Rethinking the Proposition of Privacy Engineering

Aaron Ceross
Department of Computer Science,
University of Oxford
Oxford, United Kingdom
aaron.ceross@cs.ox.ac.uk

Andrew Simpson
Department of Computer Science,
University of Oxford
Oxford, United Kingdom
andrew.simpson@cs.ox.ac.uk

ABSTRACT

The field of privacy engineering proposes a methodological framework for designing privacy-protecting information systems. Recognising that the utilisation of privacy-enhancing techniques for data storage and analysis does not address the entire scope of individual privacy, privacy engineering incorporates influences from user sentiment, legal norms and risk analysis in order to provide a holistic approach. Framed by related design principles, such as 'Privacy-by-Design', privacy engineering purports to provide a practical, deployable set of methods by which to achieve such a holistic outcome. Yet, despite this aim, there have been difficulties in adequately articulating the value proposition of privacy engineering. Without being able to adequately define privacy or map its contours, any proposed methodology or framework will be difficult to implement in practice, if not self-defeating. This paper identifies and examines the assumptions that underpin privacy engineering, linking them to shortcomings and open questions. Further, we explore possible research avenues that may give rise to alternative frameworks.

CCS CONCEPTS

• **Security and privacy → Social aspects of security and privacy**; *Economics of security and privacy*; Usability in security and privacy;

ACM Reference Format:
Aaron Ceross and Andrew Simpson. 2018. Rethinking the Proposition of Privacy Engineering. In *New Security Paradigms Workshop (NSPW '18), August 28–31, 2018, Windsor, United Kingdom.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3285002.3285006

1 INTRODUCTION

Prior to the 20th Century, the concept of privacy was an ever-present, yet secondary feature of social, political and legal thought. Today, privacy finds itself pulled from the periphery of philosophical thought [73] into the centre of debate, driven by swift developments in information technology and the consequent rapacious thirst for all manner of personal data. A result of this development has been an encroachment of systems governing aspects of social interaction where previously there were none, curtailing autonomy on how to manage one's own information.

For more than a century, the flow of data has resulted in the diminishing of what may be regarded as 'private', as this domain blurs with increasingly digital 'public' interactions, or, at the very least, what may have been conceived as such. This phenomenon, driven in part by demands of users, as well as system designers, has been recognised, described, and often lamented as an 'erosion', identifying that some fundamental transformation is occurring [53, 76, 100]. The last 50 years or so have been increasingly punctuated by policy means to halt, or at least slow, the negative effects alleged to be caused by the incorporation of information systems into virtually all aspects of individuals' lives. All the while, the original question of what it means to be private remains open. This unanswered question acts as a foundational flaw, serving only to undercut any proposed solution.

Within information system design, privacy engineering purports to provide a way out of the impasse posed by individual privacy concerns and legal obligations, offering to supplement information systems engineering to address distinct concerns. While there is consensus that privacy is a multifaceted, context-dependent concept, this quality does not seem to find sufficiently adequate articulation in modern systems.

In this paper we explore the promise of privacy engineering, identifying its goals and purpose, as well as the emergent challenges and weaknesses. Furthermore, we argue that these weaknesses are, in part, based in assumptions that have become embedded within the very proposition of privacy engineering. This is arguably traced to privacy-centric design ideals, most notably Privacy-by-Design [29], which puts forth "design principles" for ensuring informational privacy within a system. The milieu of assumptions, proclamation of principles, and recent legislative efforts have attempted to harmonise these disparate elements, without considering the fundamental goals of what privacy engineering ought to provide. This has driven research into directions that, arguably, divert attention away from holistic address of holistic privacy within information systems, focusing instead on regulatory compliance. We therefore identify and challenge these assumptions, in order to allow for the development of new methods and paradigms for reasoning about privacy within information systems.

2 BACKGROUND AND MOTIVATION

We first provide an overview, as well as definitions, to help frame our contribution. When discussing privacy, there are a number of overlapping concepts and terms that have been used interchangeably within the literature, which we argue causes confusion and misunderstandings. Briefly, we will distinguish privacy from these related, if sometimes overlapping, concepts.

NSPW '18, August 28–31, 2018, Windsor, United Kingdom

ACM ISBN 978-1-4503-6597-0/18/08…$15.00
https://doi.org/10.1145/3285002.3285006

- *Anonymity.* Anonymity is a property of identity concealment, which overlaps with privacy, but remains sufficiently distinct from it [83, 86]. The goal of anonymity is to obscure or entirely remove data that may directly or indirectly identity an individual in relation to monitoring of activity; the activity is recognised, but the identity of the individual is unknown. Díaz *et al.* [37] describe two types of anonymity: (i) data anonymity, which involves the removal of identifiers from a data subject and (ii) connection anonymity, wherein the source and destination of a data transfer is concealed or otherwise obscured. Anonymity is somewhat of a spectrum, given there are degrees of anonymity, or pseudonymous states [75].
- *Secrecy.* Secrecy is the desired hiding of information through active concealment. This intention manifests such that individual makes use of methods to block or otherwise obscure information. Warren and Laslett [117] draw distinction between the 'moral content' of privacy and secrecy, arguing that the former is consensual, whereas the latter is not. In a similar manner, Bok [23] emphasises the concealing nature of secrecy, where private matters are not necessarily hidden.
- *Confidentiality.* Confidentiality is predicated on a relationship between two or more individuals; the information is not for circulation to a wider audience and available only to authorised parties. Information that is acquired within that relationship is deemed confidential. Information security often refers to confidentiality as a core component of its triad of principles,[1] as it relates to the access control of the information, rather than the substantive nature of the information itself.

What, then, is *privacy*? Attempts to provide the term with a single, common definition have been frustrated by its nebulous, multifaceted nature, which has led to a disjointed body of research with seemingly contradictory findings [8, 62, 98, 104]. The inability to define a singular, universal concept of privacy has implications for attempts to devise effective mechanisms for its management. The nature of these definitional differences are explored in further detail in Section 3.1 but a brief, encompassing definition is that privacy, at any level, includes an ability to exclude others from participation and observation in activity and knowledge.

Nevertheless, the ubiquitous adoption of information systems that interact with or utilise personal data continues to increase, as does the appetite for increasingly fine-grained information. While there is an acknowledged requirement to 'protect' an individual's informational privacy within such systems, it is unfortunately the case that not only do salient privacy concerns remain, such as the encroachment of data collection and analysis into every facet of individual activity, social interaction quantification and decision prediction about the individual [21, 111], but there are also open questions about how to design and operate systems to manage this concept.

2.1 Engineering privacy

Responses to the challenge outlined above have coalesced around what is broadly termed *privacy engineering*, a developing, specialist field within systems engineering, focused on providing development and management methodologies for systems that have data privacy as a fundamental requirement.

Broadly, the literature on privacy engineering orbits around the fulfilment of four goals within an information system [11, 49, 79]. These goals are as follows.

(1) *Private communications*, which includes anonymisation of communications [41], as well as methods such as homomorphic encryption [112].
(2) *Private disclosures*, which includes statistical disclosure control through methods such as k-anonymity [91, 103], as well as other approaches, such as differential privacy [40].
(3) *Identity management*, allowing the user a measure of control over how personal data is being used within the system (e.g. [5, 10, 67]).
(4) *Information security*, protecting the data from unauthorised access, e.g. [92].

At the most foundational level, these system goals may be realised through the use of privacy enhancing technologies (PETs), which focus on providing *data privacy*. That is, these technologies aim to de-identify the individual to whom the data relates — amounting to anonymisation rather than what might reasonably be defined as 'privacy'. Therefore, the development and use of PETs is not the extent of privacy engineering, as the latter considers wider, more systems-level concerns, as the larger concept of privacy encompasses more complex data management processes [46]. As such, privacy engineering considers the data life cycle and suggests models that follow the same broad guidelines, most of which span from collection, to processing, then preservation, and ultimately re-use [15]. However, these is recognition that a category for destruction and disposal of data is also recognised in the light of costs associated with maintaining such data, as well as security and data protection requirements [128].

Pursuant to this, the use of PETs is often supplemented and enhanced through the provision of organisational controls within a system (e.g. a declared privacy policy as well as staff training). The goal is that these should address different aspects of data use and ultimately complement one another in preserving privacy properties within the system. However, the selection of such measures and controls are not readily prescribed and will largely rely upon the subjective determination of the system's goals and context in relation to informational privacy concerns. In addition, privacy engineering methodologies often recognise legal obligations, which may themselves have explicitly and implicitly prescribed design considerations.

2.2 Facets of privacy engineering

Aside from the technical management of data within a system, there exist other influences on privacy engineering that continue to shape the field's development. We group these influences into three distinct categories — (i) user sentiment; (ii) privacy design principles; and (iii) legal obligations — with no single category any more influential on privacy engineering than the others.

2.2.1 User sentiment and experience. The end-user plays a prominent role within privacy engineering. The user is a dynamic element within an information system: users operate the system in order to

[1] The other two are integrity and availability.

Fair information practice principles
Transparency
Individual Participation
Purpose Specification
Data Minimisation
Use Limitation
Data Quality and Integrity
Security
Accountability and Auditing

Table 1: The principles of FIPPs.

Foundational principles of Privacy-by-Design
Proactive, not reactive; Preventative not remedial
Privacy as the default setting
Privacy embedded into design
Full functionality
End-to-end security
Visibility and transparency
Respect for user privacy

Table 2: The seven principles of Privacy-by-Design.

achieve some goal [94, 116]. This operation is variable, as design of a system may require the participation and operation of some users in order to achieve the end goals of other users [123]. The operation of the system to achieve its goals has become increasingly dependent on the provision of *personal data*, which is defined as any information that may be related to an identifiable individual. The collection and storage of personal data is required for a wide variety of reasons, such as user experience customisation, as well as more innocuous information system performance optimisations. However, there are ever-present social expectations and norms of privacy, even within contexts that might apparently have no explicit privacy concerns [76], e.g. observing others' social interactions in public fora. This property can be obscured with novel information systems uses, particularly those that mimic off-line social interaction (e.g. social media applications). The subtleties of privacy expectations and norms are not necessarily primary considerations in system design.

