Developing Privacy Incident Responses to Combat Information Warfare

Sean McElroy and Lisa McKee
Dakota State University, Madison, USA

me@seanmcelroy.com
lisa.mckee@trojans.dsu.edu

Abstract: Violations of privacy harm real people, and as nation-state actors grow their information warfare capabilities, civilians suffer these harms as part of coordinated and targeted actions on objectives. When privacy harms manifest, they allow threat actors to injure data subjects by weaponizing their information to harm individuals, communities, and societies. These attacks injure civilians as the confidence of legitimate authorities, institutions, and defences is eroded, and consequences may impact national security. Distinct from cybersecurity, privacy depends upon confidentiality, integrity, and availability but encompasses a unique set of concerns. Whereas security incident response has an established practice and research history, approaches to privacy incident response, such as unauthorized disclosure, are not well researched or documented in academic literature in the unique context of privacy. By mapping privacy harm to techniques and tactics, a cohesive framework emerges to distinguish tailored mitigation strategies for each. This paper proposes a conceptual model and classification framework for privacy-related harms, tactics, techniques, and mitigation strategies to address sophisticated privacy threat actors. Using this model and framework, contingency planners can develop privacy incident response strategies to defend against the privacy harms of information warfare.

Keywords: incident response, data privacy, information warfare, disinformation

1. Introduction

Attacks on data privacy can have a paralyzing effect on a civilian population. However, while information security response strategies are well established in the literature, privacy is often only considered a secondary concern in a security incident. While many detective and preventative security controls are rooted in technical tooling and implementation, security practitioners often lack privacy awareness, privacy engineering skills, or appropriate conceptual models and action frameworks that help them distinguish the unique components of privacy incident preparation and response.

Security incident response plans partly address many privacy harms on a frequency basis, as the highest number of privacy incidents involve a data breach caused by a loss of security. However, many forms of privacy harm, both individual and societal, are not caused by a security incident and may have no technological basis for remediation. Security incident response is an insufficient mechanism for addressing privacy harm. An approach to address privacy by adding it to existing security approaches is incomplete in addressing privacy harms that do not result from data breaches. Distorted propaganda, disinformation, and incitement, for instance, may manifest as a national security concern or, when targeted at a company or individual, substantial economic or personal harm. This research proposes a conceptual model for privacy incident response that addresses the unique domain of privacy and a response framework to address the different harms enumerated in Solove's *A Taxonomy of Privacy*.

2. Literature Review

Privacy harms are distinct from information security consequences. However, the development of models and frameworks expressive enough to capture both is stymied by a false equivalency drawn between security and privacy in fundamental industry standards and practices. For example, the NIST Privacy Framework specifies response controls, but only in the context of shared "cybersecurity-related privacy events" (National Institute of Standards and Technology, 2020, p. 7), while simultaneously acknowledging that failing to address societal risks leads to an incomplete privacy risk management approach. Privacy harm to even individuals can have significant national security consequences, as the rise of widely deployed biometric profiling technology and leaks of highly sensitive personal data included in the 2015 Office of Personnel Management data breach have both harmed US intelligence missions and the country's ability to maintain and gather foreign intelligence partners. (Dorfman, 2020) On a larger scale, societal privacy harms can result in real-world dangers, as their manifestation can "demotivate [military] personnel or create unrest in the [civil] population" (Fritsch and Fischer-Hübner, 2018).

Existing academic literature attempts to identify appropriate models, frameworks, and processes to address privacy incident response. Abdullah proposes a modified conceptual framework incorporating the Forrester GRC

playbook (Abdullah, 2018), but data breaches are the only harm identified for a response. Abdullah references NIST special publication 800–61, NIST special publication 800–122, and the Nymity Privacy Management Accountability Framework as sources for further development of an incident response framework. Both 800-61 and 800-121 detail incident responses from the perspective of a technologist that is tackling a loss of confidentiality (McCallister, Grance and Scarfone, 2010; Cichonski et al., 2012), a principle of information security (McCumber, 2005) that is not inclusive or comprehensive of privacy principles. The Nymity framework includes critical operational processes that may prepare for or prevent non-breach-related privacy incidents; however, the actual incident response activities only cover data breaches (Nymity Inc., 2017). The tendency for industry best practices and guidance as well as academic models tend to address data breaches or unauthorized access resulting in privacy harms in the form of disclosure, but they do not provide a model or framework to address other forms of privacy harm, particularly those instigated by nation-state actors in long-term adversarial campaigns.

