

Christoph Bösch*, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher

Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns

Abstract: Privacy strategies and privacy patterns are fundamental concepts of the privacy-by-design engineering approach. While they support a privacy-aware development process for IT systems, the concepts used by malicious, privacy-threatening parties are generally less understood and known. We argue that understanding the “dark side”, namely how personal data is abused, is of equal importance. In this paper, we introduce the concept of privacy dark strategies and privacy dark patterns and present a framework that collects, documents, and analyzes such malicious concepts. In addition, we investigate from a psychological perspective why privacy dark strategies are effective. The resulting framework allows for a better understanding of these dark concepts, fosters awareness, and supports the development of countermeasures. We aim to contribute to an easier detection and successive removal of such approaches from the Internet to the benefit of its users.

Keywords: Privacy, Patterns

DOI 10.1515/popets-2016-0038

Received 2016-02-29; revised 2016-06-02; accepted 2016-06-02.

1 Motivation and Introduction

Over the last years, privacy research has primarily focused on (i) a better conceptual understanding of privacy, (ii) approaches for improving and enhancing privacy protection, as well as (iii) technical mechanisms for implementing these approaches.

***Corresponding Author: Christoph Bösch:** Institute of Distributed Systems, Ulm University, E-mail: christoph.boesch@uni-ulm.de

Benjamin Erb: Institute of Distributed Systems, Ulm University, E-mail: benjamin.erb@uni-ulm.de

Frank Kargl: Institute of Distributed Systems, Ulm University, E-mail: frank.kargl@uni-ulm.de

Henning Kopp: Institute of Distributed Systems, Ulm University, E-mail: henning.kopp@uni-ulm.de

Stefan Pfattheicher: Department of Social Psychology, Ulm University, E-mail: stefan.pfattheicher@uni-ulm.de

However, online service providers have become more and more sophisticated in deceiving users to hand over their personal information. Up until now, privacy research has not studied this development.

An example for this development is the Tripadvisor mobile app (depicted in Figure 1), which is a review platform for travel-related content. At first glance, the starting page asks the user to log in with a personal Google+, Facebook, or email account. Taking a closer look, one notices that a third option is given that offers the creation of a Tripadvisor account. Furthermore, a “Skip”-button is hidden in the upper right corner, which skips the login process entirely. When signing in with Facebook, Tripadvisor wants to gain access to the friend list, photos, likes, and other information (cf. Figure 1b). This is unnecessary for the main features of the service.

Skipping the login process shows the user some features which are available only after signing in (cf. Figure 1c). In addition, the “Later”-button, which finally leads to the app, is located on the left side. Placed on the right side is a “Sign in”-button which leads back to the starting page. This influencing towards logging in via Facebook/Google+ or creating a personal account gives Tripadvisor access to personal information. Figure 1 illustrates general reusable strategies for deceiving users to share more of their personal information.

In this paper, we deliberately change sides and explore the techniques used by the “bad guys” to collect privacy-sensitive data more efficiently. Similar to the collection of well-established privacy solutions (so-called privacy patterns [14, 24]) as part of the privacy-by-design strategy, we identify and collect malicious patterns that intentionally weaken or exploit the privacy of users, often by making them disclose personal data or consent against their real interest.

This approach may initially seem suspicious, as it could provide guidance for malign stakeholders such as data-driven companies or criminals. However, we believe that this shift in perspective is helpful and necessary for privacy research, as it introduces several benefits: (i) A detailed analysis and documentation of privacy dark patterns allows for a better understanding of the underlying concepts and mechanisms threatening the users’




Fig. 1. Screenshots of the Tripadvisor mobile app. (a) shows the starting page. Note the small “Skip” button in the upper right corner. (b) shows the requested personal information when logging in with Facebook. Some of the information is unnecessary for providing the service. (c) shows what happens after skipping the sign-in process.

privacy. (ii) A collection of privacy dark patterns fosters awareness and makes it easier to identify such malicious patterns in the wild. (iii) Furthermore, the documentation of a privacy dark pattern can be used as a starting point for the development of countermeasures, which disarm the pattern and re-establish privacy. The discussion is similar to IT security, where investigation and publication of vulnerabilities proved to be key for actually enhancing the security level in real systems.

1.1 Introduction to Patterns

In many disciplines recurring problems have been addressed over and over again, yielding similar and recurring solutions. The idea of a pattern is to capture an instance of a problem and a corresponding solution, abstract it from a specific use case, and shape it in a more generic way, so that it can be applied and reused in various matching scenarios.

Patterns originate from the realm of architecture, where Alexander et al. [5] released a seminal book on architectural design patterns in 1977. In this book, the authors compiled a list of archetypal designs for buildings and cities which were presented as reusable solutions for other architects. Interestingly, Alexander et al. already came up with patterns for privacy. For instance, their *Intimacy Gradient* pattern postulates a placement

of chambers in such a way that a further distance from the building’s entrance allows for increased intimacy.

In 1987, the idea of patterns was readopted by Kent and Cunningham [10] and introduced into the realm of computer science and software development. The Portland Pattern Repository of Kent and Cunningham collected patterns for programmers using object-oriented programming languages.¹ The idea of using patterns in software design gained wide acceptance in 1994, when the so-called Gang of Four released their well-known book on design patterns for reusable object-oriented software [19]. Since then, the usage of patterns has spread to various different branches of computer science and software engineering, including distributed architectures [18, 25], user interface design [46], IT security [41], and privacy [14, 22, 40, 42].

The success of patterns in software engineering has entailed novel classes of patterns with different semantics, namely anti patterns [14] and dark patterns [11]. Traditional design patterns capture a reasonable and established solution. In contrast, an anti pattern documents a solution approach that should be avoided, because it has been proven to represent a bad practice.

¹ Historical side note: the online version of the pattern repository, WikiWikiWeb (<http://c2.com/cgi/wiki/>), became the first-ever wiki on the World Wide Web

Hence, anti patterns raise awareness of sub-par solutions and advocate against their usage.

Anti patterns often target solutions that may seem obvious to the system developer at a first glance, but include a number of less obvious negative implications and consequences. Even established design patterns sometimes become obsolete and are downgraded to anti patterns due to new considerations. For instance, the Gang of Four suggested a pattern for restricting the instantiation of a class to a single instance, the so-called Singleton pattern. Only after concurrent programming and multi-core architectures became more widespread, shortcomings of this pattern eventually became apparent. Today, the Singleton pattern is widely considered an anti pattern.

The term dark pattern was first used by Brignull, who collected malicious user interface patterns [11] for better awareness. A UI dark pattern tricks users into performing unintended and unwanted actions, based on a misleading interface design. More generally speaking, a dark pattern describes an established solution for exploiting and deceiving users in a generic form.

In summary, anti patterns collect the *Don'ts* for good intentions and dark patterns collect potential *Dos* for malicious intents. In this paper, we present a first broad discussion on dark patterns in the field of privacy.

1.2 Methodology

To suggest a framework for the collection of privacy dark patterns and to compile a list of such patterns, we consider the problem from three different angles as part of a holistic approach.

First, we survey existing literature on privacy strategies and privacy patterns. We then reverse privacy strategies and adapt some of these ideas and extend them, so that they become malicious patterns. Beyond this, we have identified new types of patterns. Second, we include a psychological point of view on malevolent privacy concepts. This perspective takes into account human information processing, social cognition and motivation, as well as exploitable basic human needs. On this basis we are able to deduce additional approaches on how to reduce the power of privacy dark strategies. Third, we identify and analyze real-world examples of malicious privacy mechanisms as found on websites and in mobile applications.

Next, we integrate these findings on privacy dark patterns into a unified framework, which introduces a

general terminology for privacy dark patterns and establishes a template for documenting privacy dark patterns. Our framework suggests a list of malicious privacy strategies and psychological aspects for categorizing privacy dark patterns. Based on the pattern template of our framework, we discuss common privacy dark patterns that we extracted from real-world occurrences.