There exists a large, multidisciplinary corpus of work related to privacy attitudes and concerns. A conspicuous feature of this literature is the "privacy paradox" [19], wherein an individual may purport to care about privacy but yet engage in actions that seem to contradict this sentiment, as demonstrated in the literature [78, 127]. Acquisti and Grossklags [7] argue that this is due to a "bounded rationality" that can obscure the perception of harm and reward. Furthermore, it has been demonstrated that individuals are content to trade privacy for convenience (see generally [8] for an in-depth summary). These trade-offs, as well as the variance of the individual's perception of harm from privacy disclosure, make user sentiment a difficult requirement to adequately represent in an information system. However, some authors have argued that there may not be paradox at all, particularly when separating privacy attitudes from privacy violation concerns [38].

While the user remains the cornerstone upon which the information processing edifice is constructed, there are few validated methods by which to appropriately measure and track such sentiment. In short, privacy attitudes are difficult to identify as a social concept, let alone operationalise, track and measure for the purposes of information systems engineering.

2.2.2 Principles of privacy use and design in information systems. Information systems are designed to capture and process data. These systems are also designed to be networked together in order to share the data and outputs of data processing. These goals are inherently in opposition to the goals of privacy, which we may characterise as an exercise in autonomy, including an ability to exclude all others. This tension pits the innate human desire for context modulation with the efficacies of information processing and management.

There have been attempts to express generalisable system design and use principles which centre on the protection of informational privacy of the data subjects. These principles have been influential not only with regards to methodologies, but also with respect to the law. In this study, we highlight two sets of principles: the fair information practice principles (FIPPs) and Privacy-by-Design (PbD).

In 1973, the United States government commissioned a report on the impact of automated information processing systems and in its conclusions proposed FIPPs for appropriate use, congruent with civil liberties and public sentiment [109]. Table 1 lists these principles, which have since influenced the OECD guidelines, as well as European data protection principles.

While acknowledged as accepted privacy ideals, FIPPs have since been criticised for not being conducive to the practice of systems engineering and design [24]. To this end, PbD purports to offer a more design-focused collection of privacy-protection principles to which system designers can refer during the development process. PbD proposes that privacy should be a fundamental aspect of any system design [29] and provides seven 'foundational principles' to achieve this aim (see Table 2), which should be articulated in a system's design. The central premise of PbD is that personal data use should be minimised at each stage of the data life cycle, and, where it must be used, appropriate safeguards should be put in place [93].

PbD is characterised as a means by which to more effectively design systems that respect and appropriately manage personal data. However, as detailed in Section 3.4, PbD has been criticised as being difficult to implement in practice. Thus, while PbD has influenced privacy engineering, some might argue that it fails to offer much more than FIPPs or the OECD guidelines in relation to prescriptive system design.

2.2.3 Legal obligations. The law has been looked to, with a view to drive privacy engineering, not least because a breach of the law represents a significant risk. However, informational privacy is a contentious subject in jurisprudence, with approaches and perspectives very much informed by the different jurisdictions.

Privacy, as a discrete concept and right, has a relatively modern history in law, with scholastic focus beginning in earnest only as recently as the end of the 19th Century [118], prompted by the use of

the then novel technology of photography. Similarly, informational privacy within computer systems has been a specific concern within the law since the 1970s [28], coinciding with the growing use of such systems. It had been recognised that there was a divergence between the wider concepts of privacy and specific protections needed for informational privacy within a system. Thus began the start of national and international legal developments, with notable international landmarks including the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (1980) and the Council of Europe's Treaty 108 on Data Protection (1981).

Within information systems, the law itself is not focused on privacy as such, but, rather, on the appropriate use of personal information — a concept referred to as *data protection*.[2] Data protection focuses on the appropriate uses of data, fairness in collection, control over disclosure, and obligations for security during storage. An exercise of data protection rights, in contrast to privacy, is not necessarily focused on the impediment of the use of personal data [48]. As such, the concept of data protection has distinguished itself and has thus been seen as a pragmatic approach to avoid the theoretical difficulties of asserting rights within the umbrella of privacy [114] and, perhaps, potentially offers system designers more stable footing while traversing the capricious sump of seemingly contradictory concerns and values of individual privacy. Within this legal concept, the European Union promulgated the Personal Data Protection Directive in 1995 [2]. In the United Kingdom, this directive was implemented in the form of the Data Protection Act of 1998. Additionally, the European Union's General Data Protection Regulation (GDPR) came into force in May 2018.[3] While the United States does not have a single, overarching informational privacy law, personal data is managed and protected via a number of federal, sectoral laws, such as the Health Information Portability and Accountability Act 1996 [4], as well a number of state laws.

It is important to note that, while there is considerable legal debate regarding the transatlantic conceptions of privacy and approaches to its regulation [121], there are also growing concurrent efforts in other parts of the world, such as Asia-Pacific Economic Cooperation Privacy (APEC) Framework [13] and the African Union Convention on Cyber Security and Personal Data [9], which will undoubtedly shape the compliance of information systems in those areas, raising challenges for cross-border information flows.

2.3 Management of risks and harms

Privacy engineering aims to remove the risk of violating properties of privacy within a system. Any system is susceptible to undesired behaviour, effects or failures. The probability of these occurrences is articulated in the form of risk, which is a measure of potential loss due to undesired system behaviour. The ability to identify risk allows for the generation of associated metrics to track the potential of its occurrence, which itself informs the selection of a mitigation strategy, ultimately informing the overall design of the system.

[2] 'Data protection' is mostly used in the European Union legislative context. In US law, it is referred to as 'informational privacy'. While these concepts are not entirely congruous, for the purposes of this paper, we will treat these as equivalent.

[3] The Information Commissioner's Office has stated that the UK will adhere to the GDPR, despite the on-going negotiations to terminate membership to the European Union [55]. Additionally, national legislation has been promulgated largely incorporates the GDPR [1].

What, therefore, is the risk to data privacy within a system? The loss of privacy properties of data can result in identification, but these losses have uncertain effects outside of the system itself — resulting in what may be described as 'privacy harms', a concept that is underdeveloped for the purposes of systems engineering. Social and psychological research remains uncertain with respect to how best to measure such effects. Westin [120] posited that being observed by others was a cause of psychological distress. Margulis [69, 70] argues that social and psychological costs to privacy loss include individual stress and wider social stigmatisation. Lahlou [60] argues that privacy functions as a "face-keeping" exercise, such that individuals are able to participate within social settings without stigmatisation. Therefore, a loss of privacy diminishes the number of 'faces' that an individual may put forth.

Given the law's purview to provide redress for harms suffered, legal scholarship acts as an intuitive source of identification of such harms, which may inform risk assessment. However, even within this field there is uncertainty and debate. Descheemaeker [36], in commenting on the British legal environment, identifies four possible detriments that may be suffered by privacy violation: (i) pecuniary loss; (ii) mental distress; (iii) loss of dignity; and (iv) the tautologous loss of privacy as a value itself. Taking a different approach, Calo [26], in examining the US context, categorises privacy harms as being either *objective*, which includes information about the individual being used to his or her disadvantage, or *subjective*, which considers the anxiety and stress that comes from the "perception of unwanted observation". What becomes immediately apparent is that the harms are difficult to describe, let alone capture and model within a system, as the effects are very often external to the system itself. Furthermore, even if such events could be appropriately categorised, there still remains an open question about how to score or quantify these values. The provision of such scores or values would enable those involved in system design to mitigate appropriately. However, privacy risk analysis is a nascent area of research, fractured by approaches from different disciplines.

Despite these intrinsic difficulties, risk analysis methods have been proposed specifically for privacy (e.g. [35, 52, 79]), although it is unclear how ubiquitous these are in practice. One assessment tool that has gained a degree of currency is the Privacy Impact Assessment (PIA), which aims to identify privacy risks and prescribe mitigations based on an assessment of the proposed system process. This is often conducted through questionnaires and surveys of the system designers, and may include some consultation with the end-users.

PIAs have an acknowledged limitation, in that these are qualitative, subjective assessments, which perhaps serve more of a compliance demonstration function, being viewed as little more than "ritualised hurdles" [71]. This is in no small part due to the recognised need for metrics in order to provide more effective use of PIAs [125]. Wadhwa and Rodrigues [115] identify that the lack of follow-up after the assessment is a shortcoming of PIAs. The authors propose a method for evaluating PIAs, looking at available PIA frameworks, but, again, this remains organisation- and policy-focused, rather than holistic.

2.4 Guidance on privacy engineering

Attempts have been made in order to bring these facets into a coherent methodological framework under the umbrella of 'privacy engineering' in the form of guidance from relevant authorities. We briefly describe such guidance from the European Union Agency for Network and Information Security (ENISA) and the National Institute of Standards and Technology (NIST). These guidelines attempt to provide a means of translation from the articulation of principles and legal obligations to practical systems design, albeit via different approaches.

2.4.1 ENISA. The ENISA guidance [34] is strongly motivated by the EU's data protection legislation, particularly the aforementioned GDPR. The GDPR makes explicit reference to the term 'data protection by design' (DPbD), which can be seen as an adaptation of PbD [29]: the difference between these terms is unclear, and it may be naïve to use them interchangeably. The guidance itself provides a broad overview of available privacy-enhancing techniques. The document identifies that "the challenge for designers is to find the appropriate PET and protocols and combine them to meet the requirements of the system" [34, p.13]. In cases of requirements conflict, the document arguably minimises the difficulty, describing this, with no small measure of understatement, as "a hurdle to be overcome" [34, p.13]. ENISA does suggest that any adopted methodology should be evaluated on the basis of how well the approach meets the needs of: (i) trust assumptions; (ii) user involvement; (iii) technical constraints; and (iv) system architecture. Criteria for how these needs might be sufficiently fulfilled or appropriate measures are not addressed by the ENISA guidance; reference is instead made to the law and privacy-first design principles.