Solove's A Taxonomy of Privacy is an informative reference that outlines sixteen unique privacy harms to examine what privacy incidents other conceptual models and incident response frameworks fail to address. (Solove, 2006) Of these sixteen, the literature referenced above addresses at most two of these from Solove's taxonomy: breach of confidentiality and disclosure. Indeed, these are arguably the two most costly and prevalent forms of privacy harm for the industry, with the per-record cost of data breaches hitting an all-time high of $164 in 2022. (Ponemon Institute, 2022)

Data subjects do not experience privacy harm universally in the same way or with the same relative effect. Varying values, beliefs, and attitudes on privacy are not easily classified cross-culturally on geographic or political boundaries (Zabihzadeh et al., 2019). While Solove acknowledges this in his seminal taxonomy (Solove, 2006), the differences in how cultures value privacy and experience subsequent harms necessitate different privacy incident responses commensurate with mitigating the harm that has actually occurred to the extent it affects data subjects. Certainly, harm can be statutorily experienced by data processors and controllers in the form of punitive damages or fines, as prescribed by the EU's General Data Protection Regulation. However, the relative harm experienced by a data subject may vary by cultural, locale, and personal factors. For example, the psychological harm a privacy violation causes the data subject, which research shows is more significant among consumers residing in countries without privacy regulations (Bellman et al., 2004). This result is more significant when harms are perceived to be reputational and cause an economic future impact, such as difficulties in finding housing or employment (Anabo, Elexpuru-Albizuri and Villardón-Gallego, 2019). Blackmail, harm enumerated by Solove, may require different mitigations in a privacy incident response plan for data subjects that experience societal impacts disproportionately to data subjects in social networks that do not value reputation or suggestions of impropriety the same (Al-Saggaf, 2016).

By contrast, Decisional Interference is difficult to both identify and quantify. For example, the Mueller report on Russian interference in the 2016 United States presidential election concluded it occurred "in a sweeping and systematic fashion" (Mueller III, 2019). Subsequently, trust in elective democratic processes in America is significantly declining (Rainie, Keeter and Perrin, 2019). As quantitative macroeconomic research has concluded that lower levels of trust result in lower national economic performance and future growth (Mularska-Kucharek and Brzeziński, 2017), the costs of privacy harm other than data breaches may be both enormously underappreciated as well as lacking commensurate incident response models to protect national security interests. West acknowledges this supposition and further posits additional complications caused by the rise of encrypted communications channels that allow adversaries to conduct disinformation campaigns outside of the reach of public Internet monitoring and community analysis. (West, 2021) The rise of highly automated and expansive disinformation bots on social media coupled with direct messaging capabilities and tailored profiling of the populace may result in significant blackmail, extortion, distortion, and decisional interference Solovian privacy harms materializing. While one may think harms, such as Decisional Interference, as uniquely targeted at a single individual, big data, encryption, and cloud infrastructure may allow these to happen quickly at a vast scale, outside the visibility of intelligence and policing forces.

Returning to the subject of "trust" in the literature, repairing or rebuilding it is the subject of some data-breach-centric incident response frameworks. It is a unique characteristic not present in security incident response frameworks. Wan and Zhang noted the value of trust and analysed existing methods commercial entities use to attempt to increase it. Their analysis revealed that "apology" and "providing information", in the form of answering questions about a privacy breach, were the two most common methods used, with victim monetary or credit monitoring compensation or management action plans as the final two methods observed. (Wan and Zhang, 2014) The efficacy of these methods was not assessed, and it stands to reason that even if they

succeeded in reducing economic harm and increasing trust, they are mainly inapplicable to privacy harms in Solove's taxonomy unrelated to disclosure.