1.3 Contribution

Our contribution can be summarized as follows:

1. We introduce the *concept of privacy dark strategies and privacy dark patterns*.
2. We present a *framework for privacy dark patterns* that takes into account traditional privacy patterns, empirical evidence of malign patterns, underlying malicious strategies, and their psychological background. The resulting framework provides a template for documenting and collecting arbitrary privacy dark patterns.
3. We provide an initial *set of exemplary dark patterns* that we encountered in the wild.
4. We launched the website dark.privacypatterns.eu as an *online collection for privacy dark patterns*. Being a collaborative resource, we invite the community to submit more patterns and help to raise awareness.

2 On Privacy Strategies and Privacy Patterns

In this section, we introduce privacy patterns and corresponding privacy strategies, based on their historical development.

Until the mid-1990s, privacy was rarely considered a relevant feature of IT systems. Even if it was, the integration of privacy-preserving mechanisms was often conducted a posteriori, as an additional requirement later added to the system. The notion of “privacy as an afterthought” contradicted the cross-sectional property of privacy as part of an IT system and often yielded extensive or insufficient changes to the system.

To overcome these deficits, a joint team of the Information and Privacy Commissioner of Ontario, Canada; the Dutch Data Protection Authority; and the Netherlands Organisation for Applied Scientific Research advocated a more integral approach that included privacy considerations into the overall development cycle [27]. In 1995, they introduced the so-called *Privacy by Design*

approach,² which is postulated by the following seven foundational principles:

1. Proactive not reactive
2. Privacy as the default setting
3. Privacy embedded into design
4. Full functionality
5. End-to-end security
6. Visibility and transparency
7. Respect for user privacy

These principles have been a major milestone for the design of privacy-preserving systems as they provide general guidance. For that reason, these concepts are part of several privacy legislations today.

One frequent criticism regarding Privacy by Design and its seven principles is that they are too unspecific to be directly applied to a development process. The principles neither provide concrete advice, nor do they address the varying needs of specific domains, such as the Internet of Things, User Interface Design, or Car-2-Car communication. A system designer is thus still required to have a thorough understanding of privacy and the forces involved, to design a privacy friendly system. Clearly, more guidance and a more methodological approach is required to establish a privacy engineering process as, for example, worked on by the PRIPARE project [38].

One element of this privacy engineering approach are so-called *Privacy Patterns*. The main idea of privacy patterns is to improve the drawbacks of the Privacy by Design principles, i.e., that they are not actionable [14, 21, 48]. Privacy patterns are defined as reusable solutions for commonly occurring problems in the realm of privacy. Essentially, they are patterns for achieving or improving privacy. Privacy patterns provide guidance for engineering and development, and target the needs of specific domains, such as backend implementations or user interface design. By providing a well-structured description of a problem and its solution using a standardized template, patterns can easily be looked up and applied. Since these patterns include references to specific use-cases and possibly implementations, engineers will directly find the resources needed to implement them in their own context.

One well-known example of a privacy pattern that can be implemented in multiple domains is the stripping of metadata that is not necessary for the functionality of the service. This procedure increases privacy,

since metadata often includes personally identifiable information. Further, this solution is reusable, since it is not bound to a specific instance of a problem. Thus, stripping of metadata constitutes a privacy pattern that can be applied, e.g., to a website managing digital photographs.

A single privacy pattern addresses a problem with a limited scope. Multiple related and complementing patterns can then be compiled into a pattern catalog. Similar to the well-known design pattern catalogs from software engineering, a privacy pattern catalog collects a number of relevant problems and suitable solutions that can be applied during a privacy-aware development phase.

There are multiple collections of privacy patterns from academic research [14, 22, 40, 42] as well as online repositories³ that are more accessible for practitioners.

In a typical system development process, privacy patterns are applied during the stages of design and implementation. However, in many scenarios, privacy aspects represent fundamental system requirements that have to be considered from the very beginning. The question is, whether more general architectural building blocks exist that can be applied at an even earlier stage, i.e., during requirement analysis and architectural design. Note that this notion is a natural continuation of the Privacy by Design philosophy—to include privacy considerations into the entire development process.

These general architectural building blocks are known as *Privacy Design Strategies*. According to Hoepman [24], a privacy design strategy is on a more general level than a privacy pattern and “describes a fundamental approach to achieve a certain design goal. It has certain properties that allow it to be distinguished from other (fundamental) approaches that achieve the same goal.”

Later in the development process, a privacy design strategy can be refined with privacy patterns implementing one or more strategies. Thus, privacy design strategies provide a classification of privacy patterns. When system designers search for a privacy pattern in a collection, they are only interested in the ones implementing their chosen privacy strategy.

Hoepman [24] defines the following eight privacy design strategies.

MINIMIZE: Data minimization is a strategy which insists that the amount of personal information that is processed should be minimal. Data that is not

² <https://privacybydesign.ca/>

³ <https://privacypatterns.eu>, <http://privacypatterns.org/>

needed for the original purpose should not be collected.

HIDE: HIDE takes place after data collection. Whereas MINIMIZE forbids the collection of needless information, HIDE suggests that any personal data that is processed should be hidden from plain view.

SEPARATE: The approach of the privacy strategy SEPARATE is to process any personal information in a distributed fashion if possible. Thus, interrelationships between personal data vanish in contrast to a centralized processing.

AGGREGATE: When implementing AGGREGATE, personal information is processed at a high level of aggregation. This level should only be so high as to remain useful, however. Details that are not needed for the functionality of the service vanish. This process could include statistical aggregation such that the details of identities are blurred.

INFORM: The privacy strategy INFORM states that data subjects should be adequately informed whenever personal information is processed.

CONTROL: A common requirement of software systems is that data subjects should be in control of the processing of their personal information. Whenever this is ensured, we are dealing with the privacy strategy CONTROL. Hoepman states that he is not aware of any patterns implementing this strategy.

ENFORCE: ENFORCE states that a privacy policy that is compatible with legal requirements should be in place and should be enforced.

DEMONSTRATE: The privacy strategy DEMONSTRATE demands that data controllers are able to demonstrate compliance with their privacy policy and any applicable legal requirements. A good example for a pattern implementing this strategy is the use of audits.

In the following sections, privacy design strategies serve as the starting point for our analysis of malicious dark strategies that harm privacy. For defining and documenting malicious patterns, we adapt the idea of privacy patterns and transform it into privacy dark patterns.

3 The Dark Side

The triad of general privacy strategies for high-level privacy requirements, privacy patterns for privacy-aware design processes, and privacy-enhancing technologies

for system implementations is commonly acknowledged when building privacy-friendly IT systems.

However, there are other parties that have different agendas when building IT systems. Instead of privacy-friendly solutions, they aim for systems that purposefully and intentionally exploit their users' privacy—for instance motivated by criminal reasons or financially exploitable business strategies.

For the development of our framework, we reverse the evolution of privacy strategies and patterns: First, we define dark strategies as the high-level goals that these parties follow in order to exploit privacy. Next, we derive suitable dark patterns that implement these strategies. We then complement our framework by adding a psychological perspective on how the strategies generally achieve their deceptive and manipulative goals. Note that we do not include a counterpart to privacy-enhancing technologies as part of our framework.

As already clarified in the introduction, the resulting framework is neither intended nor structured as a construction kit for malicious parties. Instead, the framework can be used by privacy researchers and practitioners for detecting, recognizing, analyzing, and documenting malicious strategies and patterns.

When used *top-down*, the framework supports a privacy analysis of IT systems by raising awareness for malicious strategies and by uncovering corresponding mechanisms. *Bottom-up*, the framework helps to identify malicious patterns, reveals underlying strategies, and provides pointers for the development of concrete countermeasures.