2.4.2 NIST. NIST provides a framework for privacy engineering for federal systems, NISTIR 8062 [24], which aimed to devise novel approaches for the assessment of privacy risks and mitigation in information systems. The guidance makes a point to depart from FIPPs as guiding principles, as these are more akin to value statements, rather than prescriptive implementations of privacy. The guidance argues that the principles make it difficult for system designers to evaluate or compare different actions, as there is no frame of reference for such evaluation. One consequence, highlighted by the authors, is that privacy risk assessment becomes more about compliance, rather than achieving positive, measurable outcomes for informational privacy protection. This does not discount the use of law or similar principles, but is demonstrably aiming for something beyond the legal threshold for compliance.

The NIST guidance places more emphasis on risk assessment than its ENISA counterpart. To this end, the document suggests that privacy violations are not necessarily the result of adversarial activity, but something more akin to a lack of due diligence or forethought. Therefore, instead of appropriating the language and concepts of threat modelling from information security, the suggestion is that "problematic data actions" act as the measure for risk. A risk is thus composed of: (i) a data action, e.g disclosure, storage, or collection; (ii) the personal data; and (iii) the context. This is a novel conception, but is limited in that the means of how context is not prescribed.

2.4.3 The insufficient depth of the guidance. The guidance provides indications as to what privacy engineering might entail, without being especially prescriptive or definitive as to what works best. This is arguably a result of the fledgling nature of privacy engineering as a discipline. The NIST guidance is an on-going effort of consultation and refinement, with the ENISA guidance acting as a beacon that privacy engineering, especially within the regulatory context of the GDPR, is possible, albeit with strongly acknowledged lacunae in the identified techniques. These documents are useful resources in order to get an initial start within the fields of privacy engineering, as they are complete with references to appropriate literature and attempt to link these disparate findings into a single cohesive methodology — even if the cohesive nature is sparse.

2.5 Privacy engineering's proposition

The proceeding subsections provide a general overview of privacy engineering and gives sufficient foundation to outline what we describe as its proposition:

(1) Privacy is a social phenomenon, with specific distinguishing facets, that can not only be defined but can be expressed within a deterministic information system.
(2) Privacy requires unique protection measures, distinct from informational security, which can be adequately provided through the application of technical and organisational measures.
(3) The protection of privacy is commensurate with the other value propositions of the information system.

The above proposition of privacy engineering is supported by a number of assumptions. The effect of relying on these assumptions, we argue, is that the proposition of privacy engineering becomes difficult to achieve. Where privacy engineering does not sufficiently reach its stated goals, the value of utilising privacy engineering methodologies becomes less apparent, if at all. Furthermore, the proposition is encouraged by a system of governance and regulation that suggests onerous requirements — with a lack of means to evaluate and enforce such requirements.

3 DE-CONSTRUCTING PRIVACY ENGINEERING

The survey of constituent facets of and influences on privacy engineering highlights that there are gaps within understanding privacy as a concept, which has an effect on proposed methodologies for its management. This is evident even within examined guidance, suggesting that the subjective nature of privacy necessitates a subjective approach, thereby side-stepping the provision of more prescriptive solutions.

Given these gaps, what is therefore needed to advance the field from its current state? We suggest that privacy engineering must (i) examine its underlying assumptions; (ii) identify its obstacles; and (iii) firmly demarcate those aspects of privacy that may be supported by privacy engineering, as well as those that fall outside of its scope.

Value-based definitions	Cognate-based definitions
Privacy as a right	Privacy as a state
Privacy as a commodity	Privacy as control

Table 3: A classification for definitions of privacy, as provided in [98].

3.1 Privacy's multi-dimensionality and the difficulty of definition

While the literature cannot agree on a common definition of privacy, there is consensus that the concept is multifaceted, defying delineation, which is a source of difficulty when dealing with the notion in the context of information systems. The definitions within the research literature are sometimes at odds with one another, often appearing to be contradictory, or even orthogonal. We argue that many of these various definitions are not in competition with one another nor necessarily contradictory, but, rather, are related, intertwined, yet distinct concepts that have been crudely lumped into a single word: 'privacy'.

While there exists an acknowledgement within research literature that privacy challenges ought to be addressed by holistic approaches, this multiplicity of dimensions to the concept of privacy frustrate efforts to provide discrete boundaries and thus rationalise about challenges and issues regarding privacy. The fundamental challenge posed by the ambiguity of privacy has caused difficulties in other fields, especially those in need of solutions based around conceptions of privacy.[4] The underlying questions for any method that attempts to address privacy are:

(1) What is privacy?
(2) Knowing what it is, why is it required?
(3) How can it be best represented in a model?

The answers to these questions depend, in part, on the disciplinary perspective used. Broadly, the answers will divide depending on the starting point for analysis. For this, we can turn to existing attempts to categorise the existing literature in order to tease out the different concepts and ideas. We suggest that there is scope in research to map privacy risk and impact analyses onto these categories. The focus, we believe, should be on understanding how these concepts relate to and interact with one another.

A comprehensive overview of the different research directions is provided by Smith *et al.* [98], who examined the corpus of information privacy research from the 1970s to the early 2010s. The authors divided the works into two broad categories: *value-based* and *cognate-based* definitions, each with two further sub-categories, as shown in Table 3. The authors make particular note that the normative perspective of the value-based definitions are incomparable with the descriptive focus of the cognate-based ones, highlighting that researchers often do not distinguish between the two. Moreover, researchers do not identify their own perspective when describing their research. Furthermore, the authors argue that this disjointed nature of privacy research over four decades has generated an inability to identify actionable measures for its protection.

[4]For a summary of the definitional challenges of privacy within surveillance studies, see Bennett [22] and a response by Regan [88].

One of the foundational ingresses into dissecting privacy is a premise that equates privacy with an ability to exclude others from observation or participation within one's activity [47, 118]. From this perspective, Solove [99] devised a privacy taxonomy, splitting the concerns into four distinct categories, based on a type of threat or harm, each of which is itself comprised of further sub-categories: (i) information collection; (ii) information processing; (iii) dissemination of information; and (iv) invasion. Solove's taxonomy may be seen as limited, in that it is highly legalistic, something that is acknowledged by Solove. Nevertheless, it remains useful as it provides a basis for reasoning about the various means of how data can be used with the broad headings. Calo [26] is critical of this taxonomy, arguing that, under the Solove taxonomy, it is difficult to challenge sources or even add new sources for definitions of privacy. Citron and Henry [32] similarly are sceptical of the ability of the taxonomy to remain dynamic and not succumb to "ossification". Bartow [20] criticises the taxonomy as not being sufficiently thorough to address the actual harms from a privacy breach, highlighting that Solove's taxonomy provides little material or physical harm, other than unease of being watched.

There are, of course, other approaches by which privacy can be categorised, including by activity. To this end, Finn *et al.* [43] provide seven types of privacy: (i) privacy of the person; (ii) privacy of behaviour and action; (iii) privacy of communication; (iv) privacy of data and image; (v) privacy of thoughts and feeling; (vi) privacy of location and space; and (vii) privacy of association. Each of these begins to delineate the discrete concepts within privacy. The defined categories are sufficient, although there are clear overlaps and the interactions of those overlaps are not readily explained.

The determination of privacy is not static, as a desire or need for privacy may decrease with an increase in the want for some outcome of social interaction. A desire for more privacy may be motivated by a desire to conceal some fact that may diminish one's reputation [85], with less privacy required in a mutually beneficial social interaction. In the latter case, the less stringent personal thresholds for privacy may be determined by social consensus [87, 102], although, given the novelty of information system services, this may be difficult to ascertain with consistency.

Immediately at the definitional level, privacy engineering encounters a fundamental challenge in conceptually articulating and representing the concept of privacy. This goes beyond semantics, but has an impact on the operationalisation of privacy within information systems. How can it be validated that the privacy conception used within the system reflects the conception of the users? How can the operationalisation of concepts of privacy remain dynamic? These unanswered questions are compounded by the assumptions and obstacles detailed in the next two subsections.

3.2 Obstacles to privacy engineering

A number of obstacles continue to frustrate privacy engineering efforts. These obstacles exist concurrently with privacy engineering goals and are often in direct opposition to such goals.

3.2.1 Increasing demand for personal data. The tension between the demand for personal data and informational privacy is intuitively recognised within the literature across multiple disciplines. These demands include drivers such as personalisation within the

commercial sphere, reduction of data processing within the public sphere, and the promises of improved efficiency, convenience and progress through artificial intelligence and machine learning. Privacy is then pinched between ambitions for increased automation and predictability, which relies on learning from vast amounts of personal data. This appetite for personal data is not without its benefits. The outcome from increased personal data collection is increased convenience and personalisation, which, despite the protestations from a principles-based perspective, are typically welcomed by consumers. The slow erosion of any previous conception of privacy is thus offset by a new service offering.

3.2.2 Re-identification. One of privacy engineering's goals is to be able to ensure private disclosures. Re-identification, also sometimes referred to as de-anonymisation, of disclosed data occurs when an identifiable individual (or group of individuals) can be determined from an anonymised data set. The foundational cause for re-identification is typically (a) there is an additional data set that allows for 'jigsaw' reidentification and / or (b) the means of anonymisation is trivially reversible. Re-identification has been the subject of public controversy, including the 2006 America Online release of search queries [17], the Netflix competition [72] and the New York City taxi data release [82]. Within the UK context, open data has been found to have been prone to re-identification, even with simple techniques [81, 97, 107].