3. Methodology

This paper proposes a conceptual model and outlines a privacy incident response framework using a design science research methodology using the engineering lifecycle (Wieringa, 2014). Design science research allows for developing innovative artefacts to obtain a goal (Simon, 1988; Vaishnavi, 2007). It will require multiple iterations of model design and review with industry professionals and colleagues. The research will include interviews with industry experts to get their input and responses to specific questions to the research question. This question is how can privacy incident response be developed and refined with a unique ontology and approach that best addresses each privacy harm, as identified by Solove.

With a suitable conceptual model developed to approach the question, we propose a preliminary framework that maps mitigations to tactics ("high-level descriptions of behaviour") and techniques ("detailed descriptions of behaviour in the context of a tactic"), as both terms are defined by NIST (Johnson *et al.*, 2016). By leveraging a common vocabulary, the proposed framework appeals to the familiarity of threat modelling and incident response, which are well-developed concepts in information security. It establishes the unique approaches required to build privacy incident response programs.

4. Conceptual Model

Solove frames privacy harms by the underlying activity that creates the issue, culminating in a taxonomy detailing four primary classes: information collection, information processing, information dissemination, and invasion (Solove, 2006, pp. 483–490). This framework is helpful because it focuses on activities to which mitigating responses can be directly mapped. It is also a useful starting point as it focuses, in the first three, on flows into, within, and out of data holders. Privacy incident response must be implemented, at least in part, by data holders, even if at the direction of a national data protection authority.

Consider, however, that the casual elements that create harm may be intentional or accidental. Whether a specific privacy violation is legal in any given jurisdiction, that it is in any suggests intentional privacy threat actors possess *mens rea* or a culpable mental state. American law suggests four types, in descending order of culpability: "purpose, knowledge, recklessness, and negligence" (Kagan, 2021). While this research does not intend to provide a response framework that adheres to any specific legal framework, considering the diverse ways intentional privacy harm can occur is essential to model privacy threats and build a comprehensive incident response plan. Intentional actors may exhibit persistence and adaptation and be highly skilled and well-resourced if supported by a nation-state actor. The conceptual model for intentional actors causing harm is illustrated below in Figure 1. While goal formation may be less evident for reckless and negligence cases, we posit they are still discernible in a culpable actor's outcomes.

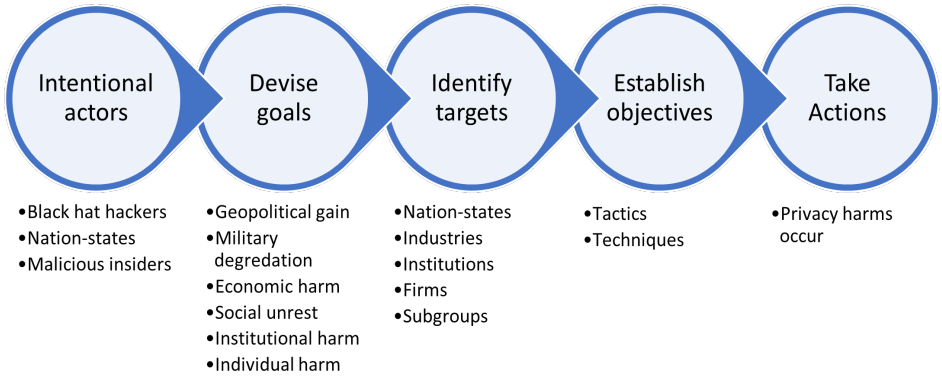

Figure 1: Conceptual Model for Intentional Privacy Harms

For accidentally created harm, we consider such accidents possible because of a failure to design a system with sufficient privacy-by-design safeguards. Cavoukian's Privacy by Design approach is further developed by Hoepman in the four technical data-oriented and process-oriented strategies (Cavoukian, 2011, 2016; Hoepman, 2020). As such, failures to minimize, separate, abstract, hide, inform, control, enforce, or demonstrate privacy in a system allow for accidental harm to occur in our conceptual model. The lack of intent does not make

accidentally created harm any less severe. While design flaws may not exhibit the same type of persistence and adaptation as intentional privacy threat actors, they may be pervasive and require considerable time and effort to design, implement, and deploy appropriate mitigations. In addition to privacy design flaws, many systems depend on security properties to deliver privacy properties. A system that correctly engineers privacy protections but has failing security controls may still result in accidental harm in this model; that is, we consider intentional attacks on security that consequently result in privacy harm within this second conceptual model. As illustrated below in Figure 2, techniques are still evident, though, without the element of intent, they represent how potential builds and not a specific planned action.

