
Information Systems Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

# From Shield to Sword: How Data Privacy Can Undermine Data Security

Alexander Gladis, Torsten-Oliver Salge, David Antons, Nicole Hartwich

To cite this article:
Alexander Gladis, Torsten-Oliver Salge, David Antons, Nicole Hartwich (2026) From Shield to Sword: How Data Privacy Can Undermine Data Security. Information Systems Research

Published online in Articles in Advance 18 Feb 2026

. https://doi.org/10.1287/isre.2023.0566

https://pubsonline.informs.org/journal/isre

INFORMATION SYSTEMS RESEARCH
*Articles in Advance*, pp. 1–18
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

# From Shield to Sword: How Data Privacy Can Undermine Data Security

**Alexander Gladis,[a,*] Torsten-Oliver Salge,[a] David Antons,[b] Nicole Hartwich[a]**

[a] Institute for Technology and Innovation Management, RWTH Aachen University, 52072 Aachen, Germany; [b] Institute for Entrepreneurship, University of Bonn, 53121 Bonn, Germany

*Corresponding author

**Contact:** alexander.gladis@rwth-aachen.de, https://orcid.org/0000-0001-5181-9671 (AG); salge@time.rwth-aachen.de, https://orcid.org/0000-0002-7801-8636 (T-OS); antons@uni-bonn.de, https://orcid.org/0000-0002-2392-0374 (DA); hartwich@time.rwth-aachen.de, https://orcid.org/0000-0001-8568-3231 (NH)

**Received:** September 22, 2023
**Revised:** July 8, 2024; April 3, 2025; July 11, 2025; September 12, 2025
**Accepted:** December 7, 2025
**Published Online in Articles in Advance:** February 18, 2026

https://doi.org/10.1287/isre.2023.0566



**Abstract.** As vital pillars for data protection in our digital age, data privacy and data security coevolve and increasingly attract the attention of researchers, practitioners, and policymakers alike. However, both concepts tend to be studied, managed, and regulated separately, leaving the precise nature of their interplay underexplored. Conceptually, we add clarity by distinguishing and explicating three mutually enriching perspectives: (1) a layered view that sees data security as a prerequisite for data privacy, (2) a complementarity view that sees both as reinforcing each other, and (3) an emerging interference view that sees one as potentially undermining the other. Empirically, we demonstrate the viability of the interference view by examining an identity theft scenario. Here, attackers weaponize individuals' right of access to their own personal data stored by organizations according to Article 15 of the European General Data Protection Regulation. Originally intended to serve as a shield enhancing data privacy, they repurpose the regulation as a sword to exfiltrate others' personal data without authorization as part of an intricate cyberattack. Our analysis of 718 such fraudulent subject access requests distributed across three cases of simulated identity thefts, as well as 21 complementary in-depth expert interviews with data protection officers, reveals that this attack strategy can be highly effective, allowing attackers to gain access even to sensitive personal data such as the victims' national identity card, bank account number, and consumption history. We provide deep insights into the characteristics, enablers, consequences, and mitigation options of this attack strategy and the broader interference view it illustrates. These insights have important implications for our understanding of the interplay between data privacy and data security as well as the unintended consequences of data protection regulations.

**History:** Youngjin Yoo, Senior Editor; Heng Xu, Associate Editor.



**Supplemental Material:** The online appendix is available at https://doi.org/10.1287/isre.2023.0566.

**Keywords:** data protection • data privacy • data security • data protection regulation • GDPR • right of access • social engineering

## 1. Introduction

As the digital transformation permeates our professional and private lives, our personal data are increasingly attractive for—and vulnerable to—exploitation. Organizations collect an ever more extensive and diverse array of personal data ranging from demographic to financial and behavioral data, often to monetize it through data-based business models (Cichy et al. 2021). At the same time, cyberattacks on data of individuals, organizations, and countries continue to grow exponentially and are estimated to cost the global economy up to $90 trillion in 2030 (World Economic Forum 2019). In 2022 in the United States alone, authorities registered more than 800,000 cyberattacks with an economic damage amounting to more than $10 billion (Federal Bureau of Investigation 2022).

As vital pillars for data protection, data privacy and data security coevolve and attract growing attention among scholars, practitioners, and policymakers alike (Smith et al. 2011, Lowry et al. 2017, Biselli and Reuter 2021), culminating in ambitious laws such as the European General Data Protection Regulation (GDPR). Somewhat surprisingly, however, data privacy and data security are still often studied, managed, and

regulated separately (Solove and Hartzog 2022). There is hence a lack of understanding regarding whether and how data privacy and data security are conceptually and practically interconnected in shaping personal data protection. Studies variously—and often implicitly—assume both to be independent, one to be a part of the other, one to be a prerequisite for the other, or both to be mutually reinforcing (Belanger et al. 2002, Biselli and Reuter 2021).

We argue that greater conceptual clarity regarding the relationship between data privacy and data security can fuel notable progress in scholarship, practice, and regulation. We hence seek to shed light on the following research question: "How do data privacy and data security interrelate in shaping personal data protection?" In our work, we add to theory and evidence on data privacy and data security as two distinct, yet interconnected concepts by systematizing prior, often implicit conceptualizations of their interplay and providing novel evidence for an emerging alternative conceptualization. More specifically, we distinguish three mutually enriching perspectives on the interplay between data privacy and data security: (1) the *layered view* that sees data security as a precondition for data privacy; (2) the *complementarity view* that sees data privacy and data security as mutually reinforcing; and (3) the emerging *interference view* that sees data privacy and data security as potentially undermining each other and unintentionally threatening personal data protection.

We provide empirical support for the interference view by demonstrating how efforts to increase individuals' data privacy can undermine their data security, when repurposed by cybercriminals as part of social engineering attacks, that is, cyberattacks exploiting the human element (Mitnick and Simon 2003). More specifically, we examine a scenario where attackers turn the European GDPR from a shield meant to foster individuals' data privacy through greater control over their personal data into a sword—and data security threat—for gaining unauthorized access to those very same data (Di Martino et al. 2019, Bufalieri et al. 2020). For this purpose, they leverage the right of access (Article 15 GDPR) that grants European citizens (*data subjects*) the right to submit to any organization (*data controller*) affected by the GDPR a so-called subject access request (SAR) to gain access to any personal data stored about themselves. We conduct three case studies of carefully simulated SAR identity thefts on three victims involving a total of 718 SARs across three attack waves and demonstrate that attackers who impersonate their victim and send spoofed SARs can exfiltrate sensitive personal data without breaching the targeted organizations' technical security measures. Having no practical experience in social engineering, possessing no prior knowledge about their victims except name and workplace, and bound by certain ethical and legal considerations, our attackers were as weak as reasonably assumable in a realistic scenario—meaning that just about anyone can replicate their attacks. Yet, even under these circumstances, they were able to exfiltrate a broad range of sensitive personal data on our three victims, including home address, phone numbers, utility bills, national identity card and bank account information, and loan financing and insurance data. Even relatively minor data leaks by few organizations (e.g., the victim's date of birth) accumulated quickly and enabled the attackers to substantially improve the persuasiveness of their SARs, thus increasing the effectiveness of subsequent attacks on the same victim. Across the three attack waves, attackers exploited both organization-induced and regulation-induced leaks to exfiltrate similar amounts of personal data for all three victims, even though they differed substantially with regards to their privacy awareness and preferences, with one victim each representing a highly privacy aware person, an average user, and a semipublic figure.

Two primary contributions emerge from these findings. First, we contribute to research on data privacy and data security (Smith et al. 2011, Lowry et al. 2017, Biselli and Reuter 2021, Solove and Hartzog 2022) not only by presenting a conceptual framework on the interplay between both concepts—encompassing the layered, complementarity, and interference views—but also by demonstrating the viability of the so far underexplored interference view. In particular, we show that the GDPR and provisions to strengthen individual data privacy contained therein (here, the right of access according to Article 15 GDPR) can be weaponized to undermine data security, posing a substantial threat to individuals and the protection of their personal data. The interplay between data privacy and data security is hence not limited to one enabling (*layered view*) or adding to (*complementarity view*) the other. Instead, they might at times destructively interfere with each other in often unintended, yet consequential ways (*interference view*). Future research has much to benefit from conceptualizing data privacy and data security as closely interconnected and from explicitly accounting for the multifaceted nature of their interplay.

Second and building on the above, our insights add to the emerging, interdisciplinary literature on the unintended consequences of increasingly prevalent and powerful data protection regulations such as the GDPR, the California Consumer Privacy Act, or the European AI Act. Previous work in this stream has pointed to the often symbolic and compliance-oriented implementation of data protection regulations within organizations (Waldman 2020b) and their negative effects on consumers, such as decreasing consumer surplus due to price increases following GDPR implementation

(Ke and Sudhir 2023, Xu et al. 2025). Our study builds on these insights and shows that data privacy provisions within the GDPR can even be weaponized to undermine data security and, ultimately, the protection of personal data. Somewhat paradoxically, a shield designed to protect personal data can thus become a sword that facilitates the exfiltration of those data during cyberattacks. We establish a lower boundary for the damage potential of this novel attack vector, determine a set of root causes enabling it, and discuss possible threat mitigation options for individuals, organizations, and regulators. This illustrates the practical significance of the interference view and implies that data privacy and data security need to be not only studied, but also managed and regulated in combination. Our insights into the underlying root causes can inform such efforts.