3.2.3 Tracking of activity. Tracking of activity is a firmly established feature of information systems [16], with added dimensions of non-online activity with regards to wearable devices and geolocation [18]. Many services are built on the value of providing summaries of such activities within the system, allowing for these services to personalise and customise experiences for users, as well as to allow the user to view their own data. However, there is a cost to the increased scrutiny and observation, especially where it is not transparent. In the now infamous Facebook 'emotional contagion' study, Kramer *et al.* [59] aimed to measure whether emotions can be spread through the social network. Putting aside discussions regarding the conclusion and significance of findings, the study highlights an important aspect of modern information system participation: personal data is viewed as a commodity to be refined. This requires that every measurable facet of participants' activity is tracked and analysed.

3.2.4 Opaque decision-making and informational asymmetry. The overall effect of tracking and information demands is that participation within information systems increasingly equates to a diminishing degree of control over informational privacy. This creates an informational asymmetry, which can lead to discriminatory practice, based upon categorisations that are not transparent to the individual, with regards to decisions being made [45, 68]. Nehf [74] provides another dimension to the concept of harm, contending that identification in a data set gives an "incomplete set of facts" and therefore may lead to mis-characterisations of that data subject. These mis-characterisations may have serious consequences within those systems that automate decision-making, whether in part or in whole, about the individual.

3.2.5 Privacy violations do not necessarily result in direct harms. Not all violations of privacy reach a sufficient level to cause harm. Here, an important distinction is required. A *privacy violation* is an objective event, wherein a guarantee of privacy is broken. This may be done by disclosure, inference, or unauthorised access. A *privacy harm* is some damage suffered because of the violation, and can be either objective and subjective (see the discussion in Section 2.3). It is unclear whether a violation of privacy is itself sufficient for a data subject to cease participation in an information system. If a subsequent harm would be necessary to cease participation, there is little evidence to suggest what such a threshold would be.

3.3 Identifying and challenging assumptions

Compounding the definitional challenge and the persistent obstacles to privacy in information systems, there are a number of implicit assumptions around privacy engineering which serve to limit the field's development. We now identify and address three such assumptions.

3.3.1 Informational privacy is reducible to a technological problem. While there is an acceptance of privacy being a challenge that requires a holistic and comprehensive approach, there is at least some implicit assumption in methodologies that the appropriate technology will be able to solve a variety of privacy concerns. The definitional issues of privacy act as a double-edged sword: at once, this is a constraint as the proposed solution is unable to generalise. Thus, if a proposed definition of privacy is accepted, then, for every application of that definition, a solution is available. However, it is common that the evaluation of the context is left virtually unaddressed, thereby leaving only an arbitrary definition of privacy, which happens to be solved by the solution proposed. What proposed technologies there are have few published studies pertaining to real-world systems.

3.3.2 The legal principles of informational privacy can be embedded into the system's architecture. Another assumption is that the legal obligations prescribed by legislation and policy can be built into the system itself, making the principles self-executing [58]. This assumption can trace its historical roots to the notion that legal regulations may be viewed as akin to software regulating the operation of society and thus those legal principles ought to be enshrined at the most granular, technical level, e.g. [65, 66]. Thus, information systems should be designed to self-execute legal principles. This is both defeatist, in that there is a diminished faith in regulatory mechanisms to manage those technologies and overly optimistic, in its implicit assumption that software development can incorporate the law such that the principles of law become self-executing in the information systems themselves. Legal scholarship may be beginning to reflect the reality that this may be an impossible feat, at least not without serious refashioning of basic legal theory such that these legal precepts may find expression in an information system [25, 124].

3.3.3 Greater penalties will drive adoption of more effective privacy engineering methodologies. While earlier work suggested that privacy protection will increase value to an entity [57], it is arguable that the benefits remain unclear, especially in relation to costs. In order to motivate the development and adoption of privacy-focused systems design, authorities, most notably the European Union, assume that stringent penalties drive better practice and

the development of cost-effective, easy to adopt technologies — but there is little evidence to support this. In addition, this seems to be the only driver for pursuing these improved methods as the other market risks pale in comparison. Despite the warnings of industry surveys that data mismanagement has dire financial consequences, such as those from the Ponemon Institute [84], this has not been observed in empirical studies. Risk of financial damage from privacy violations are limited [6, 89] and, longitudinally, the evidence suggests that there is no long-term impact in terms of user trust and use [27, 61]. Thus, while consumers have privacy concerns, there are mitigating factors when considering trust in a system. Furthermore, there are difficulties in enforcement procedures for data protection authorities, as they are limited by available staff, effective assessment procedures, and intelligence gathering mechanisms [30, 31]. This limits the scope of what a regulatory authority is able to do. Being "selective to be effective" [54] in information privacy regulation risks creating a skewed environment for penalties.

3.4 The limits of privacy engineering

The persistent challenges, in conjunction with the underlying assumptions, serve to limit the applicability and efficacy of privacy engineering. These limits are not fatal to privacy engineering's overarching goals, but, when taken together, the limits diminish privacy engineering's value proposition. As a result, there is a decrease in motivation to provide solutions to the open challenges that privacy engineering faces — with the most prominent being that the privacy design principles are insufficiently prescriptive.

There have been a number of critiques of the applicability of privacy engineering, especially insofar as the attempt to incorporate 'principles of privacy' (as prescribed by, for example, PbD) into existing systems engineering practice. Using Cavoukian's own material, Gürses *et al.* [49] struggle to provide an appropriate definition and scope to PbD, highlighting that the "vagueness" of principles of PbD is so remote from engineering applications that it truncates the utility of PbD. Similarly, van Rest *et al.* [113] criticise PbD for not acknowledging or building upon existing methods for systems engineering. The authors point out that there are no guidelines as to how to apply PbD to a particular domain, nor identification of what constitutes good PbD practice, which means that comparing PbD to other privacy-enhancing methodologies is difficult.

Thus, while it may be acknowledged that PbD, at least in a broad sense, may provide some direction towards privacy protection, the lack of detail on implementation, including how to apply these principles to legacy systems [110], means that PbD provides a strong explanation as to *what* an information system should do with regards to personal data, but is silent on *how* this may be done in harmony with other system specifications, as well as *when* the protection of privacy is adequately achieved. Within this context, we identify three limitations of privacy engineering that are at the focus of current research.

3.4.1 Technology is unable to capture the full scope of privacy. PETs are technologies that provide a function to enhance some definition of privacy within an information system. For a time, there was some consensus that by developing suitable technology, the risks to privacy might be minimised [25]. However, this supposition has been since challenged, especially in light of the evolving complexity of information systems. Ashley *et al.* [12] argue that technical privacy solutions are stymied in that the available technologies only address a "fraction of the problem", suggesting privacy risk mitigation ought to develop more holistic enterprise-level approaches. While it may be possible to define privacy as a formal, mathematical property [63], this may not adequately reflect privacy's role as a social and psychological construct. Thus, relying solely on formal guarantees of privacy provided by a system leaves more ambiguous, yet pertinent, aspects of privacy unaddressed.

There have been attempts to address this need for holistic scope. Spiekermann and Cranor [101], for example, differentiate between architectural privacy and organisational privacy, arguing that there needs to be both for an effective approach. While acknowledgement of a holistic approach is congruent with privacy-values-as-context-dependent assertions, the approaches have an implicit need for context-dependent risk assessments. Those risk assessments then need appropriate mitigations that do take into account privacy–utility trade-offs. This is missing from methodologies.

PETs are highly specialised, deterministic approaches to very specific challenges. Merely adopting PETs into system design only addresses privacy risks for which those technologies were adopted. This may impact on other system requirements. No exclusively technological solutions exist for privacy risks and the technology that does should be selected with consideration to other system requirements. Additionally, users must know how to use PETs in such a manner to maximise their efficacy. This relates to user experience and expertise in managing these technologies. PETs may interfere with the 'normal' operation of a system, prompting users to elect not to use them. This has been identified within usable cryptography [90, 96, 122], wherein the user everyman finds the experience and use of the technology cumbersome, time-consuming, and, ultimately, bothersome. Therefore, while a technique may be devised to manage some aspect of privacy within the system, its utility may be undercut by the intended audience's lack of use — especially when such utility cannot be properly motivated by the immediacy of perceived risk.

3.4.2 The difficulty in defining privacy risks. There is the possibility that personal data within an information system are at risk of being violated in some manner. However, there are few systematic approaches to assessment of such risk. One common approach within privacy engineering is to adopt information security risk assessment methodologies, e.g. the ISO 2700 family of standards [3]. However, the terminology of 'threats' and 'vulnerabilities' is not neatly applicable to privacy issues [24]. While it is evident that privacy and security have an identifiable overlap, violations of privacy extend beyond this to include non-security related failures in processes.

There is a lack of holistic methods to identify privacy risks across different contexts. Shapiro [95] finds that while qualitative, normative assessment of privacy within a system (e.g. PIAs) are ubiquitous and, perhaps, even well-defined, the development of objective, technical-driven analysis of privacy risks within such systems is "lagging". In 2013, the UK's data protection authority, the Information Commissioner's Office, commissioned a study on privacy risk management methods [106]. Of the 15 methods analysed, only three

specifically addressed privacy, with all others being general information security practices. The identified methods for privacy risk management were all qualitative, though the report acknowledged the weakness in this, arguing for more objective, quantitative methods. There exists research to suggest that qualitative assessments are time-consuming, expensive, and often suffer from ineffective problem scoping leading to poor resource allocation [33]. This is not to argue against qualitative methods, but, rather, to argue for complementary measures.