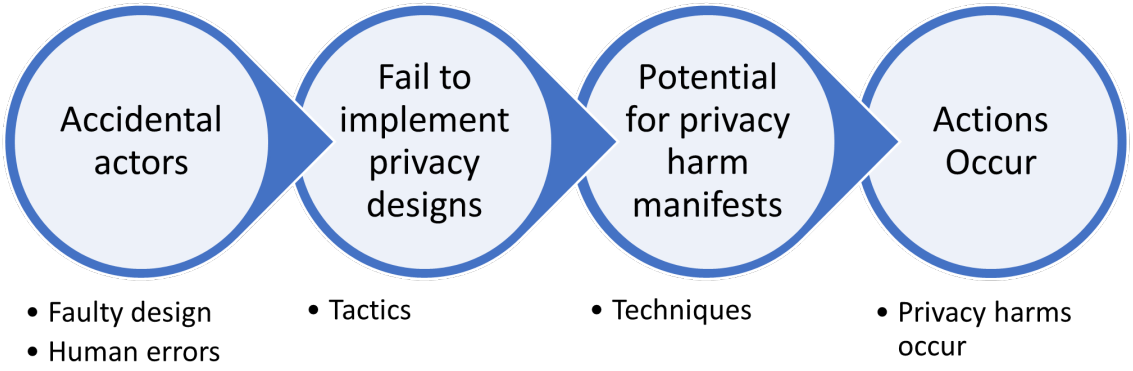

Figure 2: Conceptual Model for Accidental Privacy Harms

5. Privacy Incident Response Framework

Following the conceptual models above, a privacy incident response framework must allow for a repeatable process for incident responders to minimize privacy harm. Two popular security incident response frameworks denote similar processes in an adjacent space. The SANS Incident Response Plan enumerates six processes: preparation, identification, containment, eradication, recovery, and lessons learned (Kral, 2011; SANS Institute, 2016). By comparison, the NIST Cyber Security Framework enumerates four distinct procedural steps: (1) preparation, (2) detection and analysis, (3) containment, eradication, and recovery, and (4) post-incident activity (Cichonski *et al.*, 2012). However, unlike security incident response, where the goal is to restore the security properties of a system, the goal of privacy incident response is to minimize harm on all timescales. A Pandora's Box analogy exemplifies this fundamental difference: some privacy harms have lasting consequences that no specific activity can undo. No 'server rebuild' or gag order can restore privacy once knowledge about communication contents or associations of a data subject is disclosed to the public. Consequently, privacy incident response processes must differ to address this disparity.

5.1 Privacy Incident Response Plan (PIRP)

As discussed, privacy incident responses often do not contain harms that have already occurred, though they may contain new instances harming additional data subjects or prevent the expansion of harms (such as disclosure growing to include aggregation) from happening to the same affected subjects. For this reason, containment and eradication, terms used in security response planning, have no clear analogues in privacy incident response. Furthermore, harm recovery may take months, even years, and so privacy incident responses operate on significantly different timelines. This reality changes the nature of incident response management, resourcing, and tooling required to manage short- and long-term outcomes successfully. With this in mind, we propose the following process framework for this response, a privacy incident response plan (PIRP), illustrated below in Figure 3.