3.1 Privacy Dark Strategies

We now develop a categorization of privacy dark patterns, analogously to Hoepman's privacy design strategies [24]. As privacy design strategies can be used to *categorize* privacy patterns by their fundamental approach, the same holds for privacy dark strategies. Hoepman identified eight privacy strategies, namely MINIMIZE, HIDE, SEPARATE, AGGREGATE, INFORM, CONTROL, ENFORCE, and DEMONSTRATE. Based on these strategies, we identify the following privacy dark strategies: MAXIMIZE, PUBLISH, CENTRALIZE, PRESERVE, OBSCURE, DENY, VIOLATE, and FAKE as shown in Table 1. These are used for our categorization of privacy dark patterns in Section 5.

The privacy design strategy MINIMIZE, for example, demands the amount of processed data to be re-

Table 1. Privacy Strategies vs. Dark Strategies.

Strategies	
Hoepman	Dark Strategies
MINIMIZE	MAXIMIZE
HIDE	PUBLISH
SEPARATE	CENTRALIZE
AGGREGATE	PRESERVE
INFORM	OBSCURE
CONTROL	DENY
ENFORCE	VIOLATE
DEMONSTRATE	FAKE



stricted to the minimal amount possible. The corresponding dark strategy MAXIMIZE would collect, store, and process as much data as possible, leading to a loss of privacy. The system designer does not act out of pure maliciousness but to gain an advantage over a system with the same functionality but with stronger privacy protection. Specifically, by receiving additional personal data which can, e.g., be sold or used for personalized advertisements.

In the following we detail the eight privacy dark strategies we have developed.

Maximize. The goal of the dark strategy MAXIMIZE is to collect an inappropriate amount of data. More precisely MAXIMIZE means that...

The amount of personal data that is collected, stored, or processed is significantly higher than what is actually needed for the task.

Examples would be extensive sign-up forms with fields that are not needed for the functionality of the service. Often those unneeded fields are mandatory, maximizing the collection of personal data. Another example of a MAXIMIZE strategy are bad default privacy settings or the necessity to set up an account for the usage of a service, especially if the account is not needed for the functionality of the service.

Publish. The dark strategy PUBLISH can be characterized by the requirement that...

Personal data (not intended to be public) is not hidden from plain view.

This means that often no mechanism is in place to hide personal data from unauthorized access, such as encryption or access control. The personal data lies in the open for everyone to see. Social networks often employ

this dark strategy to encourage the sharing of personal data and thus the use of their platform. This strategy satisfies a person’s *need to belong* as will be explained in section 4.

Centralize. CENTRALIZE is the dark strategy associated to the privacy strategy SEPARATE, which mandates that personal data should be processed in a distributed way. CENTRALIZE, in contrast, enforces that...

Personal data is collected, stored, or processed at a central entity.

This strategy preserves the links between the different users and thus allows for a more complete picture of their habits and their usage of the service.

Advertising networks employ this strategy heavily by sharing pseudonymous user IDs, a practice known as cookie syncing [1]. Another common occurrence of this privacy dark strategy is the practice of flash cookies, which are cookies that are stored centrally by the flash plug-in on the file system and are thus not restricted to a specific web browser.

Preserve. The dark strategy PRESERVE requires that...

Interrelationships between different data items should not be affected by processing.

They should rather be preserved in their original state for analysis instead of storing them in a processed form, e.g., aggregation. It is not necessary to know the type of analysis in advance. A prominent example is telecommunications data retention because traffic analysis can recover the relationships between persons.

Obscure. In the dark strategy OBSCURE...

It is hard or even impossible for data subjects to learn how their personal data is collected, stored, and processed.

Users should be unable to inform themselves about what happens to their disclosed data. This can be achieved in the form of a privacy policy with many technical terms, which are difficult to understand for the average user. User Interfaces could be designed to mislead the user, leading to decisions contradicting the user’s original intent. The EFF called this particular mechanism “privacy zuckering” [28].

Deny. Patterns making use of the dark strategy DENY make a data subject lose control of their personal data. The term DENY is due to a denial of control.

Data subjects are denied control over their data.

With this dark strategy, a service provider can prevent users from taking actions that oppose that service provider's interest. An example is to not provide the functionality for deleting an account. Another example is the nonexistence of options to control sharing of information. Until recently this was the case in WhatsApp, where the online status was automatically shared with everyone who subscribed to that phone number, which has a big impact on the privacy of users [12].

Violate. The strategy VIOLATE occurs if . . .

A privacy policy presented to the user is intentionally violated.

A privacy policy is in place, shown to the user but intentionally not kept. The users are unaware of the violation; thus, this does not impact the trust put into that service if such violations are not revealed. It is hard to find concrete examples and patterns implementing this strategy since using this strategy is against the law and not publicly admitted by companies.

Fake. The privacy dark strategy FAKE means that . . .

An entity collecting, storing, or processing personal data claims to implement strong privacy protection but in fact only pretends to.

An example of this strategy are self-designed padlock icons or privacy seals, which make the user feel secure but do not have any meaning. Another example are wrong and unsubstantial claims such as an unrealistic claim on the key-size of ciphers or marketing terms like "military grade encryption".

Synthesis

Our eight Privacy Dark Strategies can be summarized as follows:

- MAXIMIZE: The amount of personal data that is collected, stored, or processed is significantly higher than what is actually needed for the task.
- PUBLISH: Personal data is published.

- CENTRALIZE: Personal data is collected, stored, or processed at a central entity.
- PRESERVE: Interrelationships between different data items should not be affected by processing.
- OBSCURE: It is hard or even impossible for data subjects to learn how their personal data is collected, stored, and processed.
- DENY: Data subjects are denied control over their data.
- VIOLATE: A privacy policy presented to the user is intentionally violated.
- FAKE: An entity collecting, storing, or processing personal data claims to implement strong privacy protection but in fact only pretends to.

3.2 Privacy Dark Patterns

After our exploration of privacy dark strategies we will now define the concept of a privacy dark pattern. As mentioned in Section 2, a pattern describes a generic, reusable building block to solve a recurring problem and hence to document best practices. They can be collected in special catalogs and allow for easy replication. Patterns fulfill the role of a common language to allow system developers and privacy engineers to communicate more efficiently.

We argue that common building blocks that are used by service providers to deceive and mislead their users exist. Some service providers use recurring patterns to increase the collection of personal data from their users. Sometimes these building blocks are used unintentionally, simply constituting usage of privacy anti patterns, but without any malicious intent. However, we claim that there are building blocks which are used on purpose, thereby yielding an advantage to the service provider. We call these building blocks *privacy dark patterns*.

Analogously to privacy patterns, privacy dark patterns can be collected in special repositories to facilitate easy access and retrievability for users and to develop countermeasures. Patterns are usually documented in a formalized template to enable system developers to easily reference and use them. Common fields in such a template include the name of the pattern, the problem the pattern is solving and references to related patterns.

However, current templates for design and privacy patterns are not suitable for documenting privacy dark patterns due to the following reasons:

1. Privacy patterns and privacy dark patterns have a different intent regarding their documentation.

Each privacy pattern solves a specific problem, which is often mentioned as a separate field in the template. Privacy patterns are documented to be copied and used. The purpose of documenting privacy dark patterns on the other hand is to create and enhance awareness about common anti-privacy techniques, since they do not solve an engineering problem. Thus, a problem-centric description is out of place.

2. The target group of privacy patterns are system designers whereas privacy dark patterns can target non-technical end-users to educate them about the strategies that are used to deceive them.

Thus, we need a different template to document privacy dark patterns.

Our Privacy Dark Pattern Template

We have developed a new template, specifically targeted towards privacy dark patterns, which we explain in detail in the following.

Name/Aliases: This field describes the name under which the privacy dark pattern is known. The name should be concise and capture the essence of the pattern.

Summary: A short summary to describe the pattern is necessary to provide an overview of the pattern and for quick reference.

Context: Context describes the scenario in which the pattern appears. e.g., online social networks or online shops.