## 2. Conceptual Background

### 2.1. Data Privacy and Data Security as Coevolving Pillars of Data Protection

Data privacy and data security are two fundamental concepts in the realm of personal data protection. Each addresses distinct yet interconnected aspects of protecting personal data (Lowry et al. 2017). Data privacy encompasses individuals' control over how data tied to their identity are gathered, stored, processed, and shared (Smith et al. 2011). Data security focuses on safeguarding those data from unauthorized access and damage (Belanger et al. 2002), following the goals of confidentiality, integrity, and availability of information (von Solms and van Niekerk 2013). As the two pillars of data protection, data privacy and data security have attracted growing attention among researchers, practitioners, and policymakers over the last decades (Lowry et al. 2017).

Research on both concepts largely evolved in parallel. Gradually, research on the nature and predictors of individual privacy (Smith et al. 1996) and security behavior (Cram et al. 2019) was complemented with studies on biases and other limitations in human behavior that might constrain data privacy (Acquisti et al. 2015) and data security (D'Arcy et al. 2014). To better understand possible countermeasures, research started investigating organizational policies and practices such as embracing data privacy and data security by design (Dehling and Sunyaev 2024), and putting in place dedicated units empowered to take full ownership of their mission (Durcikova et al. 2024). As organizations may face tradeoffs between protecting consumers' personal data and exploiting them as part of their data-driven business models (Acquisti et al. 2015), research, consumer advocates, and others have increasingly called for government intervention to ensure personal data protection by regulating data privacy and data security (Waldman 2020b). As a result, ambitious data protection regulations have emerged across the globe, with the European GDPR often seen as the current gold standard that has shaped the design of many other data protection laws worldwide. The GDPR equips Europeans with powerful privacy rights granting greater control over how their personal data are collected, stored, and processed. These rights include the right of access (Article 15 GDPR), the right to rectification (Article 16 GDPR), and the right to erasure (Article 17 GDPR). As a case in point, the right of access allows individuals to submit to any organization falling under the GDPR a SAR to gain access to any personal data stored about themselves. Organizations, in turn, are required not only to appoint a data protection officer (DPO) but also to ensure that privacy rights can be exercised and that data security mandates stipulated in the GDPR are met.

The global influence of the GDPR and its dual emphasis on data privacy and data security as pillars for data protection have attracted substantial research interest especially from information systems, economics, and legal scholars (Johnson 2022). Previous work in this interdisciplinary stream highlighted important beneficial effects, for instance, on consumer confidence in foreign digital products (Li et al. 2025) and consumer surplus in competitive, price-sensitive markets (Ke and Sudhir 2023). However, research has also started to uncover possible negative side effects of data protection regulations such as the GDPR not only on the market level (e.g., increased market and personal data concentration benefiting big tech and hurting smaller firms; Johnson et al. 2023) and firmlevel (e.g., increased compliance and litigation costs; Demirer et al. 2024), but also on the consumer level, where robust evidence is particularly scarce (Johnson 2022). Here, previous modeling work has shown that data protection regulations can trigger especially firms in monopolistic and price-insensitive markets to increase their prices in such a way that overall consumer surplus decreases despite efforts to enhance data privacy and data security (Ke and Sudhir 2023, Xu et al. 2025). Research in this stream has also highlighted that organizations tend to implement data protection regulations in a symbolic and compliance-oriented rather than a substantive and consumer-oriented way (Waldman 2020b). As part of such decoupling efforts, organizations often translate data protection regulations and the privacy and security provisions contained therein into protocols, checklists, and compliance trainings, rather than into tangible changes in product, process, and business model design. This issue is compounded by the fact that data protection professionals are often distributed across departments (e.g., legal and compliance department with responsibility for data privacy and information technology (IT) department with responsibility for

data security) and disconnected from the actual product development units (Waldman 2020b, Solove and Hartzog 2022). Because of these challenges, the original regulatory intention—namely to protect consumers and their personal data—risks being lost in translation, leaving much of the potential of data protection regulations unexploited (Waldman 2020b, Solove and Hartzog 2022).

Given the close coevolution of data privacy and data security as pillars of personal data protection in research, practice, and regulation, the question of how the two concepts interrelate in shaping personal data protection moves into the foreground.

### 2.2. Interplay Between Data Privacy and Data Security

Various—often implicit—conceptualizations of the relationship between data privacy and data security coexist (Pavlou 2011, Smith et al. 2011). Specific conceptualizations include data privacy as part of data security, data security as part of data privacy, both as two dimensions of the same construct, and both as clearly separable constructs (Biselli and Reuter 2021). The persisting conceptual ambiguity about the precise nature of the relationship between data privacy and data security is problematic in that it likely impedes advances in scholarship, practice, and policy alike (Belanger et al. 2002, Biselli and Reuter 2021). We systematize this interplay by proposing three conceptualizations or views on the interaction between data privacy and data security,[1] summarized in Table 1.

We refer to the first and still dominant conceptualization as the *layered view*. Here, data security is seen as a necessary but not sufficient precondition for data privacy (Smith et al. 2011, Acquisti et al. 2016). By ensuring data confidentiality, integrity, and availability, data security establishes the foundations needed for data privacy to build on. Conversely, a lack of data security threatens data privacy and individuals' control over their personal data (Cichy et al. 2021). Data security alone, however, is insufficient to ensure data privacy (Culnan and Williams 2009).

The *complementarity view* sees the interaction between data privacy and data security as mutually reinforcing with improvements in one likely leading to improvements in the other. As Solove and Hartzog (2022) put it, good data privacy measures strengthen data security, and vice versa. For example, data encryption enables individuals' privacy in the form of control over their personal data by ensuring that they can only be accessed by authorized parties. Vice versa, good data privacy practices may boost data security: Systematizing and minimizing the collection and storage of users' personal data, for instance, simplifies the implementation of data security measures and reduces the impact of a potential data breach (Solove and Hartzog 2022).

Our study adds a third conceptualization, which we call the *interference view*. This view sees data security and data privacy as potentially interfering with each other in a destructive and typically unintended manner with potentially far-reaching negative implications for personal data protection. It also captures unforeseen cases where one can be repurposed—with malicious intent—to undermine the other. As a case in point, there is initial evidence that data privacy rights such as the right to access personal data (Article 15 GDPR) can be weaponized for social engineering attacks to exfiltrate others' personal data (Di Martino et al. 2019, Bufalieri et al. 2020). In such an attack scenario, a malicious actor (the attacker) pretends to be another data subject (the victim) and submits a fraudulent SAR in the name of that data subject, attempting to convince the targeted organization to disclose personal data. It is the viability of this attack strategy and the *interference view* it is representing that we will examine in detail in our study.

## 3. Methods

### 3.1. Case Study Design

As Lowry et al. (2017) argue, ecological validity is essential for research in security and privacy. We therefore adopted a multiple case study design (Yin 2009) to simulate an attack that seeks to exploit Article 15 GDPR (the "right of access") under real-world conditions.

**Table 1.** Three Conceptualizations of the Interplay Between Data Privacy and Data Security

| | Layered view | Complementarity view | Interference view |
|---|---|---|---|
| Description | Data security as necessary precondition for data privacy | Data security and data privacy as mutually reinforcing | Data security and data privacy as unintentionally compromising each other |
| Illustration | Authentication techniques preventing unauthorized access as prerequisite for control over personal data | Data encryption enabling data privacy on the Internet; Structured data collection and storage facilitating the implementation of data security mechanisms | Right to access personal data being weaponized for social engineering attacks to exfiltrate personal data and conduct identity theft |
| Logic | Constitutive | Complementary | Compromising |

More specifically, we simulated identity theft through weaponized SARs—impersonating the victim and sending illegitimate ones in their name—resembling a real-world attack scenario as closely as possible. This allowed us to study the handling of SARs by organizations and to investigate the potential damage an attacker could cause. To understand if and how a victim's data privacy characteristics affect the effectiveness of such attacks and hence derive potential mitigation mechanisms for data subjects, three structurally distinct individuals volunteered as victims. We tasked a team of six attackers with gathering as much personal data on these victims as possible within the time frame and operational constraints of the case study. This attack was simulated in four stages, taking place from November 2020 until February 2021. In the first stage, which lasted a week, the attackers gathered initial data on their victims from openly available sources (e.g., social media). Afterward, they went through three iterations—roughly 33 days each—of submitting a total of 718 spoofed SARs and evaluating the responses.

**3.1.1. Realism of the Attack Scenario.** Aiming to understand if and what damages a real-world attacker could potentially cause through SAR identity theft, we tailored the design of our case study to replicate the capabilities of a realistic, yet severely constrained and weak, attacker to establish a lower boundary for this threat model. The attackers and their victims were strangers prior to our case study, such that the simulated attackers had similar knowledge constraints to real ones. They were allowed to freely select targeted organizations in analogy to a real-world attack, with restrictions only on the healthcare sector due to ethical concerns, as presented later. Because of legal and ethical constraints, the attackers were not informed about letters mailed to the victims and had no ability to intercept them, which we deemed likely to be the case in a real identity theft scenario as well. Similarly, they had no access to the victims' real email accounts or phone numbers at any time. Another operational constraint was that their frequency of interaction with organizations was limited by the iterative design of our SAR process, which will be presented later. As this limitation would not exist in a real-world setting, a real adversary would be able to react to organizations' responses more quickly and more often, likely improving their success chances.

Further, also because of legal considerations, the attackers were not allowed to falsify documents or scans thereof, and could not impersonate their victims in phone calls (from a fake number) for identity verification. Unlike our simulated attackers, a real-world adversary already in the process of committing identity theft might not hesitate to engage in such criminal activities. In their study, Di Martino et al. (2019) found that 8 of the 15 organizations that fell for malicious SARs in total did so because the adversary provided an altered identity card. We therefore believe that a significant percentage of targeted organizations could be fooled by forging a (redacted) scan of the data subject's passport or identity card. However, given our focus on simulating a highly constrained attacker, we decided to leave evaluating SAR identity theft with such capabilities as subject to future research.