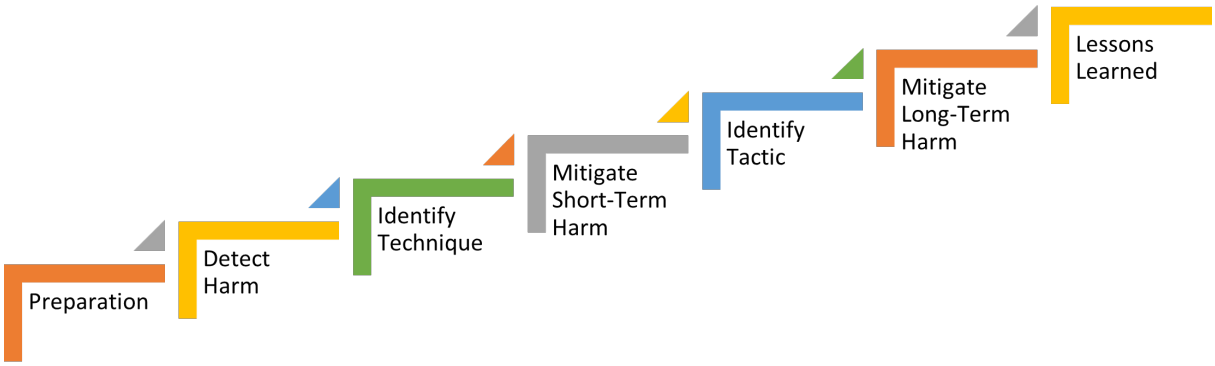

Figure 3: Privacy Incident Response Plan (PIRP) framework

5.1.1 Preparation

Preparing for using the PIRP requires knowledgeable and skilled resources who understand privacy concepts, the current legal environment, and the human impacts of specific privacy harms. PIRP plan updates and ongoing readiness training through tabletop exercises or simulated incident response scenarios prepare both incident responders and dependent legal and executive authorities to perform their duties in the face of an actual event.

5.1.2 Detect Harm

Detection mechanisms are required to initiate all subsequent plan actions. For privacy incidents, such as breach of confidentiality or unauthorized disclosure that stem from security incidents, the declaration of a security incident should also trigger a presumptive invocation of the PIRP. Other privacy-specific inputs may detect historical or ongoing harm, including direct reports from data subjects and detect leakage or use of mock or real data subject records in unexpected circumstances. Incident responders must complete all other plan steps when harm is detected and confirmed.

5.1.3 Identify Technique

Because a PIRP focuses on minimizing harm, privacy incident responders need to quickly understand the "how" before the "why" of harm occurring. An entire ontology of techniques that can cause any of the 16 Solovian privacy harms would be extensive to enumerate and describe and is outside this research's scope. However, a contingency planner could document such a technique in the scope of a nation-state PIRP, as illustrated in Table 1.

Table 1: Example technique to harm privacy

ID	Name	Description
TE0001	Disclosure Through Adjacent Queries	Adversaries may learn large amounts of private information by abusing targeted search functions in systems. *Example*: An application that allows one to find individuals on a dating website within 5 miles of a given street address could automate queries to exfiltrate the population of data subjects and pinpoint exact street addresses through triangulation of responses.

Identifying illustrative example TE0001 requires consideration in the Preparation phase of the PIRP, provided a privacy threat model determined it was a relevant privacy threat to an organization or system. Such harm detection mechanisms could be an observation of a rise in queries from a singular IP address or an anomaly in query volume that is uncharacteristic.

5.1.4 Mitigate Short-Term Harm

Individual techniques deserve specific mitigating treatments tailored to the technique. The purpose of this step is similar to 'containment' in a security incident response plan. It is to quickly stop new data subjects from being affected by this technique and continued manifestations of the technique for already-affected data subjects.

Each technique catalogued by a contingency planner and accounted for in a PIRP should have a range of mitigations. For illustrated technique TE0001, these treatments could include, as an example:

1. M0001 Rate-limit query features that return data subject object identifiers
2. M0002 Disable the query feature that returns data subject object identifiers
3. M0003 Disable access to the system providing the query feature that returns data subject object identifiers

5.1.5 Identify Tactic

Understanding the high-level descriptions of behaviour that caused privacy harm is an activity similar to root cause analysis. By viewing the specific technique that caused injury through the lens of overall behaviours, an incident responder can identify approaches for mitigating long-term harm. This is where two different conceptual models for PIRP are essential, as accidental causes are more likely related to systemic design or process problems. In contrast, intentional actors are more likely to manifest varied future harms or magnify damage through specific, additional actions.