Effect: This section explains the effects and consequences of the pattern. This should be described with sufficient granularity such that it is not too general.

Description: In this part of the template, the privacy dark pattern is described in detail. Technical language can be used if not avoidable, but it should be remembered that the main target group of the pattern are the end-users of the system in which the privacy dark pattern is applied.

Countermeasures: The countermeasures describe behaviors and tools a user can implement to negate the effects of the privacy dark pattern. These are strategies to help the “victims” of the pattern regain or maintain their privacy. This includes procedures to avoid the effects of the pattern, as well as add-ons to existing programs, e.g., web browsers, which prevent the end-user from being deceived by the pattern.

Examples/Known Uses: In this section, implementations using the dark pattern are described. Service providers applying the privacy dark pattern belong into this field. Screenshots of usage of the dark pattern can be provided where appropriate.

Related Patterns: If related privacy dark patterns exist they are referenced here.

Psychological Aspects: This field describes the psychological mechanisms that make the pattern effectively influence the users in their behavior.

Strategies: In this part of the documentation of a privacy dark pattern, the used dark strategy is provided. These are the dark strategies explained in Section 3.1.

This template can be used to systematically document different privacy dark patterns in a repository. We make use of this template later in Section 5.

4 Psychological Aspects

In the following, we address the question why privacy dark patterns do actually work. One can reasonably assume that there is, at least to some degree, awareness among a majority of users that privacy dark strategies exist and some service providers have strong incentives to violate the privacy of their users. It is similarly likely that users notice, at least sometimes, when they are being targeted by privacy dark strategies. Nevertheless, privacy dark strategies still work, as indicated by their frequent occurrence. This somewhat paradoxical situation can be explained by adopting a psychological perspective on privacy dark strategies.

Essentially, privacy dark strategies often work well because they take advantage of the psychological constitution of human beings. In this regard, we focus on the ways in which humans think and reason, i.e., humans’ cognitive information processing.

There is widespread agreement in the field of psychological research that two different cognitive systems underlie thinking and reasoning processes [29, 43, 44]. For instance, when creating a new account on a website, a user is often asked to agree to a list of general terms and conditions. Most likely, they will not read the page filled with these terms and conditions, but will agree to them quickly, intuitively, and automatically. This is an example of a System 1 thinking process; it takes place automatically, unconsciously, and with little effort [29, 43, 44].

Instead of agreeing to general terms and conditions quickly and automatically, one can take the time and make the effort to carefully read the information provided. Afterwards, one deliberately weighs the pros and cons and decides whether to agree to the conditions or not. This is an example of a System 2 thinking process; it takes place in a controlled, conscious, and effortful way. Behavior based on System 2 thinking is driven by a deliberative, effortful decision-making process, resulting in the relatively slow execution of behavior [29, 43, 44].

General terms and conditions are often not read, and agreement is typically made automatically and quickly, i.e., System 1 operates. There is thus an opportunity to fill general terms and conditions with dark ingredients. These in turn are not consciously noticed when users are in System 1 mode, as illustrated in the example. In general, we postulate that privacy dark strategies work well when individuals process information using System 1 thinking. When (dark) information is processed quickly, without much effort, and automatically, it seems likely that privacy dark strategies can unleash their full impact. In other words, in System 1 mode, subjects are likely to be less conscious of privacy dark patterns being at work and unable to deliberately act against them. On the other hand, recognizing privacy dark strategies and taking action against them requires System 2 processing.

Past research in fact shows the importance of cognitive information processing for privacy issues (e.g., [8, 32, 34]). Knijnenburg and colleagues [31], for instance, document that people automatically provide personal information on website forms when an auto-completion feature fills out forms by default with previously stored values. Reducing the impact of this automatic (System 1 based) default completion by giving users control over which forms are filled out reduces the amount of personal information provided.

A number of conditions determine whether humans rely on System 1 thinking processes and System 2 thinking processes are inhibited. There are two central aspects to consider [16, 39]. Humans engage in System 1 processing whenever they (a) have little *motivation* to think and reason in an effortful way or (b) have no *opportunity* to do so because they lack the required knowledge, ability, or time. Users, for instance, often have no motivation to read general terms and conditions. In instances where they are motivated, they often do not have the opportunity to use System 2 thinking because the language used in general terms and conditions often

is too complicated and subjects are unable to interpret this information [35].

4.1 Prompting System 1 Thinking

As argued above, privacy dark strategies are typically accompanied by System 1 thinking processes, while System 2 thinking processes are often not possible, as shown in the following analysis. Regarding the dark strategy MAXIMIZE, the amount of data that is processed is significantly higher than the data that is really needed for the task. Subjects need high motivation to resist excessive data collection. Additionally, although some users might have high motivation, they need specific knowledge and abilities to offer any resistance. However, some service providers use mandatory form fields for user registration, which renders the knowledge to circumvent the dark strategy useless if one wants to utilize the service. Thus, users often stay in System 1 mode and allow MAXIMIZE to operate.

When personal data is not hidden from plain view (PUBLISH), users need to be motivated and able to change settings. Users might lack the necessary motivation and ability to do so; thus, remaining in System 1 processing when it comes to, for instance, privacy settings.

Working against the dark strategies of centralizing personal data (CENTRALIZE) and of providers that interrelate data items (PRESERVE) requires particularly high motivation as well as extensive knowledge of and the ability to understand these strategies. It is reasonable to assume that the typical user often does not have the knowledge and ability to precisely understand the dark strategies of CENTRALIZE and PRESERVE and to work against them (e.g., taking action against data preservation). Thus, users often cannot engage in the deliberative processing that might lead to behavior that challenges these two dark strategies.

The dark strategy OBSCURE reflects the idea that it is difficult for users to learn about what happens to their personal data. This strategy implies that users must be highly motivated and able to acquire information about how their personal data is used and stored. Again, this requirement inhibits users from engaging in deliberative processing.

Analogously, when users' control of data is denied (DENY) they must be highly motivated and able to work against this strategy. DENY makes it even more difficult for users to notice the violation of privacy policies and legal requirements (VIOLATE). Here, high motiva-

tion and ability is needed to enable users to notice and to work against this dark strategy.

When certificates or other information is faked (FAKE), users need to be motivated to search for this information. Additionally, they need the ability to judge whether information has been faked or not. If motivation and ability is not present, subjects will process (fake) information using System 1 thinking and will likely not notice the privacy dark strategy.

To sum up, it is evident that privacy dark strategies work, because users often do not have the motivation or opportunity to resist them. As such, System 2 thinking processes are often absent, while System 1 thinking does accompany the use of privacy dark strategies. Building on these considerations, one can deduce suggestions on how to reduce the power of privacy dark strategies. Specifically, we argue for attempts to strengthen System 2 thinking processes by increasing motivation (e.g., through emphasizing the negative impact of privacy dark strategies) and opportunities for resistance (e.g., by increasing knowledge about privacy dark strategies as advocated by this paper, or by implementing tools that reduce automatic provision of private information [31]).

4.2 Humans’ Fundamental Need to Belong

Beyond the idea that human information processing is involved in the functioning of privacy dark strategies, humans’ fundamental needs also contribute to the effectiveness of some privacy dark strategies. Humans possess basic needs, e.g., safety and security needs, concerns about physical well-being, the need for self-esteem, and the need to belong to significant others [23]. We identified the need to belong as particularly important for why some privacy dark strategies work well. The argument that is put forward states that individuals’ need to belong forces people to disregard privacy issues.