Any model for privacy risk and associated metrics, therefore, ought to incorporate all three of these principles. Individual privacy harms are difficult to address; further, there is an associated, organisational risk of being non-compliant with legal obligations in data protection laws. While these are linked, they are not necessarily congruent. For example, a data breach may occur but escape legal liability; it does not necessarily follow that a breach will also receives a penalty.

Furthermore, there is a distinction between privacy concerns of individuals and those of groups [44]. Data protection is heavily focused on the individual and only gives scant acknowledgement of concerns of larger groups. This has implications for automated decision-making and profiling; while the individual is affected, it is based on the categorisation and membership to a particular group. This has social impacts beyond individual welfare [105] and has only had limited attention in both data protection law scholarship and privacy-enhancing technology research.

The NIST framework relies on information security guidance as applicable to privacy engineering, with an acknowledgement of those areas where it does not work, such as threat models. Instead, a novel approach to privacy risk is proposed: the risk is determined by a tri-partite combination consisting of (i) a data action that (ii) contains some manifestation of personal data in (iii) a particular context. However, the guidance is silent on how to actually use this.

It is clear that risk models for privacy should be more nuanced and expansive than 'threat modelling', as privacy risk is not only objective to the system, but is also subjective to the individual. By reducing the dimensionality of privacy to mere threats, entire definitions of privacy may be left unaddressed by the system: while the system may protect privacy in one conception, it may leave it vulnerable in another. While privacy is viewed as multidimensional, the conception of privacy harm is almost always addressed in a singular dimension, e.g. reputation, monetary or opportunity [39, 98]. Privacy risk assessment needs to have different identification processes and measures for different definitions of privacy, with an understanding of how the risk affects multiple dimensions.

3.4.3 The need for validated metrics. Privacy engineering is limited not only by the narrowly defined methodologies but also by the lack of validity of proposed solutions. Enhanced validity can provide for better metrics. Validated metrics can enhance the value proposition of privacy engineering as metrics provide a way to track whether the system is improving (and to what extent) or not. Such metrics might provide some much needed measures as to the criteria for successful privacy engineering.

Any type of engineering requires metrics in order to evaluate the effectiveness of the measures implemented in a system; this is no less true for information systems engineering. Metrics provide the ability to evaluate whether the system is meeting its objectives and, if not, how much improvement is needed and where. In order to do this, there needs to be a frame of reference and thresholds for defined values. This helps to prioritise risks and inform investment. However, in context-dependent areas, such as security and privacy, there is a challenge to provide such thresholds, due to the subjective nature of the concepts. This challenge is further compounded by the lack of availability of good data sources for metrics [51].

In their literature review of privacy engineering design patterns, Lenhard *et al.* [64] found few studies utilised empirical experimentation of the proposed methods, limiting their validity. Furthermore, the authors note that, while privacy engineering is motivated by a holistic approach to implementation, the literature does not attempt to address this in a meaningful manner. Even those risk analysis and assessment approaches that are sufficiently quantitative are subject to qualitative decision-making processes [50].

3.4.4 Imbalances in the user–information system relationship. The user is at a distinct informational disadvantage in relation to the information system; users must rely on the information about the processing and storage of their data, often provided by a privacy policy. The law requires that the information system provide some degree of transparency about these processes, additionally empowering the user with certain rights over their own data. However, despite these provisions, users are reluctant to trust the system or feel assurance in the security of their data [56, 108, 126].

3.5 The effect of limitations

The overall effect of the limitations is to reduce the overall value proposed by privacy engineering:

- If privacy cannot be adequately defined, it cannot be adequately represented in an information system.
- If privacy values cannot be effectively measured, it is difficult to improve upon its current state within the system.
- If privacy as a feature cannot be measured within a system, it cannot be correlated with the other value propositions of the system, thereby providing motivation to further improve privacy.

We argue this reduces privacy engineering to a near-exclusive compliance exercise, which fails to sufficiently distinguish privacy engineering as a distinct field from information security or wider systems engineering. One may conclude that information systems owners, designers and operators will simply only do enough to avoid a penalty as there is no clearer motivation to invest beyond this — as it is unclear what the tangible goal is.

It may very well be that the only motivation is to avoid regulatory action, the possibility and cost of which may be acceptable given the limited resources available to authorities [30]. As for economic losses from privacy violations, there is little to suggest the majority of entities suffering such events experience long-term effects [6, 61, 89]. Given that data breaches represent the largest of privacy violation actions [31, 89], there is little to distinguish privacy engineering from informational security — and therefore there is limited motivation to engage with conceptions of privacy beyond those found in common techniques for access control and confidentiality.

4 TOWARDS AN ALTERNATIVE APPROACH

The arguments outlined in the above sections are not presented in order to discredit or diminish the need for privacy engineering. Rather, the intention has been to strip away the assumptions upon which privacy engineering has been developed in order to help inform research directions. We would argue that the field should attempt to influence system design in a constructive manner that goes beyond a perfunctory salutation to legislation or a token acknowledgement of user concern.

How then to move forward? The starting point has to be an understanding of the function of privacy, especially within the system in which personal data is utilised. The design of the system and the manner by which it is engineered ought to provide added value to the stakeholders, users and operators alike. The exact form of that value is dependent on the purpose and design of the system. We argue that this value can more effectively motivate adoption of privacy engineering in contrast to regulatory penalty. Furthermore, the drive for using data effectively (i.e. maximising utility and informational privacy protection) is an attractive value proposition and, therefore, may motivate more effective risk assessments.

There is still the obstacle of *how* to understand privacy. Social theories are much less easily tested than they are generated; diverse data is difficult to collect and concepts are hard to operationalise in an appropriate manner. This has led to a multiplicity of theories of social phenomena, but no means of discernment as to which provide solutions. To this end, Watts [119] proposes something akin to a division of basic and applied sciences, in that there is space for social sciences to adopt a more solution-driven approach. More specifically, there is a need to complement wider social theory with specific solution-driven research that addresses systems design problems. Using this notion, we may, perhaps, coax ourselves away from the attractive proposition of devising a unified grand theory of privacy, and focus efforts on understanding what definitions of privacy are required by users in specific contexts. The diversification of effort may result in a more complete catalogue of privacy contours. Through such a mapping, we may distance information system design and operation away from current privacy engineering assumptions, and replace them with more actionable means. This, therefore, poses the following questions, which might characterise future privacy engineering research.

(1) How can context be identified? The question should revolve around which items of data users are required to provide, and which they prioritise in this setting.
(2) Given an identified context, what are the goals of the interactions? The value proposition to the user for participation should be made clear
(3) For the context, what harms might occur should the system fail to maintain its guarantees?
(4) For these harms, what mitigations are necessary to minimise impact? What are the limitations of these mitigations? What should be done with residual risk?
(5) Do these mitigations encourage participation within the system?

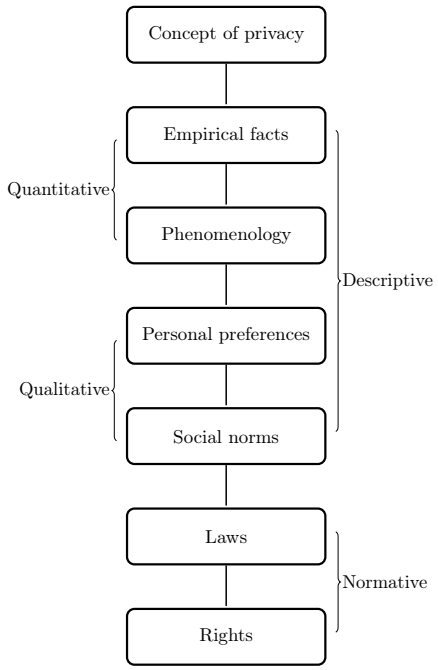

Figure 1: Annotation of types of values and definitions within the proposed division of privacy according to [80].

4.1 Operationalising the multiple dimensions of privacy in systems engineering

Using Watt's suggestion [119] that a field of solution-based approaches for social phenomena might be further developed, we consider what this might entail for privacy engineering research. To this end, we propose an outline that may shape the methodological development:

(1) Identify the social interactions with which the information system is dealing.
(2) Establish what minimum threshold is necessary for the appropriate amount of participation of users.
(3) Monitor for changes of perception and needs of participants, especially with needs for the system itself.

As a starting point, the question of *how* to better operationalise privacy for systems engineering needs to be addressed. While this remains an open question, we suggest that any operationalising ought to incorporate a multidimensional model of privacy. How might this look? We highlight an approach taken by O'Hara [80], who suggests a seven-level division of privacy that goes from simple descriptive abstraction to normative, prescriptive rights. We do not suggest that this scheme is the sole means of moving forward, but that it is illustrative of the type of thinking needed to operationalise the multiplicity of privacy definitions and uses.

O'Hara's seven levels provide a developing scale, which may provide the possibility that these levels are interlinked; each level is derived from the previous one and informs the subsequent level. This provides a bridge between the abstractions to measurable

phenomena to rights that may provide a basis for operationalising privacy within a system. Further, we suggest that the correlation between these levels may function as a potential foundation, as this would potentially provide metrics by which to track and measure privacy within the system.

For systems engineering, we may adapt the seven-level model into the following:

(1) Context: These are the social norms and expectations as well as sectoral practice. This definition has been explored by Nissenbaum [76, 77], who describes privacy as *contextual integrity*, in that the individual is not compromised by the disclosure or information being provided.
(2) Compliance: The obligations and rules surrounding the context. This goes beyond the law and can relate to more abstract principles such as FIPPs, PbD, and/or the European data protection principles.
(3) Data Flow: Here, 'data flow' is meant as the processes and paths by which data moves throughout the system. The data flow will be influenced and determined by the system's requirements. It should describe (i) what data is needed, (ii) how it will be processed, (iii) where it is stored, and (iv) the final outcome of these processes.
(4) Architecture: The architecture of a system is the implementation of the data flow. The architecture defines the specific components that will constitute the system, as well as the interactions between those components.
(5) Transactions: The atomic actions that occur on the data itself. It is at this level that formal definitions and guarantees should be focused.