While a comprehensive list of intentional and accidental tactics is beyond the scope of this research, an illustrative example of privacy harm tactics perpetuated by intentional actors is detailed below in Table 2.

Table 2: Example of intentional tactic to harm privacy

ID	Name	Description
TI0001	Disruptive Dissemination	The adversary is trying to harm the availability of a whole system, its components or features, or cause delays that degrade its service. *Example*: An actor scrapes data from a system and publishes an enriched version correlated with other data to cause embarrassment that provokes a reaction that causes a valuable service to be restricted or disabled. This tactic relates to the Solove harms "Aggregation" and "Increased Accessibility" for affected data subjects.

5.1.6 Mitigate Long-Term Harm

Because privacy harms cannot be contained and eradicated in the bounded way many security incidents can be, strategies for addressing long-term harms must be developed at both the technique and tactic levels in a PIRP. Considering a PIRP ontology that parented the example technique "Disclosure Through Adjacent Queries" from Table 1 under the tactic "Disruptive Dissemination" from Table 2, long-term harm mitigations exist for both the general tactic category and the specific technique. Long-term system design considerations at the tactical level might include preparing systems to operate with configurable reduced functionality to reduce the chances that disclosures through some features do not require the whole system to be disabled to remove access to a feature. At a technique level, reducing relationships between data or the visibility of specific fields through a query function may be a longer-term data transformation mitigation (Senarath and Arachchilage, 2019) for this particular vector of privacy harm.

6. Discussion

6.1 Analysis

Considering advanced persistent threat actors can and have sought to inflict privacy harm, we believe this approach sets the appropriate context for combatting related information warfare tactics. When privacy is considered only in terms of security incident response, it is addressed to the extent that security directly provides privacy. In the context of Solove's taxonomy of privacy harms, these are limited to what that research categorizes as 'Information Dissemination' harms. Existing incident response frameworks fail to consider the potential for well-resourced intentional actors who may persistently seek to cause and magnify privacy harms. Subsequently, they do not contemplate long-term harm mitigation strategies. Short-term harm mitigation strategies often include actions of limited efficacy: notice of disclosure, promises of change, and temporary consumer credit services to address consequential identity theft. These novel artefacts address these shortcomings and are conceptually complete in addressing the shortcomings of pre-existing approaches.

6.2 Limitations

This work establishes a conceptual model and privacy incident response framework design artefacts. However, these elements are illustrated but not comprehensively developed as a complete ontology. Incident responders need to model privacy threats and relevant harms and devise harm mitigation strategies suited to how they collect, use, process, and share information.

The model and the resulting framework have not yet been comprehensively tested with the design science methodology against documented privacy incidents to measure the magnitude of their impact or the performance of specific features of either artefact in the context of privacy harm mitigation.

6.3 Future Work

This new contribution of a conceptual model and privacy incident response plan framework would benefit from formal design science model testing. While the meta-requirements, meta-design, and theories are established in this paper, formal model testing that validates the use of these artefacts results in better privacy incident response outcomes, as proposed for this research methodology (Kuechler and Vaishnavi, 2008), for both short and long term harm reduction are essential areas for future work.

Techniques, tactics, and mitigation strategies are illustrated in this work. Developing a framework library repository of these artefact dimensions would significantly improve the usability of these artefacts by practitioners. For instance, a community-driven approach to soliciting feedback to develop these areas may yield significant coverage for each of the Solove privacy harms beyond what any single research group may develop. If informed by analysis of historical privacy incidents or future applications, subsequent technical action research may significantly add to the impact through iterative refinement.

Finally, operationalizing this model and framework for practitioners through developing PIRP awareness, education, and training materials could include tutorials and guides usable for those with limited experience in incident response or data privacy. Promulgating these findings through industry standards bodies, such as NIST and ANSI, may enhance existing privacy protection guidance and offer new incident response processes and procedures to existing standards.