The need to belong reflects humans’ desire to be an accepted member of a group. Psychological experiments (e.g., Williams et al. [49]) show that social exclusion of a subject, even by unknown other subjects in a simple ball game played on the Internet, reduced subjects’ well-being, their belief in a meaningful existence, and their self-esteem. People’s need to belong manifests as a concern for being liked and admired by others, as is evident in social networks [20]. The need to belong motivates people to accumulate social capital [9], i.e., to establish relationships with other people (e.g., in social

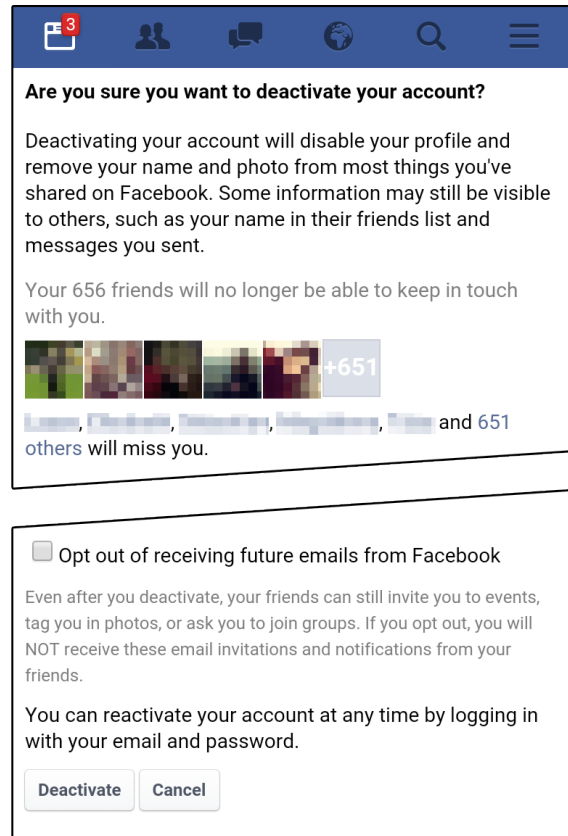


Fig. 2. Dialog of the Facebook mobile website when deactivating the account. The page shows profile pictures of contacts the user has recently interacted with and states that they will miss the user when deactivating the account. Facebook targets the user’s need to belong and provokes a reconsideration.

networks) that serve as personal resources for individuals’ well-being and functioning [7, 15, 26].

Although important for human beings [9], the need to belong might counteract privacy concerns. For example, when personal data is not hidden from plain view (PUBLISH), it can create a possibility of being liked and admired by others, which can fulfill one’s need to belong (cf. Nadkarni and Hofmann [37]). This may lead to a reduced level of privacy at the same time. Furthermore, it is hard for subjects to learn about what happens to the personal data (OBSCURE) they share based on their need to belong.

Service providers might further MAXIMIZE the amount of data based on subjects’ need to belong to gain information about their users, specifically about their social capital [15]. This information is then used to again target subjects’ need to belong, for instance when a user wants to unsubscribe. Facebook, for example, writes “Your [number] friends will no longer be able to keep in touch with you.”, and “[Name] will miss you”

(status January 16, 2016). As shown in Figure 2, users' motivation to unsubscribe is challenged by activating their need to belong and the presentation of users' social capital they would lose once they unsubscribe [7, 26].

In sum, people provide and share private information based on their need to belong. Therefore, the need to belong may run counter to high privacy standards.

4.3 Specific Mechanisms

In summary, privacy dark strategies often work well because they take advantage of human beings' psychological constitution. We argue that System 1 thinking and the need to belong are so fundamental for malicious privacy mechanisms to work, that both aspects represent the basis of psychological considerations in our framework. Furthermore, we believe that both aspects are helpful for contributors when briefly assessing potential privacy dark patterns and their psychological mechanisms.

The discussion whether a pattern is to be regarded as a dark pattern can then easily integrate a perspective of the users. This psychological perspective complements the assessments of actual impacts of the pattern and suspected motives of the service providers. This is important in order to differentiate actual privacy dark patterns with malicious intent from other forms of poorly implemented or unintended features regarding privacy.

Apart from the thinking and reasoning processes and the need to belong mentioned before, arbitrary patterns may exploit more specific psychological mechanisms which build upon these fundamental aspects. In the following, we introduce some of these mechanisms and indicate their usage for privacy dark patterns.

First, we focus on nudging, a concept for influencing decision making based on positive reinforcement and non-forced compliance [45]. Nudging has already been applied to decision making in the domain of privacy protection [3]. For instance, regular nudges that provide a user with information about data collection of smartphone applications have shown to increase awareness and motivate users to reassess the applications' permissions [6]. When the good intents of privacy nudging are replaced with malicious intents, the concept turns into a latent manipulation technique for *non-forced compliance* with weakened privacy. The dark counterpart provides choice architectures facilitating decisions that are negative to the user's privacy. For instance, the starting screen in Figure 1 does not force an account creation

and it provides a skip option. Still, the form design latently manipulates the user by encouraging the creation of a user account.

A stronger form of manipulation is achieved by applying traditional *persuasion techniques* [13]. For instance, the so-called "door in the face" technique takes advantage of the principle of reciprocity. In this technique, the refusal of a large initial request increases the likelihood of agreement to a second, smaller request. This technique has already been studied in the context of private information disclosure [4] and privacy user settings [30]. Applied to privacy dark strategies, a service provider might intentionally ask users for disproportionate amounts of personal data. By providing an option to skip the first form and then only asking for a very limited set of personal data in the second form (e.g., mail address only), users may be more willing to comply and to provide that information after all.

Closely related to the two cognitive systems, heuristics and biases provide decision guidance in case of uncertainty [47]. Although there are a lot of *heuristics and cognitive biases* related to decision making [29] that could be exploited by dark privacy patterns, we will only introduce an exemplary bias that we later use in one of our example patterns: Hyperbolic discounting [33] is a bias causing humans to inconsistently value rewards over time. Also known as present bias, this bias tricks humans into favoring a present reward over a similar reward at a later point in time. In terms of privacy, many users tend to focus on the *instant gratification* of an immediate reward, when they are forced to provide personal data to use a service. At the same time, the users discount the ramifications of privacy disclosures in the future [2].

Cognitive dissonance [17] is a state of discomfort caused by contradictory beliefs and actions. According to the theory of cognitive dissonance, the experience of inconsistency triggers a reduction of dissonance and a potential modification of the conflicting cognition. In terms of privacy dark patterns, this process can be exploited by inconspicuously providing justification arguments for sugarcoating user decisions that have negatively affected their privacy. For instance, after asking users for inappropriate amounts of personal data, a service provider would later remind the users of the high data protection standards they comply with. When a user hesitantly provides personal data although they are generally very cautious regarding personal information, a state of discomfort may emerge soon after. Such hints may then influence the dissonance resolution of the user.

5 Dark Patterns in the Wild

This section introduces patterns of frequent malicious privacy behavior. For this purpose, we surveyed popular web sites and mobile applications and gathered reports of recent privacy incidents. Next, we analyzed the underlying concepts and mechanisms regarding malicious effects on privacy, assessed their impacts, and estimated the intentionality of the service providers. Based on our framework, we then extracted a number of common dark privacy patterns and described them using our pattern template. The resulting list is not exhaustive, but illustrates the idea of privacy dark patterns based on exemplary sightings in the wild.

Of course we cannot clearly determine whether the service providers mentioned as examples in the following patterns actually had a malicious intent, and we are not claiming they did. It is still reasonable to believe that many of the companies offering free services and apps have strong motivations to gather as much data from their customers as possible and design their mobile web services and mobile applications on purpose following such privacy dark patterns. In any case, the examples are helpful to understand the mechanics of the privacy dark pattern in question.

Please note that the following patterns are shortened and use a condensed structure. The extended versions of the patterns based on our full template structure are available at our online portal dark.privacypatterns.eu.

5.1 Privacy Zuckering

The term Privacy Zuckering was first introduced by Tim Jones in an EFF article [28] for “deliberately confusing jargon and user-interfaces”, and was later used on darkpatterns.org for a UI dark pattern. For our catalog, we generalize the idea and present it as a universal privacy dark pattern.

Name/Aliases: *Privacy Zuckering*

Context: The access and usage of personal data is often governed by user-specific, modifiable privacy settings. By doing this, users can choose privacy settings that reflect their own privacy requirements.