We reason that, because privacy is a multi-dimensional construct, its protection would necessitate a similar approach. We argue that such a proposition would strengthen mitigations, as it would provide multiple levels of intervention. Furthermore, if such linkages were established and validity of these correlations could be demonstrated, this may provide for a more fruitful environment for the development of much sought-after metrics. In order to establish these linkages, however, empirical research is required.

The levels interact with one another, with actions on one level correlating with effects in the other levels, especially those that are adjacent. These correlated effects can then provide the basis for metrics, which will allow system designers and operators to understand the impact of interventions. Each of these levels then requires its own risk assessment and associated metrics. These analyses and metrics should relate to one another in some meaningful manner. This requires a research methodology that seeks to correlate and measure the impact of interventions in a system.

4.2 Improved context-identification improves trust and engagement

Improved privacy-respecting analytics and efficient anonymisation techniques will always be of interest to privacy engineering; these operate at a transactional level and provide one of the foundational bases for data privacy. However, given the focus of context and proposal of methodologies for its identification within this paper, we highlight a special case of risk within information systems that may not receive as much attention: trust between the data subject and the system itself. We argue that this facet is related to the psychological harms that a privacy violation can cause. However, the level of this trust is difficult to measure. The intuitive answer is to rely on external authorities to manage the behaviour of such entities.

Improvement of trust within information systems does require, in part, a more effective regulatory regime. (The intricacies of privacy regulation are outside the scope of this paper; however, concerns about the effectiveness of regulation have been expressed throughout.) In addition to this, we would argue that privacy engineering adopt techniques to improve trust. This can be achieved by improving the clarity of how privacy operates within a system, extending beyond privacy policies to give consideration to the user experience, and providing multi-layered mitigations for the multidimensional nature of privacy harms.

The preceding two sections provide a framework for the identification of risk and harms on a more holistic understanding. A consequent question is how to manage and evaluate those harms and select appropriate mitigations. For this, we can utilise some of the principles described by both NIST and ENISA and discussed in Section 2.4. The different elements can be amalgamated into a privacy triad, which might help inform the type of controls that should be considered. For the purposes of this section, we adopt the concepts from the NIST guidance [24]:

- *Manageability* refers to the degree of intervenability by either the operators or designers of a system or by the user themself.
- *Predictability* embodies a concept of transparency, which allows for the knowable and expected function of an information system.
- *Unlinkability* is a characteristic of dissociability of an identity to a particular instance of data. The ability to unlink information may not always be possible.

The use of this triad may serve to evaluate how intervention at each level operates. Additionally, this triad may be a vehicle by which to traverse the different levels of privacy within an information system. Research within this framework may focus more intently on psychological protections for users, providing mechanisms that more effectively ensure users can retreat from interaction with a system. This is seldom addressed by current methodologies. Nevertheless, there are examples within the literature that illustrate possible means of incorporation into system design. For example, Egelman and Peer [42] propose a "psychographic targetting" approach to privacy and security mitigations, wherein the focus is to individualise mitigations and controls for users. This may alleviate some privacy concerns as well as lead to more effective mitigations. There is scope for research into dynamic controls for users, especially when paired with appropriate privacy semiotics. This has the potential to improve user trust and confidence within information systems, which is characterised by distrustful impersonal relationships between these parties [14].

4.3 Other directions for future development

While the above proposal addresses some of the presumptions and limitations identified in previous sections, there are additional privacy concerns that remain unaddressed.

Linking different levels of privacy together. If the premise of different privacy levels is accepted, then there is a necessity to understand how these different levels interact within an information system. The understanding of these interactions will inform the development and implementation of more holistic privacy control mechanisms. These levels could be integrated into existing risk assessment methodologies (e.g. privacy impact assessments). At present, the notion of differing privacy levels remains speculative and theoretical and thus in need of empirical assessment.

Validation of privacy risk identification methods. While the above may be intuitive to some degree, experimental research is required to test the constructs presented. Empirical research ought to be undertaken to investigate whether the links between the different levels correlate with one another and whether intervention in one has some effect on another. This will provide the necessary feedback to help hone a system design methodology. The metrics devised will shape methodologies more suited to the achievement of the purported value proposition of privacy engineering, especially if the purpose is something beyond a mere compliance exercise.

Asymmetries in vertical privacy relationships. Communications services often provide horizontal privacy controls for users, that is, controls to minimise information dissemination to *other* users. However, there are no controls with regards to the service itself, which we characterise as a vertical privacy relationship. This is a complicated tension to manage, especially where the information system is the exclusive service offering, such as those entities that either enjoy a dominant position in the market or government services. The user very often has no recourse to manage this relationship, and must solely rely on the regulatory framework for informational privacy as well as the remedies available via the law. The potential for abuse and manipulation is high and current means of regulation are insufficient for wide-scale monitoring and intervention.

Temporal management of privacy requirements and harms. Privacy expectations and effects of privacy violations are thought to be temporally influenced. Privacy engineering methodologies most often treat these requirements and effects as static, although there is some recognition that these should be periodically revisited. To date, this area has been insufficiently explored; further research in this area would help inform the context identification focus proposed in this paper.

5 CONCLUSIONS

Informational privacy is not solely a property regarding the content of data but also a characteristic of a contextual environment in which that data exists. If we accept that privacy is multi-faceted, then it follows that addressing the concept within a system requires different approaches for different aspects. As such, technological solutions can only address specific properties of data and guarantee very specific aspects about a system. Many privacy engineering methodologies and guides recognise the importance of a holistic approach, although it remains unclear what constitutes such an approach.

We argue that the many dimensions of privacy are left insufficiently addressed within privacy engineering, most likely due to the field's lack of validated methods by which to appropriately identify and model such dimension. Furthermore, without clearer benefits, stricter penalties from regulators will not necessarily drive innovation in this field. This is because specific goals and thresholds are unavailable for privacy engineering.

There is also an assumption in law that adherence to such principles sufficiently addresses informational privacy concerns of users. This reduces the concept of privacy to a narrowly defined legal abstraction, which may ignore its function in social interactions.

A framework for privacy engineering should therefore accept that, because of privacy's multiple definitions — many of which have a strongly subjective quality — a generalisable solution for every instance of information flow is impossible. Furthermore, such a solution is undesirable on some level as it avoids accepting the possibility that privacy encompasses many different concepts. The user experience, regulatory obligations and system functionality must be linked. If informational privacy protections give rise to better functioning of other requirements, or even have an effect on user engagement, then there will inevitably be greater motivation to engage with these methods.

ACKNOWLEDGEMENTS

The authors would like to thank the reviewers for their constructive and thoughtful comments. AC's research is supported by the Centre for Doctoral Training in Cyber Security, which is funded by EPSRC grant number (EP/P00881X/1), as well as the Oxford Radcliffe Scholarship provided by University College, University of Oxford.

REFERENCES

[1] Data Protection Act 2018. https://www.legislation.gov.uk/ukpga/2018/12/contents/enacted.

[2] 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. http://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX%3A31995L0046.

[3] ISO/IEC 27001:2013 Information technology – Security techniques – Information security management systems – Requirements. https://www.iso.org/isoiec-27001-information-security.html.

[4] Health Insurance Portability and Accountability Act. https://www.gpo.gov/fdsys/pkg/PLAW-104publ191/html/PLAW-104publ191.htm, 1996.

[5] Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L. F., Komanduri, S., Leon, P. G., Sadeh, N., Schaub, F., Sleeper, M., Wang, Y., and Wilson, S. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR) 50*, 3 (2017), 44.

[6] Acquisti, A., Friedman, A., and Telang, R. Is there a cost to privacy breaches? An event study. *Proceedings of the International Conference on Information Systems (ICIS)* (2006), 94.

[7] Acquisti, A., and Grossklags, J. Privacy and rationality in individual decision making. *IEEE Security & Privacy 3*, 1 (2005), 26–33.

[8] Acquisti, A., Taylor, C., and Wagman, L. The economics of privacy. *Journal of Economic Literature 54*, 2 (2016), 442–492.

[9] African Union. African Union Convention on Cyber Security and Personal Data Protection. https://au.int/en/treaties/african-union-convention-cyber-security-and-personal-data-protection, 2014.

[10] Ahn, G. J., Ko, M., and Shehab, M. Privacy-enhanced user-centric identity management. In *IEEE International Conference on Communications* (June 2009), pp. 1–5.

[11] Alshammari, M., and Simpson, A. C. Towards a principled approach for engineering privacy by design. In *Annual Privacy Forum* (2017), E. Schweighofer, H. Leitold, A. Mitrakas, and K. Rannenberg, Eds., vol. 10518 of *Lecture Notes in Computer Science (LNCS)*, Springer, pp. 161–177.

[12] Ashley, P., Powers, C., and Schunter, M. From privacy promises to privacy management: A new approach for enforcing privacy throughout an enterprise. In *Proceedings of the 2002 Workshop on New Security Paradigms* (2002), NSPW '02, ACM, pp. 43–50.

[13] Asia-Pacific Economic Cooperation. APEC Privacy Framework. https://www.apec.org/Publications/2017/08/APEC-Privacy-Framework-(2015), 2015.

[14] Bachmann, R., Gillespie, N., and Priem, R. Repairing trust in organizations and institutions: Toward a conceptual framework. *Organization Studies 36*, 9 (2015), 1123–1142.

[15] Ball, A. Review of data management lifecycle models. http://opus.bath.ac.uk/28587/, 2012.