Outside of the strictly regulated interactions between victims and attackers according to our SAR process presented later, there was no communication between the two parties and the victims took a passive role providing neither assistance nor additional information to the attackers. To prevent biases through local hearsay, for example, word of the study getting out to a DPO prior to them receiving a fraudulent SAR, all involved parties agreed to keep our study secret until after it had concluded.

**3.1.2. Legal and Ethical Considerations.** A vital pillar for our study design was protecting the victims, the attackers, the organizations, and the individuals therein who handled our SARs. Guided by legal restrictions and established standards for ethical research, we derived a set of operational constraints for our attack simulation. First, we required that the simulated victims (data subjects) were kept informed about the state of the attack, retained control over their personal data, and were able to withdraw their consent at all times. The simulated attackers (data requesters) received a legally binding guarantee that they could not be held liable for their attack as long as they followed rules mutually agreed upon with their victims a priori. Given that they, by design, were unfamiliar with their victims prior to the simulated attack, we asked both parties to sign a contract outlining the case study protocol and establishing rules. Additionally, as doing so may violate German law, the attackers were prohibited from forging any documents, impersonating their victim in phone calls, or attempting to intercept postal mail for the purpose of identity verification.

Second, we ensured that our study causes no harm to the targeted organizations (data controllers). To prevent them from becoming liable to legal penalties according to the GDPR, we designed our SAR submission procedure in such a way that organizations technically did not transmit any data to an unauthorized third party, even if they would have done so in a real attack due to insufficiently verifying the data subject's identity. We also did not report any such incidents to governmental data protection agencies. To prevent reputational damage to affected organizations, we do not disclose their names here nor give descriptions detailed enough to deduce their identity. Additionally, we took care to not disrupt any organization nor waste an

overproportional amount of organizational resources. For example, the total number of SARs sent to a small sports club would have been restricted to one across all attack victims, whereas a large corporation could have received one request for each victim. SARs to healthcare professionals (e.g., local doctor's offices) were limited to a few instances to avoid overburdening their resources already strained by the COVID-19 pandemic.

Third, we took care to protect the individuals that improperly handled our attackers' SARs from repercussions. From our outside perspective, we were unable to judge if flawed organizational policies or individual errors were at fault for a successful attack in some instances. Hence, we decided on a case-by-case basis to responsibly disclose vulnerabilities only if we were able to rule out individual error. When doing so, we emphasized constructive advice on improving policies rather than putting blame on the individuals that execute them. For organizations with an external DPO, we chose to address our disclosure to the external entity rather than the appointing organization. This way, they could improve their SAR processes without being at risk for having their appointment revoked.

**3.1.3. Victim Personas and Privacy Characteristics.** Our victims were selected based on structural differences in their privacy characteristics. We restrict our use of identifying information in this paper to protect their anonymity, for example, by using gender-neutral pronouns.

***3.1.3.1. VictimA.*** As a university professor, this person has interacted with a large number of organizations throughout their professional career and private life. Although aware of the resulting privacy implications, VictimA tends to disclose their real name, date of birth, and more when signing up on websites. They use their work email and phone number for many interactions with organizations, do not follow the principle of data minimization, and have little regard for recommended security practices such as using a password manager and a unique password for each account. A variety of key identifiers (e.g., date and place of birth) and a detailed CV are publicly accessible. Their research interests and parts of their professional network can be identified based on publications. VictimA has accounts on multiple social media platforms; however, most information shared on there is not publicly visible. In our case study, they represent a semipublic figure with limited privacy awareness.

***3.1.3.2. VictimB.*** Similar to VictimA, this person uses nonpseudonymized data when interacting with organizations and does not practice data minimization. Hence, many organizations store VictimB's personal data. As a research associate, however, they are less publicly exposed and have fewer key identifiers in their public CV. Despite being mostly private, their social media presences disclose some data such as age (implying year of birth) and a few hobbies. For our case study, VictimB represents an average user with some privacy awareness, for example, setting most social media private.

***3.1.3.3. VictimC.*** Being proficient in cybersecurity and working as a research associate, this person is highly aware of their digital footprint and privacy. As such, they try to minimize interactions with organizations that require disclosing personal data. VictimC uses pseudonymized data for interactions with organizations whenever possible and maintains multiple personal email addresses for creating accounts on websites, for example, one email when it is necessary to give their real name, and a different one when using pseudonymized data. Further, they use random passwords and a password manager and keep track of reported data breaches to react accordingly. When no longer interested in interacting with an organization, VictimC makes use of the right to erasure (Article 17 GDPR) to force that organization to delete their personal data. Their use of social media is limited to professional networks, where they take care to minimize disclosure of personal data. However, some details such as a profile picture and their higher education are publicly visible on social media and the website of their employer. Within our case study, VictimC represents a highly privacy aware person.

**3.1.4. Initial Knowledge and Open-Source Intelligence.** The attackers were provided with only their victims' names and workplaces, reflecting a bare minimum set of identifying information a real-world adversary would possess. During the first week of the attack and prior to submitting any SARs, they expanded their knowledge through open-source intelligence (OSINT), that is, "intelligence that is produced from publicly available information" (U.S. House 109th Congress 2006, Section 931). To this end, they scoured the victims' social media presences,[2] their employers' websites, newspaper archives, and more for information. Data validity and integrity were established by conservatively filtering out ambiguous findings based on a combination of reasonable assumptions (e.g., that a victim lives in a city close to their workplace) and cross-referencing data from multiple sources. We expect real-world adversaries to leverage OSINT for initial reconnaissance in a similar manner, given that it constitutes a promising yet easily and covertly accessible source of information.

**3.1.5. SAR Process.** Our case study design was guided by striking a balance in the tradeoff between our goal of simulating a realistic attack by a malicious

adversary on the one hand and fulfilling the previously outlined legal and ethical requirements on the other hand. We developed an iterative process for coordinating the SARs between attackers and victims, as depicted in Figure 1.

Before starting the first iteration of our simulated attack, each victim created a "fake" email account in their own name with a popular free email provider.[3] Having the victims create an email account themselves helped avoid potential legal implications for organizations, because any reply from organizations, including those leaking data without proper identity verification, would be received by an account that, on paper, belongs to the true data subject. In a real identity theft scenario, the attackers would have created this email account themselves in order to impersonate their victim. A credible email address not only boosts the legitimacy of the fraudulent SARs but also provides the attackers with a plausible pretext for lacking access to the victim's real email, which the targeted organization might store and use for proof of ownership. For example, if an organization sends a password to the victim's real email in response to the SAR, the attackers could ask for it to be sent to the victim's "new" (fake) email instead, pretending to have lost access to the real account.

At the start of each iteration, the victims granted the attackers access to their fake email, but never to their real email accounts. This does not impact realism, as the attackers would have created this email account themselves in a real scenario and thus have full access to it. The attackers then had 24 hours to draft, but not send, SAR emails to organizations of their choosing. As presented in more detail later, they strategically selected and approached organizations in different ways to best fit their current knowledge about each victim.[4] SAR emails were authored in German or English and addressed to the data protection inquiry email or, as a fallback solution, generic customer support.

Subsequently, the victims revoked the attackers' access to their fake email by changing the account password. Then, they sent out the attackers' drafts without being allowed to alter them. This proxied design was crucial to fulfilling our legal and ethical requirements because it ensured that all SARs were technically submitted by the true data subject (the victim) rather than an impersonator (the attackers), eliminating potential legal issues for all involved parties. Further, it also guaranteed that the victims reaffirm their consent to each individual SAR being sent. The attackers had no access to the fake email account for the entire legally allowed response time of one month. Once it had elapsed, the victims audited all received emails and could prevent any personal data from reaching the attackers by deleting individual responses,[5] in line with our ethical guidelines. Finally, the attackers were again granted access to the fake email account, analyzed all replies, and aggregated newly gained knowledge to be weaponized. They also decided if and how to follow up on authentication requests by DPOs and prepared a list of new organizations to contact in the next iteration.

**3.1.6. Use of Social Engineering.** Social engineering refers to exploiting the human element for cyberattacks (Mitnick and Simon 2003). It is commonly studied in the context of phishing (Williams et al. 2018), which