Description: A service provider allows users to change their privacy settings. However, the settings are unnecessary complex, overly fine-grained, or incomprehensible to the user. As a result, the user either gives up, or

makes unintended changes to their privacy settings.

Effect: While the service provider will claim that users have full control over their privacy settings, the presentation, terminology and user experience will highly discourage users from making changes. When combined with the *Bad Defaults* pattern, these patterns facilitate the enforcement of privacy settings suggested by the service provider. Privacy Zuckering could lead to unintentional changes of privacy settings, when the complexity of the settings does not align with the user’s perception, and hence prevents originally intended preference adjustments.

Countermeasures: When service providers apply *Privacy Zuckering*, users require help of third parties that clarify the settings and guide them through the intended preferences.

Examples/Known Uses: In the past, Facebook has been accused of applying *Privacy Zuckering* to their users’ privacy setting pages, which termed the mechanism in the first place [11]. For instance, in August 2010, an updated privacy settings page of Facebook allowed for highly customized settings, but required users to change dozens of settings on multiple pages to maximize personal privacy.

Related Patterns: When *Bad Defaults* are in place, *Privacy Zuckering* prevents changes and increases the number of retained default settings.

Psychological Aspects: Overly complex settings and inappropriate terminology requires *System 2 thinking*. When a user is *motivated* to change their settings, but is overwhelmed at the same time, and hence lacks the *opportunity* to do so purposefully, the user may either switch back to *System 1 thinking* and make vague changes, or the user may refrain from doing so at all.

Strategies OBSCURE

5.2 Bad Defaults

Name/Aliases: *Bad Defaults*

Context: This dark pattern is used mainly on websites, by applications, or in social networks. For *Bad Defaults* to have an effect it is often necessary that the system has some form of user accounts.

Description: When creating an account at a service provider the default options are sometimes chosen badly in the sense that they ease or encourage the sharing of personal information. Most users will be too busy to look through all the options and configure their account properly. Thus, they often unknowingly share more personal information than they intend to.




Fig. 3. Facebook default settings from 2010. The graph shows which information can by default be accessed by You, your Friends, Friends of Friends (FoF), all Facebook user (FBU), and the whole Internet (Inet). For the source and more details we refer to <http://mattmckeeon.com/facebook-privacy/>.

Effect: This pattern causes the user to share more information with the system or other users than the user intends to do. This includes but is not limited to which sites the user visits, parts of his user profile, and his online status.

Countermeasures: Users need to be educated to develop more awareness of bad default settings so that they become self-motivated to configure their accounts properly. However this is hard to achieve.

Examples/Known Uses: Facebook Default Privacy Settings (cf. Figure 3).

Related Patterns: *Privacy Zuckering* demotivates users from changing the defaults.

Psychological Aspects: When users are not aware of the defaults that are in effect, a *deliberative processing* of this information is inhibited.

Strategies: OBSCURE

5.3 Forced Registration

Name/Aliases: *Forced Registration*

Context: This pattern can be applied in nearly every service which provides some functionality to users. When the functionality technically requires an account, e.g., in online social networks, this pattern degenerates. In this case we are not speaking of a privacy dark pattern anymore since without an account the service cannot be provided in the intended way.

Description: A user wants to use some functionality of a service which is only accessible after registration. Sometimes this is necessary to use the service in a

meaningful way or prevent misbehavior. But very often this is unnecessary and serves the interest of the service provider by giving him access to (unneeded) personal data. The personal information collected regularly includes an e-mail address, since this is required for creating the account, but is often augmented by birthdates, home addresses, etc.

Effect: The effect of this pattern is that the user is forced to register an account at the service provider, thereby allowing the service provider to track user behavior on his platform. Additionally the registration process often requires an e-mail address and other personal identifiable information. Since the user does not want to have an account in the first place, the user is unlikely to configure the settings properly, thereby possibly revealing even more personal information not intended for disclosure.

Countermeasures: One countermeasure is to create a new account and fill it with random data. Often, one can use an anonymous one-time e-mail address⁴ during registration to receive the activation link for the account.

Another countermeasure is provided by the service BugMeNot⁵. They enable users to bypass the forced registration by allowing many users to share their account details creating a large anonymity set. A user can try accounts published at BugMeNot for using the service. BugMeNot allows users to create new accounts and share them with other users of BugMeNot. It can even be used as a browser extension by some web browsers.

Examples/Known Uses: As of Feb. 2016, the popular question-and-answer website Quora.com requires external visitors to sign up and log in when opening a question page. While the page is rendered initially, it is then blocked by pop-up modal dialog that forces visitors to register, even for one-time, read-only access.

Related Patterns: When a user is required to register, an *Immortal Account* will prevent the later cancellation of the account. Forced accounts can come with *Bad Defaults*.

Psychological Aspects: As the user’s original goal is prevented by the necessary registration, account creation often happens as part of an automatic behavior for achieving that goal. This gives the user an *instant gratification*, and critical and deliberative thoughts are inhibited.

Strategies: MAXIMIZE

⁴ e.g., <http://10minutemail.com>

⁵ <http://bugmenot.com/>

5.4 Hidden Legalese Stipulations

Name/Aliases: *Hidden Legalese Stipulations*

Context: This pattern can be used by all systems which incorporate a document describing the terms and conditions of using the service.

Description: Terms and conditions are mandatory by law. Nevertheless, most users do not read them, since they are often long and written in a complicated legal jargon. This legal jargon is necessary to provide succinctness and clarity, but is not user-friendly.

The inability of the user to grasp the legal jargon puts him in a vulnerable state, since the policy is legally binding. If this vulnerability is exploited, the policy turns into an instance of a privacy dark pattern. Service providers can hide stipulations in the policies which target the privacy of the user. Often the user will not notice this, not reading the terms and conditions or being unable to understand their implications. Some service providers state that they will change their policies without further notice, preventing the user even further from learning what happens to his data.

Effect: Usage of this pattern leads to the service provider being able to hide his malicious deeds from the user without necessarily violating legal regulations.

Countermeasures: There are various proposals for easier communication of legal conditions.

One solution is to make the legal conditions machine-readable. This was the approach that P3P, the Platform for Privacy Preferences Project, followed. P3P is a standard by the W3C⁶ for a machine-readable rendering of privacy policies. The basic idea is that an XML-file specifying the privacy policy can be retrieved from any participating web pages. This policy can automatically be checked against the preferences of the user by the browser.

The Privacy Bird⁷, for example, was a tool which could show the P3P description as an icon, namely a bird. The color of the bird, i.e., red or green, signified if the policy of the site matched the users' preferences.

The drawback of this approach is, that the service provider needs to provide the machine-readable P3P description. A malicious service provider who wants to trick his users with hidden legal stipulations will of course not provide such a description. Since this countermeasure depends on the collaboration with the service provider it is not effective.

⁶ <https://www.w3.org/P3P/>

⁷ <http://www.privacybird.org/>

Another approach is the one followed by the Terms of Service; Didn't Read (TOSDR⁸) webpage. This is a community-driven repository of ratings of privacy policies. TOSDR is available as a browser add-on and shows the rating of the terms of service of the current web page as a small icon. When clicking on the icon one can see the positive and negative points of the terms of service in an easily understandable language.

Examples/Known Uses: In 2000, the then-popular instant messenger service ICQ introduced a "Terms Of Service — Acceptable Use Policy"⁹ which granted the service operators the copyright on all information posted by their users. Hidden in this legalese, the operators granted further rights of use "including, but not limited to, publishing the material or distributing it".

The British firm GameStation owns the souls of 7,500 online shoppers, thanks to an "immortal soul clause"¹⁰ in the terms and conditions. This April Fool's gag reveals the effectiveness of this pattern and shows that companies can hide everything in their online terms and conditions. Please note that McDonald et al. [36] calculated that reading the privacy policies you encounter in a year would take 76 work days.