[16] Banse, C., Herrmann, D., and Federrath, H. Tracking users on the internet with behavioral patterns: Evaluation of its practical feasibility. In *Information Security and Privacy Research* (2012), D. Gritzalis, S. Furnell, and M. Theoharidou, Eds., Springer Berlin Heidelberg, pp. 235–248.

[17] Barbaro, M., Zeller, T., and Hansell, S. A face is exposed for AOL searcher No. 4417749. *The New York Times* (August 9 2006). https://www.nytimes.com/2006/08/09/technology/09aol.html.

[18] Barkhuus, L., and Dey, A. K. Location-based services for mobile telephony: a study of users' privacy concerns. In *Interact* (2003), vol. 3, pp. 702–712.

[19] Barnes, S. B. A privacy paradox: Social networking in the United States. *First Monday 11*, 9 (2006).

[20] Bartow, A. A feeling of unease about privacy law. *University of Pennsylvania Law Review 155* (2006), 52–63.

[21] Bélanger, F., and Xu, H. The role of information systems research in shaping the future of information privacy. *Information Systems Journal 25*, 6 (2015), 573–578.

[22] Bennett, C. J. In defence of privacy: The concept and the regime. *Surveillance & Society 8*, 4 (2011), 485.

[23] Bok, S. *Secrets: On the ethics of concealment and revelation.* Oxford University Press, 1984.

[24] Brooks, S., Garcia, M., Lefkovitz, N., Lightman, S., and Nadeau, E. NISTIR 8062: An introduction to privacy engineering and risk managment in federal systems. https://doi.org/10.6028/NIST.IR.8062, January 2017.

[25] Bygrave, L. A. Hardwiring privacy. In *The Oxford Handbook of Law, Regulation, and Technology*, R. Brownsword, E. Scotford, and K. Yeung, Eds. Oxford University Press, 2017, ch. 31, pp. 754–775.

[26] Calo, R. The boundaries of privacy harm. *Indiana Law Journal 86* (2011), 1131–1162.

[27] Campbell, K., Gordon, L. A., Loeb, M. P., and Zhou, L. The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security 11*, 3 (2003), 431–448.

[28] Cate, F. H. The EU data protection directive, information privacy, and the public interest. *Iowa Law Review 80* (1994), 431–443.

[29] Cavoukian, A., Taylor, S., and Abrams, M. E. Privacy by design: essential for organizational accountability and strong business practices. *Identity in the Information Society 3*, 2 (2010), 405–413.

[30] Ceross, A. Examining data protection enforcement actions through qualitative interviews and data exploration. *International Review of Law, Computers & Technology 32*, 1 (2018), 99–117.

[31] Ceross, A., and Simpson, A. C. The use of data protection regulatory actions as a data source for privacy economics. In *Computer Safety, Reliability, and Security (SAFECOMP)* (2017), S. Tonetta, E. Schoitsch, and F. Bitsch, Eds., vol. 10489 of *Lecture Notes in Computer Science (LNCS)*, Springer, pp. 350–360.

[32] Citron, D. K., Henry, L. M., and Solove, D. J. Visionary pragmatism and the value of privacy in the twenty-first century. *Michigan Law Review 108* (2010), 1107–1126.

[33] Cox, L. A. What's wrong with risk matrices? *Risk Analysis 28*, 2 (2008), 497–512.

[34] Danezis, G., Domingo-Ferrer, J., Hansen, M., Hoepman, J.-H., Metayer, D. L., Tirtea, R., and Schiffner, S. Privacy and data protection by design – from policy to engineering. https://www.enisa.europa.eu/publications/privacy-and-data-protection-by-design, January 2015.

[35] Deng, M., Wuyts, K., Scandariato, R., Preneel, B., and Joosen, W. A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering 16*, 1 (2011), 3–32.

[36] Descheemaeker, E. The harms of privacy. *Journal of Media Law 7*, 2 (2015), 278–306.

[37] Díaz, C., Seys, S., Claessens, J., and Preneel, B. Towards measuring anonymity. In *Privacy Enhancing Technologies* (2003), R. Dingledine and P. Syverson, Eds., vol. 2482 of *Lecture Notes in Computer Science (LNCS)*, Springer, pp. 54–68.

[38] Dienlin, T., and Trepte, S. Is the privacy paradox a relic of the past? An in-depth analysis of privacy attitudes and privacy behaviors. *European Journal of Social Psychology 45*, 3 (2015), 285–297.

[39] Dinev, T., and Hart, P. An extended privacy calculus model for e-commerce transactions. *Information Systems Research 17*, 1 (2006), 61–80.

[40] Dwork, C. Differential privacy. In *Automata, Languages and Programming*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., vol. 4052 of *Lecture Notes in Computer Science (LNCS)*. Springer, 2006, pp. 1–12.

[41] Edman, M., and Yener, B. On anonymity in an electronic society: A survey of anonymous communication systems. *ACM Computing Surveys (CSUR) 42*, 1 (2009), 5.

[42] Egelman, S., and Peer, E. The myth of the average user: Improving privacy and security systems through individualization. In *Proceedings of the 2015 New Security Paradigms Workshop* (2015), NSPW '15, ACM, pp. 16–28.

[43] Finn, R. L., Wright, D., and Friedewald, M. Seven types of privacy. In *European Data Protection: Coming of age*, S. Gutwirth, R. Leenes, P. de Hert, and Y. Poullet, Eds. Springer, 2013, pp. 3–32.

[44] Floridi, L. Open data, data protection, and group privacy. *Philosophy & Technology 27*, 1 (2014), 1–3.

[45] Gandy Jr, O. H. *The Panoptic Sort: A Political Economy of Personal Information. Critical Studies in Communication and in the Cultural Industries.* ERIC, 1993.

[46] Garfinkel, S. L. De-identification of personal information. *National Institute of Science and Technology Internal Report 8053* (2015).

[47] Gavison, R. Privacy and the limits of law. *The Yale Law Journal 89*, 3 (1980), 421–471.

[48] Greenleaf, G. Data protection in a globalised network. In *Research Handbook on Governance of the Internet*, I. Brown, Ed. Edward Elgar Publishing, 2013, pp. 221–259.

[49] Gürses, S., Troncoso, C., and Diaz, C. Engineering privacy by design. *Computers, Privacy & Data Protection 14* (2011), 3.

[50] Hansson, S. O., and Aven, T. Is risk analysis scientific? *Risk Analysis 34*, 7 (2014), 1173–1183.

[51] Heitzenrater, C. D., and Simpson, A. C. Policy, statistics and questions: Reflections on UK cyber security disclosures. *Journal of Cybersecurity 2*, 1 (2016), 43–56.

[52] Hong, J. I., Ng, J. D., Lederer, S., and Landay, J. A. Privacy risk models for designing privacy-sensitive ubiquitous computing systems. In *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques* (2004), DIS '04, ACM, pp. 91–100.

[53] Houghton, D. J., and Joinson, A. N. Privacy, social network sites, and social relations. *Journal of Technology in Human Services 28*, 1-2 (2010), 74–94.

[54] Hustinx, P. The role of data protection authorities. In *Reinventing Data Protection?*, S. Gutwirth, Y. Poullet, P. De Hert, C. de Terwange, and S. Nouwt, Eds. Springer, 2009, pp. 131–137.

[55] Information Commissioner's Office. Guide to the General Data Protection Regulation (GDPR). https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/, February 2018.

[56] Kehr, F., Kowatsch, T., Wentzel, D., and Fleisch, E. Blissfully ignorant: The effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal 25*, 6 (2015), 607–635.

[57] Kenny, S., and Borking, J. The value of privacy engineering. *The Journal of Information, Law and Technology*, 1 (2002).

[58] Koops, B.-J., and Leenes, R. Privacy regulation cannot be hardcoded. a critical comment on the 'privacy by design'provision in data-protection law. *International Review of Law, Computers & Technology 28*, 2 (2014), 159–171.

[59] Kramer, A. D. I., Guillory, J. E., and Hancock, J. T. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences 111*, 24 (2014), 8788–8790.

[60] Lahlou, S. Identity, social status, privacy and face-keeping in digital society. *Social Science Information 47*, 3 (2008), 299–330.

[61] Lange, R., and Burger, E. W. Long-term market implications of data breaches, not. *Journal of Information Privacy and Security 13*, 4 (2017), 186–206.

[62] Laufer, R. S., and Wolfe, M. Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of Social Issues 33*, 3 (1977), 22–42.

[63] Le Métayer, D. A formal privacy management framework. In *Formal Aspects in Security and Trust (FAST)* (2008), P. Degano, J. Guttman, and F. Martinelli, Eds., vol. 5491 of *Lecture Notes in Computer Science (LNCS)*, Springer, pp. 162–176.

[64] Lenhard, J., Fritsch, L., and Herold, S. A literature study on privacy patterns research. In *2017 43rd Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (Aug 2017), pp. 194–201.

[65] Lessig, L. Code is law. *The Industry Standard 18* (1999).

[66] Li, W., Azar, P., Larochelle, D., Hill, P., and Lo, A. W. Law is code: A software engineering approach to analyzing the United States Code. *Journal of Business & Technology Law 10* (2015), 297.

[67] Liu, Y., Gummadi, K. P., Krishnamurthy, B., and Mislove, A. Analyzing Facebook privacy settings: User expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference* (2011), IMC '11, ACM, pp. 61–70.

[68] Lyon, D. Surveillance as social sorting: Computer codes and mobile bodies. In *Surveillance as Social Sorting*, D. Lyon, Ed. Routledge, 2005, ch. 1, pp. 13–30.

[69] Margulis, S. T. Conceptions of privacy: Current status and next steps. *Journal of Social Issues 33*, 3 (1977), 5–21.