**Figure 1.** SAR Process of the Case Study

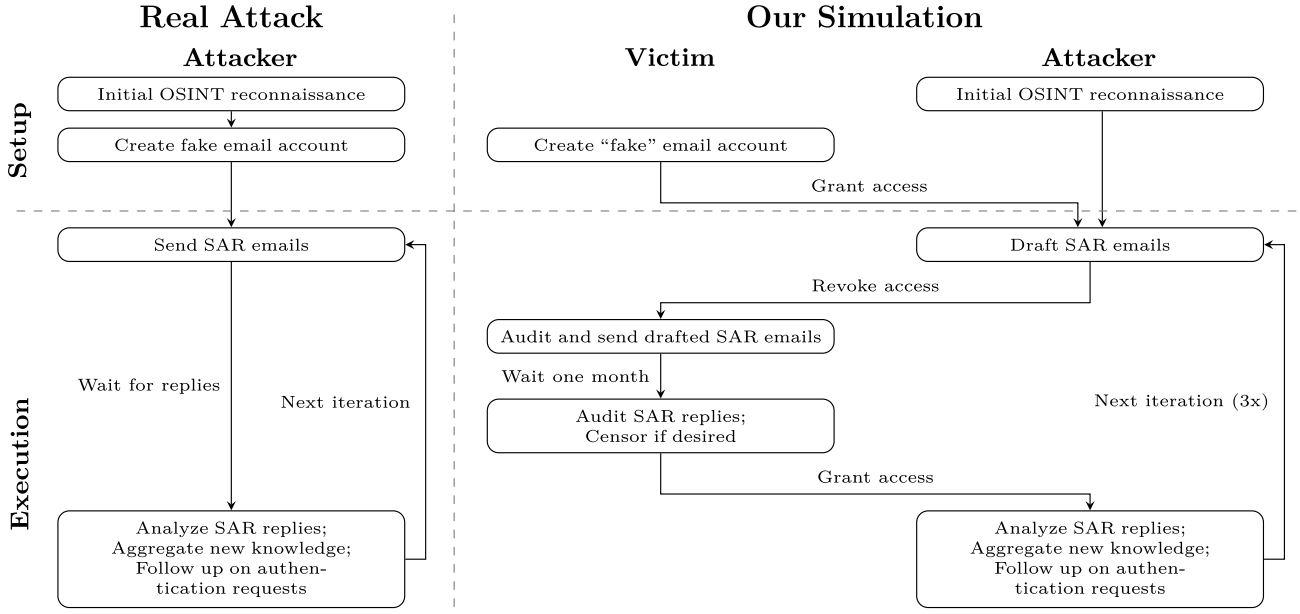

exhibits many similarities to our fraudulent SARs. Recognizing this, our attackers read up on the basics to prepare for interacting with DPOs, given that they had no experience in social engineering. They formulated SARs using techniques shown to boost phishing susceptibility such as credible impersonation (Algarni et al. 2017), contextualization (Goel et al. 2017), and influence techniques (Cialdini 2009, Wright et al. 2014). For example, they emphasized the mandatory response within one month to instill a sense of authority and urgency, hoping that the resulting susceptibility to phishing also applies to our scenario (Williams et al. 2018).

### 3.2. Data Collection

Our data sample comprehensively depicts the entire chronology of the simulated attacks, comprising 718 SARs submitted by the attackers in total. Throughout the study, we systematically recorded the attackers' strategies, including their reasoning for choosing them, as well as how organizations reacted to submitted SARs. In three focus group interviews (one prior to each iteration), we acquired reports about the attackers' initial OSINT findings and their discussions regarding attack strategy, target selection, and development of highly targeted SARs sent to specific organizations. Further, we recorded all communication with organizations (e.g., received emails and letters) with detailed metadata such as request and response timestamps, recipient (e.g., a response sent to the victim's true email rather than the fake one) and whether an external DPO handled the request.

### 3.3. Data Analysis

We analyzed our data sample in a three-stage process. First, we focused on the individual victims as isolated cases in a within-case analysis. We investigated what data were leaked by organizations and how, as well as especially effective attack strategies and key challenges for the attackers. Second, a comparison of the three cases in a cross-case analysis yielded common patterns in organizations' SAR processing. We also examined the effect of our victims' distinct privacy characteristics on the attack by contrasting scope and sensitivity of leaked personal data acquired by the attackers. Third, we conducted an in-depth root cause analysis to reveal the enablers of such attacks.

### 3.4. Supplemental Expert Interviews with DPOs

We enriched our case study with 21 in-depth expert interviews with DPOs of organizations targeted by our attackers and beyond.[6] They were conducted after the simulated attack had concluded and comprised a total interview duration of over 15 hours, resulting in more than 250 pages of transcripts for qualitative analysis. We selected a broad sample of interview partners to cover different industries and organizations of different sizes. In contrast to our case study that approaches SAR identity theft from the attackers' perspective, these complementary interviews provided valuable insights from the receiving end of such attacks, allowing us to assume a more complete perspective in our analyses.

## 4. Results

According to our three-stage data analysis, we present our findings in three parts. First, an isolated analysis of how the three victims were attacked reveals successful data leaks and attack strategies. We then compare the three cases to identify common behavior by organizations and to investigate the impact of our victims' privacy characteristics on data leaked to the attackers. Finally, we enrich these insights with expert interviews and an inspection of the GDPR to unveil and systematize the underlying root causes enabling our attack.

### 4.1. Within-Case Analysis

Starting with varying levels of OSINT knowledge about their victims, our attackers tailored different strategies for the first iteration and adapted them after each iteration to account for newly acquired and now weaponizable information. We will present their approaches, successful attacks resulting in important data leaks, and other notable incidents in an anecdotal manner on a case-by-case basis.[7]

**4.1.1. VictimA.** Prior to the first iteration of our case study, the attackers had already collected some key data commonly used for weak forms of knowledge-based authentication about VictimA through OSINT, such as date and place of birth, as well as a likely home address. Through educated guesses based on the aggregated data sources, in particular the victim's public CV, they were able to reconstruct a coarse timeline of places that VictimA had lived in since birth.

***4.1.1.1. First Iteration.*** Utilizing this a priori knowledge, the attackers decided to pursue a dual-track strategy, submitting SARs to a total of 118 organizations in the first iteration.

On the one hand, they submitted highly targeted SARs to select organizations that were known to store data about VictimA (e.g., the former school) or were likely to do so (e.g., sports clubs matching hobbies in cities they have likely lived in). When approaching the victim's former school, for example, our attackers weaponized data from VictimA's public CV, summarizing their professional career and referencing the year of graduation to convey legitimacy by disclosing supposedly nonpublic knowledge. Additionally, they plausibly contextualized the SAR with an academic research project on the GDPR, which—ironically—was not far from the truth. This proved successful, yielding

our attackers scanned documents containing sensitive data from VictimA's childhood, such as their parents' names and address, their religion, and the name of their primary school. Any of these could have been requested for proof of identity, which would have rendered the attack unsuccessful.

On the other hand, the attackers picked industries that almost everyone interacts with but that also have relatively few key players, such as airlines or insurances operating in Germany, and sent a generic SAR to all organizations within. Because the GDPR mandates a SAR response even if the requester is unknown to the organization, the attackers hoped to find out which organizations store data on VictimA by principle of exclusion—if all organizations except one reply that they store no data, it is highly likely that this one organization does. Through such "organization enumeration attacks," the attackers acquired knowledge that helped them narrow down the scope of subsequent efforts, without necessarily being able to fulfill identity verification requests by the organization yet.

With this dual-track strategy, the attackers caused four significant data leaks already during the first iteration. The knowledge acquired this way confirmed previous assumptions on key information usable for proof of identity in subsequent iterations, such as VictimA's current home address. They also gathered new key identifiers such as a private email address used by the victim and learned a variety of organization-specific information (e.g., customer numbers). Although potentially useful to a real-world adversary engaging in further attacks beyond submitting SARs, organization-specific data were of little use to our attackers as they could not be weaponized against other organizations.

#### 4.1.1.2. Second Iteration.

Having exhausted all organizations they knew to store data on VictimA, the attackers focused on educated guesses to target 126 new organizations. Among others, they succeeded in exfiltrating personal data from a multinational conglomerate in the furniture industry, a multinational car rental company, a comparison shopping website, and multiple scientific publishers. In addition to organization-specific data such as customer numbers or purchase histories, they acquired more key identifiers including VictimA's private mobile phone number, former home address, national identity card and driver's license numbers, and a personal bank account number.

Further, they gained some insights into the victim's insurances and a building loan, presumably for their current home. This was caused by flaws in how the organization secured the data exchange with the requester, allowing the attackers to access names, but not contents, of files sent in the SAR reply. These file names, however, contained relevant metadata such as dates and descriptions of individual insurance policies or loan financing plans, for example, "20170821_lifeinsurance_contract.pdf."

#### 4.1.1.3. Third Iteration.

The attackers decided to leverage newly gained knowledge to target a total of 54 more specialized organizations, for example, government agencies from the victim's hometown, a winery, and a lottery. Despite some of them storing data on VictimA, no further leaks were achieved.

### 4.1.2. VictimB.

In the case of VictimB, the attackers were unable to find certain key data commonly used for proof of identity, such as date and place of birth or home address, through OSINT.

#### 4.1.2.1. First Iteration.

Afraid of being unable to fulfill a potential request for proof of identity due to their limited knowledge, the attackers decided to not immediately contact the victim's former school (as known from a public CV) in the first iteration. Instead, because their knowledge was mostly related to the victim's professional career, they decided to focus on submitting SARs to 26 organizations that they deemed likely to have had interactions with a person from VictimB's line of work. For example, they contacted select European railway companies, as well as hotel chains popular for work-related travel throughout Germany and Europe. Out of these organizations, one hotel chain provided them with the victim's home address and entire booking history.

#### 4.1.2.2. Second Iteration.

The attackers leveraged their newly acquired data for a total of 126 SARs and to follow up on requests for further authentication from the first iteration. Given their limited success thus far, the attackers adapted their targeting strategy to a combination of educated guesses and broader organization enumeration attacks, similar to VictimA. This proved successful, resulting in six major leaks yielding key information including VictimB's date of birth, former and current private mobile phone number, former home address, and personal bank account number—all obtained from multiple organizations, thus strengthening data validity through cross-references.

One of the most glaring security flaws was an automobile association with millions of members leaking VictimB's date of birth, home address, current private phone number, bank account number, and more. Their DPO sent an encrypted ZIP file attached to an email stating that VictimB's date of birth was used as password. This implies that they considered knowledge of the data subject's date of birth, plus supplied name and work email matching their records, to suffice for accessing data such as bank account numbers. Although questionable as to whether this is GDPR compliant, their approach also exhibits a trivial vulnerability

rendering the requirement to know the data subject's date of birth void: Using any date of birth to encrypt a file archive transmitted to the attackers (implying losing the ability to rate limit attempts at cracking the password) provides no security. The attackers simply tried out all valid dates of birth[8] within milliseconds using a computer program.