Related Patterns: n/a

Psychological Aspects: Even if the user is motivated to read terms and conditions, *missing opportunity* to fully comprehend all details makes a *System 1-based processing* more probable.

Strategies: OBSCURE

5.5 Immortal Accounts

Name/Aliases: *Immortal Accounts*

Context: Many services require user accounts, either because they are necessary for service fulfilment, or because user accounts represent a benefit for the service.

Description: The service provider requires new users to sign up for accounts to use the service. Once users decide to stop using the service, they might want to delete their accounts and associated data. However, the service provider prevents the user from doing so by either—unnecessarily complicating the account deletion experience, or by not providing any account deletion option

⁸ <https://tosdr.org/>

⁹ <https://web.archive.org/web/20001204110500/http://www.icq.com/legal/policy.html>

¹⁰ <http://www.foxnews.com/tech/2010/04/15/online-shoppers-unknowingly-sold-souls.html>

at all. Additionally, the service provider might trick the user in the deletion process by pretending to delete the entire account, while still retaining (some) account data.

Effect: When the user interface makes the account deletion options hard to access, the barrier to delete the account is increased. If the users are required to call the customer support, the process is even more cumbersome. Both of these deliberately inconvenient user experiences may cause the user to reconsider the actual deletion decision. A deletion process where the service provider claims to remove the account, but instead just flags the user records as deleted while still keeping the data gives the user a false feeling of deletion.

Countermeasures: Online resources such as justdelete.me¹¹ or accountkiller.com¹² curate a list of service providers and their policies towards account removal. They provide step-by-step tutorials for users how to delete an account at those providers. If the service to be used is known for a non-delete policy but requires a user account, the usage of a throwaway account with incorrect data should be considered.

Examples/Known Uses: As of February 2016, the community-curated data set of justdelete.me lists 474 services. 75 services thereof do not provide the possibility to delete the account at all and 100 services require contacting the customer support. From the remaining 299 services listed, another 31 services have a non-trivial deletion process that requires additional steps.

Related Patterns: The creation of accounts can be required due to *Forced Registration*.

Psychological Aspects: When the service provider renders the user experience for account deletion deliberately painful, users might struggle in the process. If the user wants to delete the account, but fails to do so, *cognitive dissonance* may emerge. As a result, the user could then reduce the inconsistent mental state by reconsidering their original intent and deciding not to delete the account.

Strategies: DENY, OBSCURE

5.6 Address Book Leeching

Name/Aliases: *Address Book Leeching*

Context A service provider offers users to upload or import their address books to connect with known contacts on that service.

Description: When the user imports the list, the service executes a lookup against its own database. It then provides suggestions for connections to the user. However, the service provider stores the list of all contacts as internal data records for further processing—including purposes that have not been initially declared.

Effect: Using an import feature may lead to exposing unwanted information, specifically the contents of personal address books to third parties. A potential usage of such information is the dispatch of invitations or other advertisements, at worst even in the name of the original uploader without consent. Service provider may misuse such data for profiling and tracking individuals that do not yet possess a user account.

Countermeasures: If it is unknown or unclear how a service provider is handling and processing imported contact lists, such a feature should be avoided. Many mobile and desktop operating systems allow users to deny applications access to address book data. Users should routinely click on deny unless it is definitely required or in their interest to share those data.

Examples/Known Uses: In 2008, the social book cataloging website goodreads.com attracted negative attention for unsolicited invite emails based on the address book import feature. The experiences of customers and reactions of the service providers are still available on a customer support page¹³. Based on a misleading upload form design, users thought they would only provide contacts for matching against goodreads' user base. Instead, goodreads sent invite emails to persons which had mail addresses not yet registered at goodreads, thereby referring to the user who provided the address.

Related Patterns: This pattern is a potential source of information for *Shadow User Profiles*.

Psychological Aspects: Trading personal information for instant connections to friends or known contacts is motivated by the *need to belong*.

Strategies: MAXIMIZE, PRESERVE

5.7 Shadow User Profiles

Name/Aliases: *Shadow User Profiles*

Context: A service provider tracks personal information about individuals.

Description: While registered users have deliberately

¹¹ <http://justdelete.me/>

¹² <http://www.accountkiller.com/>

¹³ https://getsatisfaction.com/goodreads/topics/why_did_goodreads_trick_me_into_spamming_my_entire_address_book

opted in for a user account and an associated profile, the service provider may collect information and keep records about individuals that do not use the service. For instance, in a social network, the social graph can be supplemented with persons that are not members of the network, but are known to the network based on data from members (e.g., imported address books, content metadata, or mentions). Such non-members enrich the graph and improve the quality of algorithms such as contact suggestions.

Effect: The service provider stores and processes information on individuals without their knowledge or consent. The affected individuals are not aware of personal data records they have accidentally created or that have been provided by third parties.

Countermeasures: While it is possible to minimize the own data trail, the accidental release of personal data through third parties cannot always be prevented.

Examples/Known Uses: The basic mechanism of shadow user profiles fuels the entire online advertisement industry. Although not verifiable, social networks may store informations of non-users. This notion is based on the experiences of newly registered users of social networks who received accurate friendship suggestions without having ever interacted with these persons on the social network before.

Related Patterns: *Address Book Leeching* is a potential source of information for this pattern.

Psychological Aspects: Given the fact that this pattern operates without any knowledge of the affected users, it is not targeting any psychological aspects.

Strategies: MAXIMIZE, PRESERVE, CENTRALIZE

6 Conclusions

In this paper, we introduce the concepts of privacy dark strategies and privacy dark patterns. Both are based on the idea that actors intentionally manipulate people to provide their personal data for collection, storage, and processing against their original intent and interest.

Documenting such strategies and patterns is a vital first step towards a better recognition of such activities, e.g., in the Internet or in mobile apps. Our eight privacy dark strategies MAXIMIZE, PUBLISH, CENTRALIZE, PRESERVE, OBSCURE, DENY, VIOLATE, and FAKE provide a coarse categorization for the subsequent patterns. Privacy dark patterns are documented using a uniform template. Beyond a mere description of the pattern, the

template contains countermeasures and a psychological viewpoint that explains why the pattern is effective.

We extensively discussed psychological aspects in Section 4. Understanding those psychological mechanisms triggered by privacy dark patterns is of crucial importance as it will allow affected users to take appropriate countermeasures.

Based on our privacy dark pattern framework and the extensive discussion of the related concepts, we briefly presented seven of such patterns including some concrete examples. These patterns and more are available in an extended form via an online privacy dark pattern portal dark.privacypatterns.eu. We have set up this portal for the community to study and discuss existing patterns and contribute new ones.

Acknowledgements

The authors like to thank the anonymous reviewers for their valuable comments and suggestions to improve the quality of the paper. They are also grateful to Yllka Thaqi and Florian Oberlies for insightful remarks and fruitful discussions.

References

- [1] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz, "The Web never forgets: Persistent tracking mechanisms in the wild," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 674–689.
- [2] A. Acquisti, "Privacy in electronic commerce and the economics of immediate gratification," in *Proceedings of the 5th ACM conference on Electronic commerce*. ACM, 2004, pp. 21–29.
- [3] —, "Nudging privacy: The behavioral economics of personal information." *IEEE Security & Privacy*, vol. 7, no. 6, pp. 82–85, 2009.
- [4] A. Acquisti, L. K. John, and G. Loewenstein, "The impact of relative standards on the propensity to disclose," *Journal of Marketing Research*, vol. 49, no. 2, pp. 160–174, 2012.
- [5] C. Alexander, S. Ishikawa, and M. Silverstein, *A Pattern Language: Towns, Buildings, Construction (Center for Environmental Structure Series)*. Oxford University Press, 1977.
- [6] H. Almuhammedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal, "Your location has been shared 5,398 times!: A field study on mobile app privacy nudging," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '15. New York, NY, USA: ACM, 2015, pp. 787–796.