7. Conclusions

Privacy incident response as a holistic program that addresses long-term and short-term harms using a systematic approach is not established in the literature apart from the context of security incident response or “disclosure” privacy harm. Indeed, data breaches are increasing in frequency and magnitude. However, security incident response plans often neglect the long-term effects of such disclosure, which include blackmail, extortion, distortion, and decisional interference. Using Solove’s ‘Taxonomy of Privacy’ provides a valuable framework to identify harms unaccounted for by existing incident response plans. Further, considering data privacy harms are not only caused by accidents but could be the underlying strategic objective of intentional threat actors, mapping tactics, techniques, and harms and enumerating specific short and long-term response strategies is necessary to address the lasting consequences of privacy harm. The conceptual model and privacy incident response framework presented in this design science research offer new tools that allow incident responders to address data privacy concerns directly.

References

Abdullah, H. (2018) ‘A Structural and Navigational Method for Integrated Organizational Information Privacy Protection’, in *2018 International Conference on Advances in Big Data, Computing and Data Communication Systems (icABCD). 2018 International Conference on Advances in Big Data, Computing and Data Communication Systems (icABCD)*, Durban, South Africa: IEEE, pp. 1–9. Available at: https://doi.org/10.1109/ICABCD.2018.8465138.

Al-Saggaf, Y. (2016) ‘An Exploratory Study of Attitudes Towards Privacy in Social Media and the Threat of Blackmail: The Views of a Group of Saudi Women’, *The Electronic Journal of Information Systems in Developing Countries*, 75(1), pp. 1–16. Available at: https://doi.org/10.1002/j.1681-4835.2016.tb00549.x.

Anabo, I.F., Elexpuru-Albizuri, I. and Villardón-Gallego, L. (2019) ‘Revisiting the Belmont Report’s ethical principles in internet-mediated research: perspectives from disciplinary associations in the social sciences’, *Ethics and Information Technology*, 21(2), pp. 137–149. Available at: https://doi.org/10.1007/s10676-018-9495-z.

Bellman, S. *et al.* (2004) ‘International Differences in Information Privacy Concerns: A Global Survey of Consumers’, *The Information Society*, 20(5), pp. 313–324. Available at: https://doi.org/10.1080/01972240490507956.

Cavoukian, A. (2011) 'The 7 Foundational Principles - Implementation and Mapping of Fair Information Practices'. Information and Privacy Commissioner of Ontario. Available at: https://iapp.org/media/pdf/resource_center/pbd_implement_7found_principles.pdf.

Cavoukian, A. (2016) 'The Time to Embed Privacy, by Design is Now: Into IoT, AI, and Big Data'. *SAIConference*, London, England, 4 August. Available at: https://www.youtube.com/watch?v=mzk9Ijw9O_o (Accessed: 30 August 2021).

Cichonski, P. *et al.* (2012) *Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology*. NIST SP 800-61r2. National Institute of Standards and Technology, p. NIST SP 800-61r2. Available at: https://doi.org/10.6028/NIST.SP.800-61r2.

Dorfman, Z. (2020) 'China Used Stolen Data to Expose CIA Operatives in Africa and Europe', *Foreign Policy*, 21 December. Available at: https://foreignpolicy.com/2020/12/21/china-stolen-us-data-exposed-cia-operatives-spy-networks/.

Fritsch, L. and Fischer-Hübner, S. (2018) 'Implications of Privacy & Security Research for the Upcoming Battlefield of Things', *Journal of Information Warfare*, 17(4), pp. 72–87.

Hoepman, J.-H. (2020) *Privacy Design Strategies (The Little Blue Book)*.

Johnson, C.S. *et al.* (2016) *Guide to Cyber Threat Information Sharing*. NIST SP 800-150. National Institute of Standards and Technology, p. NIST SP 800-150. Available at: https://doi.org/10.6028/NIST.SP.800-150.

Kagan, E. (2021) *Charles Borden, Jr. v. United States*.

Kral, P. (2011) 'Incident Handler's Handbook'. SANS Institute. Available at: https://sansorg.egnyte.com/dl/6Btqoa63at (Accessed: 1 November 2021).