Notably, a former insurer of VictimB initially replied via postal mail to the home address they stored. Given that intercepting letters was beyond our attackers' capabilities, we would consider this sufficiently secure. However, the letter was returned as undeliverable because the stored address was not up to date. In an act of helpfulness, the employee in charge contacted the attackers via the fake email, providing all sensitive personal data electronically without encryption or proof of identity. This incident highlights that even if a reasonably secure SAR response workflow exists in an organization for everyday cases, the security concept is at risk for falling apart in extraordinary situations.

***4.1.2.3. Third Iteration.*** The attackers decided to continue their strategy and contacted 54 new organizations, including some previously identified as candidates for systematic flaws like the multinational car rental company that leaked VictimA's data. In analogy to that interaction, the company also transmitted VictimB's private phone numbers, date and place of birth, national identity card and driver's license numbers, and bank account number without any proof of identity. Further—now confident in being able to authenticate themselves—the attackers targeted the victim's former school, whose DPO disclosed data on the victim's family, religion, as well as grades and courses throughout their school years. Having found a photograph of VictimB wearing glasses via OSINT, the attackers also targeted optical store chains. Despite storing the victim's personal email and phone number, both usable for reasonably secure authentication, one such chain with a significant share of the German eyewear market disclosed the victim's visual acuity measurements and purchase history.

**4.1.3. VictimC.** The attackers collected the least amount of data on VictimC via OSINT. More importantly, they were unable to find any key identifiers, such as a personal email address, beyond work-related ones published on the website of the victim's employer. Further, they could not find any hints regarding VictimC's hobbies or affiliations with organizations other than a former school.

***4.1.3.1. First Iteration.*** Given the limited amount of information the attackers knew about VictimC, they chose to pursue only organization enumeration attacks in the first iteration, submitting SARs to 99 organizations (e.g., airlines or supermarket chains) in total. Even if they were likely unable to follow up on requests for proof of identity, they hoped to gain some insights into what organizations stored data about VictimC. However, their efforts were of limited success because none of the targeted organizations leaked personal data, and only three organizations revealed that they stored data on the victim in their authentication requests. Unlike the schools contacted impersonating VictimA and VictimB, a scanned national identity card was requested for verification by the employee handling the SAR submitted to VictimC's secondary school. A noteworthy observation during this iteration is that several of the contacted organizations denied storing any data on VictimC, despite doing so. We suspect that these organizations look up data for SARs by email address rather than name, as the victim had used only their personal email address—which was unknown to the attackers at this point and hence not included in the SAR—for previous interactions with them.

***4.1.3.2. Second Iteration.*** Having gained barely any new knowledge on VictimC, the attackers targeted 94 organizations in the second iteration without changing their strategy. A local energy provider operating in the vicinity of VictimC's workplace leaked a broad spectrum of personal data, such as home address, personal email address, bank account number, customer number, and electricity bills. None of the identifiers supplied by the attackers, such as work email or work phone number, were known to the energy provider. This implies that the only data their DPO could have used for proof of identity was the victim's name, which is certainly in breach of the GDPR. Concerningly, this organization had secure options for identity verification at their disposal, such as requesting VictimC's customer number plus data from their latest electricity bill for reasonably secure knowledge-based authentication. They could also have leveraged the victim's personal email address or home address for proof of ownership. This is particularly noteworthy because the SAR was handled by an externally appointed DPO specialized in providing GDPR compliant services. The attackers hence identified this organization as a likely candidate for a systematically flawed SAR workflow to be exploited for the other victims.

***4.1.3.3. Third Iteration.*** The attackers followed up on SARs from previous iterations, now able to provide identifiers such as VictimC's personal email. Additionally, they submitted SARs to 21 new organizations, most of which they previously identified as systematically vulnerable. The attackers were unable to obtain more personal data from these organizations because none of them stored any data on VictimC. However, a

newly contacted multinational video game and consumer electronics retailer disclosed VictimC's date of birth. This organization could have leveraged VictimC's personal email for proof of ownership but processed the SAR without any further authentication. Even worse, they added the attackers' fake email as a legitimate alternative one to the victim's user account.

## 4.2. Cross-Case Analysis

**4.2.1. Privacy Characteristics, Leaked Data, and Potential Damage.** The three victims of our case study are structurally distinct in their public exposure and privacy preferences. For example, whereas VictimA is a semipublic figure with little regard for their digital footprint, VictimC tries to minimize publicly available personal data and uses pseudonyms whenever possible. This reduced the effectiveness of our simulated SAR identity theft in two ways: First, the attackers were unable to find key identifiers such as VictimC's date of birth via OSINT while preparing their attack. Second, the number of organizations that could have leaked data on this victim's real identity was comparatively low due to most of them storing only pseudonymized data.

Despite VictimC's efforts, they were unable to protect themselves against SAR identity theft, as shown in Table 2. Rather than fully mitigating this attack, we observed that an increased level of privacy awareness resulted in a time shift (in iterations) of data known to the attackers. Here, multiple data leaks across three iterations were required for their knowledge on VictimC to reach a level slightly above what they knew about VictimA from OSINT. Arguably, this convergence of known data would have continued throughout further iterations, ultimately nullifying VictimC's advantage. This is because simply contacting more organizations increases the likelihood of data leaks and enables organization enumeration attacks, without any diminishing returns or increased risk for the attackers other than more likely alarming their victim of the ongoing attack.[9]

**Table 2.** Iterations Needed to Acquire Victims' Key Identifiers

| Data | VictimA | VictimB | VictimC |
|---|---|---|---|
| First and last name | Provided | Provided | Provided |
| Workplace | Provided | Provided | Provided |
| Workplace email address | OSINT | OSINT | OSINT |
| Workplace phone number | OSINT | OSINT | OSINT |
| Home address | OSINT | 1 | 2 |
| Date of birth | OSINT | 2[a] | 3 |
| Place of birth | OSINT | 3 | — |
| Personal email address | 1 | 2 | 2 |
| Personal phone number | 2 | 2 | — |
| Bank account number | 2 | 2 | 2 |
| National identity card number | 2 | 3 | — |
| Driver's license number | 2 | 3 | — |

[a]Year of birth known from OSINT prior to the first iteration.

One of the most prevalent uses of stolen identities by criminals is stealing money through credit card fraud or similar (Willox et al. 2004). Relying solely on the personal data exfiltrated by our attackers, such goals could have been accomplished for all three victims. For instance, using the victims' bank account data, a criminal actor could have purchased goods in online shops via direct debit. A more sophisticated attacker could have even attempted to forge a national identity card using VictimA's data to, for example, register a credit card in their name or cross state borders under a false identity for the purpose of drug trafficking. On the bright side, attackers without knowledge about which organizations store data on their victim need to submit SARs to a large number of organizations. Although the process of sending generic emails can be automated easily, the replies our attackers received were mostly unstructured, highly individualized, and sometimes implied clues for educated guesses even when not directly leaking personal data. Based on our observations, we consider it difficult to fully automate evaluating these replies as of today. This implies limited scalability, meaning that even though SAR identity theft is feasible for targeted attacks on select individuals, it does not yet pose a highly automated large-scale threat such as phishing emails. However, attackers may soon be able to leverage advances in artificial intelligence (AI) to automate parts of SAR identity theft that previously required human input, such as researching organizations to contact, authoring SAR emails,[10] or extracting and aggregating information from unstructured replies.

**4.2.2. Key Identifiers Requested by Organizations.** Organizations commonly requested certain types of data either for knowledge-based proof of identity or for identifying the data subject in their database. These key identifiers can be split into two groups: organization-specific (e.g., customer number) and generic (e.g., date of birth) personal information.

In case of a data leak, organization-specific data typically cannot be exploited for SARs submitted to other organizations. As doing so is key to our identity theft scenario, such data were of little value to the attackers. Additionally, DPOs insisting on organization-specific data for identity verification—rather than being content with generic data—constituted a roadblock given the capabilities of our simulated attackers. We therefore consider this approach an improvement over requesting generic personal information. Nevertheless, a more sophisticated attacker in a real-world scenario might still be able to obtain organization-specific data through additional, potentially illegal, means.

Leaked generic personal information, however, was highly useful to our attackers as they could weaponize it against further organizations to exfiltrate even more data. Consequently, we observed that even small leaks of such data, seemingly unimportant when seen in isolation, played a significant role when accumulated in a large-scale attack targeting hundreds of organizations. Despite its widespread use, knowledge-based authentication using generic data (e.g., date of birth) can therefore not be considered sufficiently secure for verifying a SAR data subject's identity.

Further, we identified certain "critical mass" thresholds of exfiltrated generic data. Once crossed, we observed a significant increase in the number of organizations able and willing to process the attackers' requests. For example, knowing a broad range of personal contact information (e.g., home address, email address, and phone number) enabled most organizations to locate the victim's records in their data storage. Based on Di Martino et al. (2019), we believe that an even more significant critical mass effect can be observed once the attackers are capable of convincingly forging a national identity card scan. Although our attackers managed to acquire all data necessary for such an endeavor, they did not attempt to falsify any documents for legal and ethical reasons.

## 4.3. Root Cause Analysis

To identify the root causes for our observations, we investigated salient patterns that emerged when analyzing how organizations reacted to the 718 submitted SARs. We supplemented this data with insights from 21 in-depth expert interviews with DPOs of organizations targeted in our case study and beyond, as well as with a close inspection of the GDPR and the role it played in our attack scenario. The results of our analysis are summarized in Table 3 and will be presented in the following sections.