[70] Margulis, S. T. Privacy as a social issue and behavioral concept. *Journal of Social Issues 59*, 2 (2003), 243–261.

[71] Marx, G. T. Privacy is not quite like the weather. In *Privacy Impact Assessment* (2012), D. Wright and P. de Hert, Eds., Dordrecht Springer, pp. v–xiv.

[72] Narayanan, A., and Shmatikov, V. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy* (May 2008), pp. 111–125.

[73] Negley, G. Philosophical views on the value of privacy. *Law and Contemporary Problems 31*, 2 (1966), 319–325.

[74] Nehf, J. P. Recognizing the societal value in information privacy. *Washington Law Review 78*, 1 (2003), 1–92.

[75] Nissenbaum, H. The meaning of anonymity in an information age. *The Information Society 15*, 2 (1999), 141–144.

[76] Nissenbaum, H. Privacy as contextual integrity. *Washington Law Review 79*, 1 (2004), 119–158.

[77] Nissenbaum, H. A contextual approach to privacy online. *Daedalus 140*, 4 (2011), 32–48.

[78] Nosko, A., Wood, E., and Molema, S. All about me: Disclosure in online social networking profiles: The case of Facebook. *Computers in Human Behavior 26*, 3 (2010), 406–418.

[79] Notario, N., Crespo, A., Martín, Y. S., Alamo, J. M. D., Métayer, D. L., Antignac, T., Kung, A., Kroener, I., and Wright, D. PRIPARE: Integrating privacy best practices into a privacy engineering methodology. In *2015 IEEE Security and Privacy Workshops* (May 2015), pp. 151–158.

[80] O'Hara, K. The seven veils of privacy. *IEEE Internet Computing 20*, 2 (2016), 86–91.

[81] O'Hara, K., Whitley, E., and Whittall, P. Avoiding the jigsaw effect: Experiences with Ministry of Justice reoffending data. https://eprints.lse.ac.uk/45214/, 2011.

[82] Pandurangan, V. On taxis and rainbows. https://tech.vijayp.ca/of-taxis-and-rainbows-f6bc289679a1#.wq2gtd7ot, June 2014.

[83] Pfitzmann, A., and Hansen, M. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management. http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, Aug. 2010. v0.34.

[84] Ponemon Institute. Cost of Data Breach Study: United Kingdom. https://www-03.ibm.com/security/uk-en/data-breach/, 2017.

[85] Posner, R. A. Privacy, secrecy, and reputation. *Buffalo Law Review 28*, 1 (1978), 1–56.

[86] Qian, H., and Scott, C. R. Anonymity and self-disclosure on weblogs. *Journal of Computer-Mediated Communication 12*, 4 (2007), 1428–1451.

[87] Regan, P. M. *Legislating privacy: Technology, social values, and public policy.* Univ of North Carolina Press, 1995.

[88] Regan, P. M. Response to Bennett: Also in defense of privacy. *Surveillance & Society 8*, 4 (2011), 497–499.

[89] Romanosky, S. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity 2*, 2 (2016), 121–135.

[90] Ruoti, S., Kim, N., Burgon, B., van der Horst, T., and Seamons, K. Confused Johnny: When automatic encryption leads to confusion and mistakes. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (2013), SOUPS '13, ACM, pp. 5:1–5:12.

[91] Samarati, P., and Sweeney, L. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Tech. rep., Technical report, SRI International, 1998.

[92] Sandhu, R. S., Coyne, E. J., Feinstein, H. L., and Youman, C. E. Role-based access control models. *Computer 29*, 2 (1996), 38–47.

[93] Schaar, P. Privacy by design. *Identity in the Information Society 3*, 2 (2010), 267–274.

[94] Schmidt, A. Implicit human computer interaction through context. *Personal Technologies 4*, 2-3 (2000), 191–199.

[95] Shapiro, S. S. Privacy risk analysis based on system control structures: Adapting system-theoretic process analysis for privacy engineering. In *IEEE Security and Privacy Workshops (SPW)* (May 2016), pp. 17–24.

[96] Sheng, S., Broderick, L., Koranda, C. A., and Hyland, J. J. Why Johnny still can't encrypt: Evaluating the usability of email encryption software. In *Proceedings of the 2006 Symposium On Usable Privacy and Security* (2006), SOUPS '06, pp. 3–4.

[97] Simpson, A. C. On privacy and public data: A study of data.gov.uk. *Journal of Privacy & Confidentiality 3*, 1 (2011), 51–65.

[98] Smith, H. J., Dinev, T., and Xu, H. Information privacy research: an interdisciplinary review. *MIS Quarterly 35*, 4 (2011), 989–1016.

[99] Solove, D. J. A taxonomy of privacy. *University of Pennsylvania Law Review* (2006), 477–564.

[100] Spencer, S. B. Reasonable expectations and the erosion of privacy. *San Diego Law Review 39* (2002), 843.

[101] Spiekermann, S., and Cranor, L. F. Engineering privacy. *IEEE Transactions on Software Engineering 35*, 1 (2009), 67–82.

[102] Steeves, V. M. *Lessons from the Identity Trail : Anonymity, Privacy and Identity in a Networked Society.* Oxford University Press, 2009, ch. 11. Reclaiming the Social Value of Privacy, pp. 191–208.

[103] Sweeney, L. *k*-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems 10*, 05 (2002), 557–570.

[104] Tavani, H. T. Philosophical theories of privacy: Implications for an adequate online privacy policy. *Metaphilosophy 38*, 1 (2007), 1–22.

[105] Taylor, L., Floridi, L., and Van der Sloot, B. *Group privacy: New challenges of data technologies*, vol. 126. Springer, 2016.

[106] Trilateral Research & Consulting. Privacy impact assessment and risk managment: Report for the Information Commissioner's Office. https://ico.org.uk/media/1042196/trilateral-full-report.pdf, May 2013.

[107] Tudor, C., Cornish, G., and Spicer, K. Intruder testing on the 2011 UK census: Providing practical evidence for disclosure protection. *Journal of Privacy and Confidentiality 5*, 2 (2014), 3.

[108] Turow, J., and Hennessy, M. Internet privacy and institutional trust: Insights from a national survey. *New Media & Society 9*, 2 (2007), 300–318.

[109] US Secretary's Advisory Committee on Automated Personal Data Systems. Records, computers and the rights of citizens. Tech. rep., Office of the Assistant Secretary for Planning and Evaluation, 1973. https://aspe.hhs.gov/report/records-computers-and-rights-citizens.

[110] Van Audenhove, L., Constantelou, A., Poel, M., van Lieshout, M., Kool, L., van Schoonhoven, B., and de Jonge, M. Privacy by design: an alternative to existing practice in safeguarding privacy. *info 13*, 6 (2011), 55–68.

[111] Van Dijck, J. Datafication, dataism and dataveillance: Big data between scientific paradigm and ideology. *Surveillance & Society 12*, 2 (2014), 197.

[112] Van Dijk, M., Gentry, C., Halevi, S., and Vaikuntanathan, V. Fully homomorphic encryption over the integers. In *Advances in Cryptology – EUROCRYPT 2010* (2010), H. Gilbert, Ed., vol. 6110 of *Lecture Notes in Computer Science (LNCS)*, Springer, pp. 24–43.

[113] van Rest, J., Boonstra, D., Everts, M., van Rijn, M., and van Paassen, R. Designing privacy-by-design. In *Privacy Technologies and Policy: First Annual Privacy Forum (AFP)* (2014), B. Preneel and D. Ikonomou, Eds., vol. 8319 of *Lecture Notes in Computer Science (LNCS)*, Springer, pp. 55–72.

[114] Wacks, R. *Privacy and Media Freedom.* Oxford University Press, 2013.

[115] Wadhwa, K., and Rodrigues, R. Evaluating privacy impact assessments. *Innovation: The European Journal of Social Science Research 26*, 1-2 (2013), 161–180.

[116] Walton, R. E. Social choice in the development of advanced information technology. *Human Relations 35*, 12 (1982), 1073–1083.

[117] Warren, C., and Laslett, B. Privacy and secrecy: A conceptual comparison. *Journal of Social Issues 33*, 3 (1977), 43–51.

[118] Warren, S. D., and Brandeis, L. D. The right to privacy. *Harvard Law Review 4* (1890), 193–220.

[119] Watts, D. Should social science be more solution-oriented? *Nature Human Behaviour 1* (2017), 0015.

[120] Westin, A. F. *Privacy and Freedom.* The Bodley Head, 1967.

[121] Whitman, J. Q. The two western cultures of privacy: Dignity versus liberty. *Yale Law Journal 113*, 6 (2004), 1151–1221.

[122] Whitten, A., Tygar, J. D., Whitten, A., and Tygar, J. D. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the 8th Conference on USENIX Security Symposium - Volume 8* (Berkeley, CA, USA, 1999), SSYM'99, USENIX Association, p. 14.

[123] Wiener, N. *The human use of human beings: Cybernetics and society.* No. 320. Perseus Books Group, 1988.

[124] Wiese Schartum, D. Making privacy by design operative. *International Journal of Law and Information Technology 24*, 2 (2016), 151–175.

[125] Wright, D. Should privacy impact assessments be mandatory? *Communications of the ACM 54*, 8 (2011), 121–131.

[126] Xu, H., Dinev, T., Smith, J., and Hart, P. Information privacy concerns: Linking individual perceptions with institutional privacy assurances. *Journal of the Association for Information Systems 12*, 12 (2011), 798.

[127] Yao, M. Z., Rice, R. E., and Wallis, K. Predicting user concerns about online privacy. *Journal of the Association for Information Science and Technology 58*, 5 (2007), 710–722.

[128] Yu, X., and Wen, Q. A view about cloud data security from data life cycle. In *2010 International Conference on Computational Intelligence and Software Engineering* (Dec 2010), pp. 1–4.