- [7] Y. Amichai-Hamburger and E. Ben-Artzi, "Loneliness and internet use," *Computers in Human Behavior*, vol. 19, no. 1, pp. 71–80, 2003.
- [8] C. M. Angst and R. Agarwal, "Adoption of electronic health records in the presence of privacy concerns: The elaboration likelihood model and individual persuasion," *MIS quarterly*, vol. 33, no. 2, pp. 339–370, 2009.
- [9] R. F. Baumeister and M. R. Leary, "The need to belong: desire for interpersonal attachments as a fundamental human motivation," *Psychological Bulletin*, vol. 117, no. 3, pp. 497–529, 1995.
- [10] K. Beck and W. Cunningham, "Using pattern languages for object oriented programs," in *Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 1987.
- [11] H. Brignull, "Dark Patterns: fighting user deception worldwide," <http://darkpatterns.org/>, accessed: 2016-01-24.
- [12] A. Buchenscheit, B. Könings, A. Neubert, F. Schaub, M. Schneider, and F. Kargl, "Privacy implications of presence sharing in mobile messaging applications," in *Proceedings of the 13th International Conference on Mobile and Ubiquitous Multimedia*. ACM, 2014, pp. 20–21.
- [13] R. Cialdini, *Influence: the psychology of persuasion*. New York: Morrow, 1993.
- [14] N. Doty and M. Gupta, "Privacy Design Patterns and Anti-Patterns," in *Trustbusters Workshop at the Symposium on Usable Privacy and Security*, 2013.
- [15] N. B. Ellison, C. Steinfield, and C. Lampe, "The benefits of facebook "friends": social capital and college students' use of online social network sites," *Journal of Computer-Mediated Communication*, vol. 12, no. 4, pp. 1143–1168, 2007.
- [16] R. H. Fazio, "Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework," *Advances in Experimental Social Psychology*, vol. 23, pp. 75–109, 1990.
- [17] L. Festinger, *A theory of cognitive dissonance*. Stanford university press, 1962, vol. 2.
- [18] M. Fowler, *Patterns of Enterprise Application Architecture*. Boston: Addison-Wesley Professional, 2003.
- [19] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design patterns: elements of reusable object-oriented software*. Pearson Education, 1994.
- [20] H. Gangadharbatla, "Facebook me: Collective self-esteem, need to belong, and internet self-efficacy as predictors of the igeneration's attitudes toward social networking sites," *Journal of interactive advertising*, vol. 8, no. 2, pp. 5–15, 2008.
- [21] S. Gürses, C. Troncoso, and C. Diaz, "Engineering privacy by design," *Computers, Privacy & Data Protection*, vol. 14, 2011.
- [22] M. Hafiz, "A collection of privacy design patterns," in *Proceedings of the 2006 conference on Pattern languages of programs*. ACM, 2006, p. 7.
- [23] E. T. Higgins, *Beyond pleasure and pain: How motivation works*. Oxford University Press, 2011.
- [24] J.-H. Hoepman, "Privacy Design Strategies," *CoRR*, vol. abs/1210.6621, 2012.
- [25] G. Hohpe and B. Woolf, *Enterprise Integration Patterns - Designing, Building, and Deploying Messaging Solutions*, 1st ed. Boston: Addison-Wesley Professional, 2004.
- [26] D. J. Hughes, M. Rowe, M. Batey, and A. Lee, "A tale of two sites: Twitter vs. Facebook and the personality predictors of social media usage," *Computers in Human Behavior*, vol. 28, no. 2, pp. 561–569, 2012.
- [27] P. Hustinx, "Privacy by design: delivering the promises," *Identity in the Information Society*, vol. 3, no. 2, pp. 253–255, 2010.
- [28] T. Jones, "Facebook's "evil interfaces"," <https://www.eff.org/de/deeplinks/2010/04/facebooks-evil-interfaces>, accessed: 2016-02-25.
- [29] D. Kahneman, *Thinking, fast and slow*. Macmillan, 2011.
- [30] B. P. Knijnenburg and A. Kobsa, "Increasing sharing tendency without reducing satisfaction: Finding the best privacy-settings user interface for social networks," in *Proceedings of the International Conference on Information Systems - Building a Better World through Information Systems, ICIS 2014, Auckland, New Zealand, December 14-17, 2014*, 2014.
- [31] B. P. Knijnenburg, A. Kobsa, and H. Jin, "Counteracting the negative effect of form auto-completion on the privacy calculus," in *Thirty Fourth International Conference on Information Systems, Milan*, 2013.
- [32] A. Kobsa, H. Cho, and B. P. Knijnenburg, "The effect of personalization provider characteristics on privacy attitudes and behaviors: An elaboration likelihood model approach," *Journal of the Association for Information Science and Technology*, 2016, in press.
- [33] D. Laibson, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, vol. 112, no. 2, pp. 443–478, 1997.
- [34] P. B. Lowry, G. Moody, A. Vance, M. Jensen, J. Jenkins, and T. Wells, "Using an elaboration likelihood approach to better understand the persuasiveness of website privacy assurance cues for online consumers," *Journal of the American Society for Information Science and Technology*, vol. 63, no. 4, pp. 755–776, 2012.
- [35] E. Luger, S. Moran, and T. Rodden, "Consent for all: revealing the hidden complexity of terms and conditions," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2013, pp. 2687–2696.
- [36] A. M. McDonald and L. F. Cranor, "Cost of reading privacy policies, the," *ISJLP*, vol. 4, p. 543, 2008.
- [37] A. Nadkarni and S. G. Hofmann, "Why do people use Facebook?" *Personality and Individual Differences*, vol. 52, no. 3, pp. 243–249, 2012.
- [38] N. Notario, A. Crespo, Y.-S. Martín, J. M. Del Alamo, D. Le Métayer, T. Antignac, A. Kung, I. Kroener, and D. Wright, "PRIPARE: Integrating Privacy Best Practices into a Privacy Engineering Methodology," in *Security and Privacy Workshops (SPW), 2015 IEEE*. IEEE, 2015, pp. 151–158.
- [39] R. E. Petty and J. T. Cacioppo, *The elaboration likelihood model of persuasion*. Springer, 1986.
- [40] S. Romanosky, A. Acquisti, J. Hong, L. F. Cranor, and B. Friedman, "Privacy patterns for online interactions," in *Proceedings of the 2006 conference on Pattern languages of programs*. ACM, 2006, p. 12.
- [41] M. Schumacher, "Security patterns and security standards," in *EuroPLoP*, 2002, pp. 289–300.

- [42] T. Schümmer, "The public privacy-patterns for filtering personal information in collaborative systems," in *CHI2004: Proceedings of the Conference on Human Factors in Computing Systems*, 2004.
- [43] K. E. Stanovich and R. F. West, "Advancing the rationality debate," *Behavioral and Brain Sciences*, vol. 23, no. 05, pp. 701–717, 2000.
- [44] F. Strack and R. Deutsch, "Reflective and impulsive determinants of social behavior," *Personality and Social Psychology Review*, vol. 8, no. 3, pp. 220–247, 2004.
- [45] R. Thaler, *Nudge : improving decisions about health, wealth, and happiness*. New York: Penguin Books, 2009.
- [46] J. Tidwell, *Designing Interfaces*. Sebastopol: "O'Reilly Media, Inc.", 2010.
- [47] A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," *science*, vol. 185, no. 4157, pp. 1124–1131, 1974.
- [48] J. van Rest, D. Boonstra, M. Everts, M. van Rijn, and R. van Paassen, *Designing privacy-by-design*. Springer, 2014, pp. 55–72.
- [49] K. D. Williams, C. K. Cheung, and W. Choi, "Cyberostracism: effects of being ignored over the internet." *Journal of Personality and Social Psychology*, vol. 79, no. 5, pp. 748–762, 2000.