**4.3.1. Organization-Induced Leaks.** We first zoomed in on the level of isolated SARs and disregarded interdependencies with SARs sent to other organizations. Throughout our study, most data were leaked here either due to nonsystematic errors made by the individual employees handling our attackers' requests or due to systematic flaws in the processes implemented by organizations for handling SARs. Hence, we named such instances "organization-induced leaks." Negligent disclosure of personal data to an unauthorized third party does not comply with the GDPR and can lead to administrative action and/or fines.

***4.3.1.1. Nonsystematic Flaws.*** We observed a broad range of errors made by DPOs processing our attackers' fraudulent SARs, which resulted in significant data leaks. On the one hand, some of them were simple human errors, for example, a large automobile manufacturer mailing VictimA detailed data of an unrelated person because documents were accidentally placed in the wrong envelope. On the other hand, many of the DPOs' errors were deliberately caused by our attackers' use of social engineering techniques. This comes as a surprise because DPOs must have "expert knowledge of data protection law and practices" (Article 37(5) GDPR), implying a certain level of resilience against precisely such social engineering attacks. We explain our attackers' success with four attributes unique to the GDPR that ironically boost the efficacy of social engineering in our scenario.

First, credibly contextualizing the interaction with the victim is typically challenging but highly convincing in other social engineering attacks such as phishing (Goel et al. 2017). In our attack, this context is readily provided by the GDPR and the pretense of the impersonated data subject exercising their right of access. The principal of VictimA's school, for example, did not question the legitimacy of the SAR due to the credible context of a former student, now professor, conducting research. Our expert interviews corroborate this, as all DPOs stated that they associate incoming SARs with a credible interaction context by default. For instance, the DPO of an airline requests additional proof of identity

**Table 3.** Types of Data Leaks Observed in Our Study

| | Organization-induced leaks | Regulation-induced leaks |
|---|---|---|
| Description | Comprehensive data leak due to GDPR noncompliant behavior | Smaller data leak despite GDPR compliant behavior |
| Examples | Insecure identity verification policies; DPO manipulated via social engineering | Organization enumeration enabled by mandatory null responses |
| Severity | Main source of data leaks | Supplemental source of data leaks |
| Scope | Isolated (individual SAR) | Sequential (aggregated SARs) |
| Responsibility | Organization (systematic flaws); individual DPO (nonsystematic flaws) | Regulators/GDPR |
| Characterization | Systematic or nonsystematic GDPR implementation failure | Systemic issue in the GDPR right of access |
| Liability | Administrative action/fines | None |
| Root causes | Insecure processes; insufficient training; lack of social engineering resilience | Shift of scope from isolated to sequential attacks is not covered by the GDPR |

only if "there are signs that something is not quite right" (I17).

Second, the GDPR mandates that "[t]he controller shall facilitate the exercise of data subject rights" (Article 12(2) GDPR). Multiple interviewees reported difficulties reconciling this with rigid identity verification, as perhaps best summarized by the DPO of a large tourism company: "Identity verification in [the GDPR] is a double-edged sword, because you are not allowed to raise the hurdles too far" (I21). It also nudges DPOs toward a helpful and customer-oriented mindset, which our attackers exploited with social engineering (Wright et al. 2014). For example, they persuaded DPOs into accepting weaker proof of identity by arguing that their initial request was excessive.

Third, most interviewees admitted to uncertainty regarding GDPR compliance due to generic formulations in the law. Per the DPO of a large pharmaceutical company, "no organization can truly say that they are absolutely one hundred percent compliant" (I4). Our attackers exploited this using distraction and authority cues (Cialdini 2009). For example, the law vaguely stipulates that "[t]he controller should use all reasonable measures to verify the identity of a data subject who requests access" (Recital 64 GDPR). Accordingly, they argued about the definition of "reasonable measures" and thereby negotiated what identifiers they had to provide.

Fourth, organizations are obliged to process SARs "without undue delay and in any event within one month of receipt" (Article 12(3) GDPR). Throughout our interviews, this was frequently mentioned as a challenge by DPOs who face complicated data gathering due to unstructured, decentralized, and sometimes even analog data storage. Our attackers exploited this via urgency cues (Cialdini 2009) in combination with deliberately complex SARs, hoping that DPOs were forced to reallocate time away from a thorough identity verification to meet the deadline.

#### 4.3.1.2. Systematic Flaws.

We explain the systematic flaws identified throughout our study with organizations being inadequately prepared for handling SARs that are not standard operating procedure, such as ours submitted with malicious intent. When told that SARs are exploitable for identity theft, for example, the DPO of a national hostel chain felt "embarrassed about having to admit that [they] had not even considered that" (I20). Another indicator for this is the aforementioned difficulty fulfilling our deliberately complex SARs within the mandatory deadline. Further, all DPOs we interviewed had completed some form of cybersecurity and data protection training to comply with the GDPR mandating expert knowledge of data protection practices. Most of them had also consulted external (e.g., law firms) and/or internal (e.g., Chief Information Security Officer or IT department) experts for SAR process development. Despite this, we identified trivial vulnerabilities resembling those observed in our simulated attacks when we discussed these processes with the DPOs. Based on our observations, we cluster such systematic vulnerabilities into two groups, neither of them compliant with the GDPR.

On the one hand, many organizations implemented an insecure identity verification process, or had no such process at all. For example, some disclosed sensitive personal data despite confirming the legitimacy of the SAR only through relatively easy to acquire key identifiers (e.g., date of birth). As another example, a local energy provider leaked a broad set of personal data on both VictimB and VictimC without any authentication. Almost all interviewees reported that a censored scan of the data subject's national identity card would suffice as proof of identity. Quoting the DPO of a large company in the German healthcare sector, "[an attacker] might as well falsify a national identity card. But this criminal energy, that usually does not happen, because we would become suspicious" (I10). Di Martino et al. (2019), however, demonstrated such scans to be convincingly falsified with little effort,[11] so simple low-quality scans should not be relied on for identity verification.

On the other hand, certain organizations failed to implement a secure data exchange with the requester in response to SARs. For example, some organizations transmitted personal data as plain text in an unencrypted email, implementing no security measures protecting against eavesdropping by a malicious third party. Similar issues came to light in our interviews, such as a large hospital network considering the transmission of sensitive medical data via unencrypted email to be sufficiently secure (I12). Other organizations attempted to securely transmit personal data to the SAR subject but failed to do so. A common approach was to transmit data contained in an encrypted ZIP archive via email and disclose the corresponding password separately as means of authentication. For example, one organization indicated that the archive was encrypted using the data subject's date of birth, which provided no security because all valid passwords were easily tried by a computer program. Even with a sufficiently complex and securely transmitted password, this approach implies losing control over cracking attempts and therefore enables "harvest now, decrypt later" strategies.

### 4.3.2. Regulation-Induced Leaks.

Zooming back out to investigate our identity theft scenario as a whole and observing effects that emerge from the chain of multiple cascading attacks on many organizations, we enter territories largely uncharted in literature. Contrary to organization-induced leaks where the organization was

clearly at fault, we observed "regulation-induced leaks" where even organizations with—if seen in isolation—reasonably secure and GDPR compliant authentication leaked small amounts of information to our attackers. We identified two distinct types of such leaks:

First, consider an exemplary organization that stores four identifiers (name, date of birth, place of birth, and bank account number). It would be reasonably secure and GDPR compliant for that organization to request three of them (name, date of birth, and bank account number) for authentication. If an attacker possesses these, more data (place of birth) would be disclosed to them than what they were required to provide. Although this may appear unproblematic considering only this single organization in isolation, our study demonstrates that even such small bits of data may be highly valuable to an attacker when accumulated for large-scale sequential attacks on many organizations.

Second, organizations are obliged to process SARs even if they do not store data regarding the requester. Compliant with the GDPR, many organizations in our sample briefly replied that no data were stored about the impersonated victim. This by itself constitutes novel information, which in some industries revealed to our attackers organizations that had likely interacted with the victim. Ultimately, such null responses enabled the organization enumeration attack strategy that we presented earlier and allowed our attackers to narrow down their efforts to specific organizations.

Given that such leaks occurred despite GDPR compliance, our observations hint at a systemic issue rooted in the GDPR itself. It sufficiently covers SAR security on the level of isolated organizations: Organization-induced leaks are caused exclusively by flawed implementations and not by flaws in the GDPR. However, the cascading nature of our specific scenario allowed attacking the integrated system of all organizations as a whole rather than just multiple organizations isolated from each other. We argue that this shift from isolated to sequential attack scope—broader than what is regulated by the GDPR—resulted in the observed emergence of regulation-induced leaks as a novel vulnerability.

## 5. Discussion

Our study sought to shed new light on how data privacy and data security interrelate in shaping personal data protection. To do so, we systematized prior, often implicit conceptualizations of this interplay and advanced what we call the interference view as a novel perspective highlighting tensions between the two concepts (Figure 2). We illustrated this view by exploring an identity theft scenario, as part of which attackers weaponized individuals' right of access to their personal data and the associated instrument of subject access requests (Article 15 GDPR) to steal sensitive personal data including the victims' national identity card, bank account number, and consumption history from organizations. Paradoxically, a shield for data privacy could be turned into a sword that threatens data security and is even effective in exfiltrating the personal data of highly privacy aware individuals. As highlighted in Figure 2, interferences between data privacy and data security are inherently multilayered in nature, in that they can manifest at the conceptual, regulatory, or organizational level. Importantly, interferences at one level (e.g., mismatching scopes of legal mandates at the regulatory level) can cascade up to the next level (e.g., organizations struggling to fulfill diametrical demands at the organizational level), accumulating into a far-reaching threat to data protection (e.g., enabling cybercriminals to steal identities by exploiting regulation-induced leaks combined with organization-induced leaks). Below, we discuss the implications of these findings for research and practice, highlight the limitations of our study, and suggest directions for future research.

