R. Gellman. *Fair Information Practices: A Basic History*. Working Paper, Version 2.18, 2017.
- [22] P. M. Schwartz and D. J. Solove. *The PII Problem: Privacy and a New Concept of Personally Identifiable Information*. 86 N.Y.U. L. Rev. 1814 (2011).
- [23] 201 Code Mass. Regs. § 17.02.
- [24] 45 C.F.R. § 164.514.
- [25] Family Education Rights and Privacy, 73 Fed. Reg. 74,805, 74,853 (Dec. 9, 2008).
- [26] H. Nissenbaum, *Privacy In Context: Technology, Policy, and the Integrity of Social Life* (Stanford University Press, 2009).
- [27] Obama Administration, Administration Discussion Draft: Consumer Privacy Bill of Rights Act of 2015 (2015).
- [28] A. Barth, A. Datta, J. Mitchell, and H. Nissenbaum. *Privacy and Contextual Integrity: Framework and Applications*. Proc. of the 2006 IEEE Symp. on Security and Privacy, pp. 184–198, 2006.
- [29] N. R. Adam and J. C. Worthmann. *Security-control methods for statistical databases: a comparative study*. ACM Computing Surveys (CSUR) Vol. 21, no. 4, pp. 515–556, 1989.
- [30] Federal Committee on Statistical Methodology. *Report on Statistical Disclosure Limitation Methodology*. Statistical Policy Working Paper 22, 2005.
- [31] Recommendations to Identify and Combat Privacy Problems in the Commonwealth: Hearing on H.R. 351 Before the House Select Committee on Information Security, 189th Sess. (Pa. 2005) (statement of Latanya Sweeney, Associate Professor, Carnegie Mellon University).

- [32] L. Sweeney. *Uniqueness of Simple Demographics in the U.S. Population*. Laboratory for Int'l Data Privacy, Working Paper LIDAP-WP4, 2000.
- [33] A. Narayanan and V. Shmatikov. *Robust De-anonymization of Large Sparse Datasets*. Proceedings of the 2008 IEEE Symposium on Research in Security and Privacy, p. 111, 2008.
- [34] Y. A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel. *Unique in the Crowd: The Privacy Bounds of Human Mobility*, 3 NATURE SCI. REP. 1376 (2013).
- [35] P. Ohm. *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*. UCLA Law Review, Vol. 57, p. 1701, 2010.
- [36] A. Narayanan and V. Shmatikov. *Myths and Fallacies of "Personally Identifiable Information"*. Communications of the ACM 24, 26, 2010.
- [37] L. Sweeney. *k-anonymity: A model for protecting privacy*. International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, 10 (5), 2002; 557-570.
- [38] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. *l-diversity: Privacy Beyond K-anonymity*. ACM Trans. Knowl. Discov. Data. 1 (1), March 2007.
- [39] N. Li, T. Li, and S. Venkatasubramanian. *t-Closeness: Privacy Beyond k-Anonymity and l-Diversity*. Proceedings of the 23rd International Conference on Data Engineering, ICDE, pp. 106–115, 2007.
- [40] S. Goldwasser and S. Micali. *Probabilistic encryption*. Journal of Computer and System Sciences, Vol. 28, Is. 2, pp. 270–299, 1984.
- [41] T. Dalenius. *Towards a methodology for statistical disclosure control*. Statistik Tidskrift 15, pp. 429–444, 1977.
- [42] C. Dwork and M. Naor. *On the Difficulties of Disclosure Prevention in Statistical Databases or The Case for Differential Privacy*. Journal of Privacy and Confidentiality, vol. 2, no. 1, Sep. 2010.
- [43] Y. Lindell, and B. Pinkas. *Secure Multiparty Computation for Privacy-Preserving Data Mining*. Journal of Privacy and Confidentiality, Vol. 1, No. 1, 2009.
- [44] C. Gentry. *Computing arbitrary functions of encrypted data*. Communications of the ACM, 53(3), pp.97-105, 2010.
- [45] E. Hayden. *Extreme cryptography paves way to personalized medicine*. Nature News, 519(7544), p.400, 2015.
- [46] W. Diffie and M. Hellman. *New Directions in Cryptography*. IEEE Trans. on Information Theory, Vol. IT-22, No. 6, Nov. 1976.
- [47] I. Dinur and K. Nissim. *Revealing information while preserving privacy*. Proceedings of the 22nd symposium on Principles of database systems (PODS), pp. 202-210, 2003.
- [48] C. Dwork, F. McSherry, K. Nissim, and A. Smith. *Calibrating Noise to Sensitivity in Private Data Analysis*. Journal of Privacy and Confidentiality, vol. 7, no. 3, pp. 17-51, May 2017.

- [49] S. Haney, A. Machanavajjhala, J. M. Abowd, M. Graham, M. Kutzbach, and L. Vilhuber. *Utility Cost of Formal Privacy for Releasing National Employer-Employee Statistics*. SIGMOD Conference pp. 1339-1354, 2017.
- [50] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang. *Privacy Loss in Apple's Implementation of Differential Privacy on MacOS 10.12*. Working paper, 2017.
- [51] K. Nissim, T. Steinke, A. Wood, M. Altman, A. Bembenek, M. Bun, M. Gaboardi, D. R. O'Brien, and S. Vadhan. *Differential Privacy: A Primer for a Non-technical Audience*. Vanderbilt Journal of Entertainment & Technology Law, 2018 forthcoming.
- [52] K. Nissim, A. Bembenek, A. Wood, M. Bun, M. Gaboardi, U. Gasser, D. R. O'Brien, T. Steinke, and S. Vadhan. *Bridging the Gap between Computer Science and Legal Approaches to Privacy*. Harvard Journal of Law & Technology, vol. 31, no. 2, pp. 687-780, 2018.
- [53] Article 29 Data Protection Working Party, Opinion 05/2014 on Anonymisation Techniques (2014).