Kuechler, B. and Vaishnavi, V. (2008) 'On theory development in design science research: anatomy of a research project', *European Journal of Information Systems*, 17(5), pp. 489–504. Available at: https://doi.org/10.1057/ejis.2008.40.

McCallister, E., Grance, T. and Scarfone, K.A. (2010) *Guide to protecting the confidentiality of Personally Identifiable Information (PII)*. 0 edn. NIST SP 800-122. Gaithersburg, MD: National Institute of Standards and Technology, p. NIST SP 800-122. Available at: https://doi.org/10.6028/NIST.SP.800-122.

McCumber, J. (2005) *Assessing and managing security risk in IT systems: a structured methodology*. Boca Raton, FL: Auerbach Publications.

Mueller III, R.S. (2019) *Report on the Investigation into Russian Interference in the 2016 Presidential Election*. U.S. Department of Justice, p. 448. Available at: https://www.justice.gov/archives/sco/file/1373816/download (Accessed: 9 September 2021).

Mularska-Kucharek, M. and Brzeziński, K. (2017) 'The Economic Dimension of Social Trust', *European Spatial Research and Policy*, 23(2), pp. 83–95. Available at: https://doi.org/10.1515/esrp-2016-0012.

National Institute of Standards and Technology (2020) *NIST Privacy Framework: A Tool for Improving Privacy Through Enterprise Risk Management Version 1.0*. NIST CSWP 01162020. Gaithersburg, MD: National Institute of Standards and Technology. Available at: https://doi.org/10.6028/NIST.CSWP.01162020.

Nymity Inc. (2017) 'Nymity Privacy Management Accountability Framework'. Available at: https://vvena.nl/wp-content/uploads/2018/04/PMAF-Poster.pdf (Accessed: 11 September 2022).

Ponemon Institute (2022) *Cost of a Data Breach Report 2022*. Armonk, NY: IBM Security, p. 59. Available at: https://www.ibm.com/downloads/cas/3R8N1DZJ.

Rainie, L., Keeter, S. and Perrin, A. (2019) 'Trust and Distrust in America', *Pew Research Center - U.S. Politics & Policy*, 22 July. Available at: https://www.pewresearch.org/politics/2019/07/22/trust-and-distrust-in-america/.

SANS Institute (2016) 'SANS 504-B Incident Response Cycle: Cheat Sheet'. Available at: https://www.sans.org/media/score/504-incident-response-cycle.pdf (Accessed: 19 January 2022).

Senarath, A. and Arachchilage, N.A.G. (2019) 'A data minimization model for embedding privacy into software systems', *Computers & Security*, 87, p. 101605. Available at: https://doi.org/10.1016/j.cose.2019.101605.

Simon, H.A. (1988) 'The Science of Design: Creating the Artificial', *Design Issues*, 4(1/2), p. 67. Available at: https://doi.org/10.2307/1511391.

Solove, D.J. (2006) 'A Taxonomy of Privacy', *University of Pennsylvania Law Review*, 154(3), p. 477.

Vaishnavi, V.K. (2007) *Design Science Research Methods and Patterns: Innovating Information and Communication Technology*. 1st Edition. Auerbach Publications. Available at: https://doi.org/10.1201/9781420059335.

Wan, L. and Zhang, C. (2014) 'Responses to trust repair after privacy breach incidents', *Journal of Service Science Research*, 6(2), pp. 193–224. Available at: https://doi.org/10.1007/s12927-014-0008-2.

West, L.B. (2021) 'Beyond Fighting Words: Reconceptualizing Information Warfare and its Legal Barriers', *National Security Law Journal*, 8(2), p. 162.

Wieringa, R.J. (2014) *Design science methodology for information systems and software engineering*. New York, NY: Springer Berlin Heidelberg.

Zabihzadeh, A. *et al.* (2019) 'Cultural differences in conceptual representation of "Privacy": A comparison between Iran and the United States', *The Journal of Social Psychology*, 159(4), pp. 357–370. Available at: https://doi.org/10.1080/00224545.2018.1493676.