### 5.1. Implications for Research

We see two main contributions for research. First, we advance knowledge on the interplay between data privacy and data security, the two constituents of data protection in our digital age. Despite their coevolution and substantial attention from researchers, practitioners, and policymakers (Lowry et al. 2017), surprisingly, both concepts still tend to be studied, managed, and regulated separately, leaving the precise nature of their interplay underexplored (Belanger et al. 2002, Smith et al. 2011, Biselli and Reuter 2021). It is against this backdrop that our study adds conceptual clarity by distinguishing and explicating three perspectives. These include the still dominant layered view that conceives of data security as a necessary but not sufficient precondition for data privacy (Smith et al. 2011), as well as the emerging complementarity view that sees them as mutually reinforcing (Solove and Hartzog 2022). We add an underexplored, yet conceptually and empirically meaningful third perspective that we call the interference view, where data privacy can compromise data security or vice versa. We enrich this perspective with empirical evidence by identifying interferences between the two concepts in the GDPR right of access on multiple levels: conceptual, regulatory, and organizational. As highlighted in Figure 2, their adverse effects accumulate and contribute to the unexpected severity of the SAR identity theft attacks we simulated. Our study thus illustrates how progress in data privacy might inadvertently trigger regress in data security, ultimately threatening personal data protection. These insights challenge the assumption underpinning the layered view and the complementarity

**Figure 2.** Interference View: Exemplary Tensions Between Data Privacy and Data Security

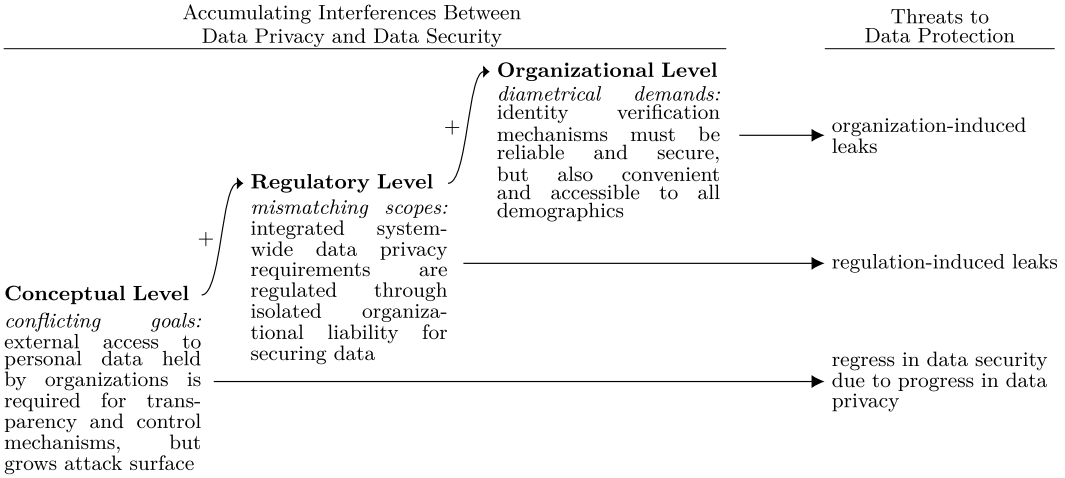


view that data privacy and data security are necessarily prerequisites for—or complements of—each other (Acquisti et al. 2016, Biselli and Reuter 2021). Instead, under certain conditions, they can have a mutually compromising relationship with one exhibiting the potential to undermine—or interfere with—the other. Importantly, we do not argue for the superiority of the interference view or indeed any of the two others. Rather, we see these views as coexisting and mutually enriching. In many settings, the interplay between data privacy and data security will exhibit facets of all three views.

Second, our theory and evidence contribute to the still small, but growing literature stream on the effectiveness and unintended consequences of data protection regulations such as the GDPR (Johnson 2022). Our study complements previous modeling work that focuses on potential unintended negative consequences of the GDPR with regards to pricing and ultimately consumer surplus (Ke and Sudhir 2023, Xu et al. 2025). We do so by showing empirically and across a broad range of markets that data privacy provisions contained in data protection regulations can, as an unintended consequence, even undermine data security with potentially far-reaching negative implications for personal data protection and consumer surplus. More specifically, we demonstrate that the right of access to personal data according to Article 15 GDPR can be weaponized by malicious actors to threaten consumers' data security and, ultimately, their personal data protection—a risk earlier studies had pointed to (Di Martino et al. 2019, Bufalieri et al. 2020). These insights and the interference view they exemplify illustrate that data protection regulations may contain provisions that paradoxically can be turned from a shield into a sword to attack what the regulation was designed to protect in the first place.

Our insights into both organization-induced and regulation-induced leaks that enabled such attacks to succeed also add to the literature on how data protection regulations are enacted and translated into organizational and individual practices. In line with previous research in this space (Waldman 2020a, b; Solove and Hartzog 2022), our study documents that the original regulatory intention—namely to protect consumers and their personal data—risks being lost in enactment and translation. Going beyond prior research, we show that these translation challenges can be exploited by malicious actors during cyberattacks to make the regulation and individual provisions contained therein not only less productive than anticipated, but even potentially counterproductive for protecting personal data. Our newly introduced conceptual framework and its interference view expose drivers for these translation challenges, as illustrated in Figure 2. For example, the GDPR poses diametrical demands on identity verification, which contributed to enabling organization-induced leaks that result from ineffective data protection policies and/or individual noncompliance (Parks et al. 2017). Additionally, we unearthed regulation-induced leaks that occur despite compliant behavior and are rooted in imperfections of the regulatory texts, which leave ample room for interpretation or even allow malicious exploitation through sophisticated system-level attack strategies (Waldman 2020a). At a deeper level, our analyses reveal that cyberattacks can—implicitly or explicitly—exploit imperfections of the regulation (e.g., right to make an unlimited number of SARs), the organization (e.g., disconnect between data privacy and data security professionals in GDPR process design and execution), and the individual handling the SAR (e.g., lack of situational threat awareness) to exfiltrate sensitive personal data. This highlights that data protection regulations must be designed to

withstand not only implementation by organizations with conflicting commercial interests (Waldman 2020b), but also misuse by malicious actors in search for new attack strategies.

### 5.2. Implications for Practice

Overall, we demonstrate the importance of managing and regulating data privacy and data security in an integrated manner to account for their interdependence. Given the multifaceted interplay between the two, treating them as separate or loosely connected concepts might leave potential complementarities unexploited and create unintended interferences on the regulatory or on the organizational level, such as the exemplary ones identified in our study and illustrated in Figure 2. For organizations, an integrated approach requires close coordination between, e.g., the IT department typically handling data security and the legal/compliance department often responsible for data privacy. Regulators crafting integrated data protection laws, in turn, could benefit from the three views on the interplay between data privacy and data security we advance. A deeper understanding of this interplay also appears relevant for legal scholars seeking to evaluate ambitious rights, such as the GDPR's right of access, and their potential consequences. More specifically, our study highlights the need for action by organizations and regulators to mitigate SAR identity theft and the damage it could cause to victims (Willox et al. 2004). Our insights into the root causes of both organization-induced and regulation-induced data leaks appear instructive in this regard.

To prevent organization-induced leaks, it is vital to raise awareness of this indirect, sequential attack strategy among DPOs, to educate them on how to recognize and address it, and to motivate them to engage in secure behavior. What appears to be needed is deeper social engineering resilience training that builds an understanding of the underlying techniques rather than focusing solely on their typical application in phishing emails (Abbasi et al. 2021). Likewise, organizations need to carefully navigate the tension between accessibility and security of their SAR identity verification mechanisms, implement resilient policies for processing SARs, reduce the room for human error, and ensure a secure data exchange with the requester. Regulators, in turn, could contribute to weakening the effectiveness of urgency and authority cues employed by attackers by showing leniency if an organization is found in breach of the GDPR but acted in good faith. On the technical side, national or international identity verification systems, such as the European Digital Identity, could be linked to organizations' databases for more robust authentication.

To contain regulation-induced leaks, measures are needed that can interrupt the chain of cascading attacks against many organizations, for instance, by constraining the use of personal data that are not organization-specific (e.g., date of birth) for SAR identity verification. If possible, proof of ownership (e.g., replying only to a stored email address) should be preferred over proof of knowledge for authentication. Governments could assist organizations with authentication and simultaneously impede organization enumeration attacks by limiting the number of SARs a data subject is allowed to submit within a given time frame. For example, a centralized clearinghouse system, accessed using government-issued credentials, could issue a limited number of one-time tokens to be submitted alongside a SAR and to be consumed by the receiving organization. However, such a system in itself might raise new privacy concerns and could become a high-value target for cybercriminals.

### 5.3. Limitations, Future Work, and Conclusion

We carefully designed our exploratory multiple case study to fulfill the quality criteria outlined by Yin (2009) and to achieve ecological validity (Lowry et al. 2017). This includes executing simulated attacks that closely mimic a real-world scenario, observing similar outcomes for three structurally distinct victims, contacting a large number of different organizations, and presenting a detailed case study protocol. Yet, like any research, our study is not without limitations that future research might want to address. External validity could be further enhanced by replicating the study in other geographical contexts or by imposing fewer constraints on the attackers' capabilities, for example, allowing them to simulate document scan forgery or to interact with organizations through additional communication channels such as phone calls. This would yield further insights into the damage a more determined and sophisticated attacker with no regard for the law could cause. To investigate the scalability of the identity theft scenario, the attackers could attempt to automate tasks such as authoring SARs or evaluating replies utilizing AI.

Overall, the theory and evidence we provide on how data privacy and data security interrelate in shaping personal data protection might affect how we study, manage, and regulate data privacy and data security at the individual, organizational, and broader societal level. Most fundamentally, scholars, managers, and regulators might need to attend to both concepts increasingly in combination rather than in isolation to better account for the mutually constitutive, complementary, and at times interfering nature of this relationship.

## Acknowledgments

The authors thank the senior editor, associate editor, and anonymous reviewers for guidance and advice that led to

significant improvements of this paper; Ekaterina Korneeva for immensely helpful support during the planning and execution of our study; and the nine anonymous participants in our simulated identity thefts, who sacrificed their valuable time to make this research project possible.

## Endnotes

[1] There are related but conceptually distinct arguments about tradeoffs between individuals' privacy and *public* security at the societal level, such as public surveillance initiatives that necessarily invade individual privacy (Pavone and Esposti 2010). Such tradeoffs between the individual-level interest in privacy and the societal-level interest in public security are outside the scope of our conceptualization of the interplay between data privacy and *data* security on the individual level.

[2] In order to not risk alarming the victims, this reconnaissance step was purely passive. For example, no friend requests were sent with the goal of gaining access to more sensitive data. Instead, only publicly visible data were gathered.

[3] For example, FirstName.LastName@gmail.com or LastName_FirstCharacterOfFirstName@outlook.com.

[4] Exemplary emails authored by the attackers are provided in the Online Appendix.

[5] Made use of only once to conceal sensitive data about religion, family, childhood, and more. As this occurred in the final iteration of our study, it had no effect on its outcome because there were no subsequent for the data to be deployed in. The victim precisely described what types of data were leaked so this incident could be included in our evaluation.

[6] Details on the interview protocol, interview partners, and interview guideline are presented in the Online Appendix.

[7] An aggregated overview over all data leaks that occurred throughout the iterations is provided in the Online Appendix.

[8] A reasonable range from 1900/01/01 to 2020/01/01 contains only approximately 45,000 passwords to test.

[9] The Online Appendix contains a brief treatise on our victims' awareness of the ongoing attacks.

[10] We asked ChatGPT to author a convincing email using social engineering techniques. We found the result to be just as persuasive as our attackers' emails. Like them, the AI also leveraged the GDPR for persuasiveness.

[11] After our case study had concluded, we made a genuine black-and-white scan of one victim's German national identity card—using a low image resolution on purpose—and redacted all information except for name, date of birth, and address. Because of the poor scan quality and heavily redacted information, this image would have been trivial to forge even for an unsophisticated adversary. All data visible had previously been gathered from interactions with multiple organizations by our attackers. Yet, this document persuaded a DPO who had initially denied the SAR for identity verification reasons, even though that request had already contained all information not redacted in the scan.

## References

Abbasi A, Dobolyi D, Vance A, Zahedi FM (2021) The phishing funnel model: A design artifact to predict user susceptibility to phishing websites. *Inform. Systems Res.* 32(2):410–436.

Acquisti A, Brandimarte L, Loewenstein G (2015) Privacy and human behavior in the age of information. *Science* 347(6221):509–514.

Acquisti A, Taylor C, Wagman L (2016) The economics of privacy. *J. Econom. Literature* 54(2):442–492.

Algarni A, Xu Y, Chan T (2017) An empirical study on the susceptibility to social engineering in social networking sites: The case of Facebook. *Eur. J. Inform. Systems* 26(6, SI):661–687.

Belanger F, Hiller JS, Smith WJ (2002) Trustworthiness in electronic commerce: The role of privacy, security, and site attributes. *J. Strategic Inform. Systems* 11(3–4):245–270.

Biselli T, Reuter C (2021) On the relationship between IT privacy and security behavior: A survey among German private users. Ahlemann F, Schütte R, Stieglitz S, eds. *Innovation Through Information Systems* (Springer International Publishing, Cham, Switzerland), 388–404.

Bufalieri L, Morgia ML, Mei A, Stefa J (2020) GDPR: When the right to access personal data becomes a threat. *Proc. 2020 IEEE Internat. Conf. Web Services* (IEEE, Piscataway, NJ), 75–83.

Cialdini RB (2009) *Influence: The Psychology of Persuasion* (HarperCollins, New York).

Cichy P, Salge TO, Kohli R (2021) Privacy concerns and data sharing in the internet of things: Mixed methods evidence from connected cars. *MIS Quart.* 45(4):1863–1892.

Cram WA, D'Arcy J, Proudfoot JG (2019) Seeing the forest and the trees: A meta-analysis of the antecedents to information security policy compliance. *MIS Quart.* 43(2):525–554.

Culnan MJ, Williams CC (2009) How ethics can enhance organizational privacy: Lessons from the choicepoint and TJX data breaches. *MIS Quart.* 33(4):673–687.

D'Arcy J, Herath T, Shoss MK (2014) Understanding employee responses to stressful information security requirements: A coping perspective. *J. Management Inform. Systems* 31(2):285–318.

Dehling T, Sunyaev A (2024) A design theory for transparency of information privacy practices. *Inform. Systems Res.* 35(3):956–977.

Demirer M, Hernández DJJ, Li D, Peng S (2024) *Data, Privacy Laws and Firm Production: Evidence from the GDPR* (National Bureau of Economic Research, Cambridge, MA).

Di Martino M, Robyns P, Weyts W, Quax P, Lamotte W, Andries K (2019) Personal information leakage by abusing the GDPR 'right of access'. *Proc. 15th Sympos. Usable Privacy Security* (USENIX Association, Santa Clara, CA), 371–385.

Durcikova A, Miranda SM, Jensen ML, Wright R (2024) United we stand, divided we fall: An autogenic perspective on empowering cybersecurity in organizations. *MIS Quart.* 48(4):1503–1536.

Federal Bureau of Investigation (2022) Internet crime report 2022. Accessed September 20, 2023, https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf.

Goel S, Williams K, Dincelli E (2017) Got phished? Internet security and human vulnerability. *J. Assoc. Inform. Systems* 18(1):22–44.

Johnson GA (2022) *Economic Research on Privacy Regulation: Lessons from the GDPR and Beyond* (University of Chicago Press, Chicago).

Johnson GA, Shriver SK, Goldberg SG (2023) Privacy and market concentration: Intended and unintended consequences of the GDPR. *Management Sci.* 69(10):5695–5721.

Ke TT, Sudhir K (2023) Privacy rights and data security: GDPR and personal data markets. *Management Sci.* 69(8):4389–4412.

Li Z, Lee G, Raghu TS, Shi ZM (2025) Impact of the general data protection regulation on the global mobile app market: Digital trade implications of data protection and privacy regulations. *Inform. Systems Res.* 36(2):669–689.

Lowry PB, Dinev T, Willison R (2017) Why security and privacy research lies at the centre of the information systems (IS) artefact: Proposing a bold research agenda. *Eur. J. Inform. Systems* 26(6):546–563.

Mitnick KD, Simon WL (2003) *The Art of Deception: Controlling the Human Element of Security* (John Wiley & Sons, Hoboken, NJ).

Parks R, Xu H, Chu CH, Lowry PB (2017) Examining the intended and unintended consequences of organisational privacy safeguards. *Eur. J. Inform. Systems* 26(1):37–65.

Pavlou PA (2011) State of the information privacy literature: Where are we now and where should we go? *MIS Quart.* 35(4):977.

Pavone V, Esposti SD (2010) Public assessment of new surveillance-oriented security technologies: Beyond the trade-off between privacy and security. *Public Understanding Sci.* 21(5):556–572.

Smith HJ, Dinev T, Xu H (2011) Information privacy research: An interdisciplinary review. *MIS Quart.* 35(4):989.

Smith HJ, Milberg SJ, Burke SJ (1996) Information privacy: Measuring individuals' concerns about organizational practices. *MIS Quart.* 20(2):167.

Solove DJ, Hartzog W (2022) *Breached!: Why Data Security Law Fails and How to Improve It* (Oxford University Press, Cary, NC).

U.S. House 109th Congress (2006) National Defense Authorization Act for Fiscal Year 2006, H.R. 1815, https://www.congress.gov/bill/109th-congress/house-bill/1815.

von Solms R, van Niekerk J (2013) From information security to cyber security. *Comput. Security* 38(1):97–102.

Waldman AE (2020a) Data protection by design? A critique of Article 25 of the GDPR. *Cornell Internat. Law J.* 53(1):147–167.

Waldman AE (2020b) Privacy law's false promise. *Washington University Law Rev.* 97(3):773–834.

Williams EJ, Hinds J, Joinson AN (2018) Exploring susceptibility to phishing in the workplace. *Internat. J. Human-Comput. Stud.* 120(1):1–13.

Willox NA Jr, Gordon GR, Regan TM, Rebovich DJ, Gordon JB (2004) Identity fraud: A critical national and global threat. *J. Econom. Crime Management* 2(1):3–48.

World Economic Forum (2019) The cybersecurity guide for leaders in today's digital world. https://www3.weforum.org/docs/WEF_Cybersecurity_Guide_for_Leaders.pdf.

Wright RT, Jensen ML, Thatcher JB, Dinger M, Marett K (2014) Influence techniques in phishing attacks: An examination of vulnerability and resistance. *Inform. Systems Res.* 25(2):385–400.

Xu F, Wang X, Zhang F (2025) Consumer privacy in online retail supply chains. *Management Sci.* 71(10):8371–8389.

Yin RK (2009) *Case Study Research: Design and Methods*, vol. 5 (SAGE, Thousand Oaks, CA).