Boosting with Friction

The Effects of Design Friction on Deliberation in the Context of Privacy Decisions

Lotte van Elteren
s4487478

Supervisors: Hanna Schraffenberger and Franc Grootjen

A thesis presented for the Bachelor's degree of Artificial Intelligence

Radboud University
Nijmegen, the Netherlands

Faculty of Social Sciences
Radboud University
The Netherlands
July 2020

Abstract

People face privacy decisions every day. Digital consent forms aim to acquire consent by a user, but often fail in this aspect by nudging a user towards a certain direction. The user is pushed to make a non-deliberate choice, often resulting in a privacy-unfriendly decision. While quite a lot is known about design patterns that can foster this type of behavior, not much is known about design patterns that can boost decision making, to empower a user to make a more deliberate decision. This study investigates the use of a design pattern called friction to bring forth this deliberation. We looked at the privacy decision, deliberation and user experience of our participants when interacting with a friction pattern compared to a neutral design. Even though no significant results were found for these three factors, we do have ground for further research on the matter.

1 Introduction

When designing a digital interface, the success of a design is often measured by the number of users performing an action like creating, sharing, or purchasing. It is therefore vital that a user can navigate through a design with ease, meaning that the design does not obstruct them from achieving their goal. Sometimes however, a digital design can capitalize on a user by having maximum control over a user's actions. A 'dark pattern' in UX design is a purposefully designed interface that tries to convince a customer that it is in their best interest to do something that benefits the designer more than them (Nodder, 2013). A common example of this is the way Privacy Consent Statements, such as consent queries for placing cookies in web browsers, nudge the user to make an irrational or uninformed decision (Forbrukerrådet, 2018). While quite a lot is known about such interfaces, which nudge a user to make an irrational decision without much deliberation, not much is known about digital designs that have the opposite effect. With this study we aim to fill this gap in research, and investigate whether a design pattern called 'friction' can empower a user to make a more deliberate decision in the context of privacy. Graßl (2019) states that recent changes in the GDPR raise demand for a design pattern that can guarantee the informed consent of a user. This study will investigate why many commonly used design patterns based on nudging, by nature, fail to acquire informed consent. We will first explain what Dark and Bright Patterns are. Then we will explain the foundations of meaningful control and consent, and why commonly used design patterns fail in acquiring these. We then will explain a new policy called Boosting, and a way to implement this policy in a design, called friction. We will then investigate whether this design pattern boosts a user's deliberation.

1.1 Dark Patterns and Bright Patterns

Designing interfaces such as consent statements often involve the use of a design pattern. Design patterns are repeatable solutions to recurring design problems

(Norman & Draper, 1986). Many design patterns are used to nudge a user in a certain direction. For example, bright colors, arrows, or icons are used to indicate to the user how to navigate through the design. Some of these patterns are specifically used to nudge a user to make a certain decision that does not benefit the user. These patterns are known as Dark Patterns. Dark Patterns are design patterns that are specifically designed to nudge a user towards a decision that is not in their best interest, but that has value for the maker of the design. They are for instance commonly found in cookie statements, in order to nudge the user to accept cookies used for advertisement. In the context of a dark pattern, the nudging is always towards a privacy-unfriendly option (Graßl, 2019). A recent study may show the potential use of the opposite of these patterns, which have a more virtuous purpose.[1] These patterns, named Bright Patterns by Graßl, are less commonly seen, and they are used to nudge the user to make a rational, often privacy-friendly, decision (Graßl, 2019).

There are various popular design patterns. For privacy consent statements, these patterns typically include at least two options for the user to choose: agreeing or disagreeing to the matter at hand. Commonly used patterns that can be used for both bright and dark purposes include Aesthetic Manipulation, Obstruction, and Default (Graßl, 2019; Fansher, Chivukula, & Gray, 2018). Aesthetic Manipulation can be done by highlighting one of two options in a privacy consent statement: giving one option visual or interactive precedence over others. Oftentimes, one of the buttons is colorful and attractive, while the other is simply grey. Obstruction makes choosing one option over the other more difficult than it needs to be. This is often done by not presenting the opt-out option next to the opt-in option in cookie-consent statements, or hiding the opt-out option in a menu while the opt-in option can be chosen in one click. The Bright version of this pattern would do the opposite(Graßl, 2019). The Default pattern pre-selects one of the available options, which means an additional action is required to choose the other option. The default option in a Dark pattern will then be the privacy-unfriendly option (Graßl, 2019).

[1] Research will be published on this at the end of July 2020, in the work "Dark and Bright Patterns in Consent Requests" by Paul Graßl

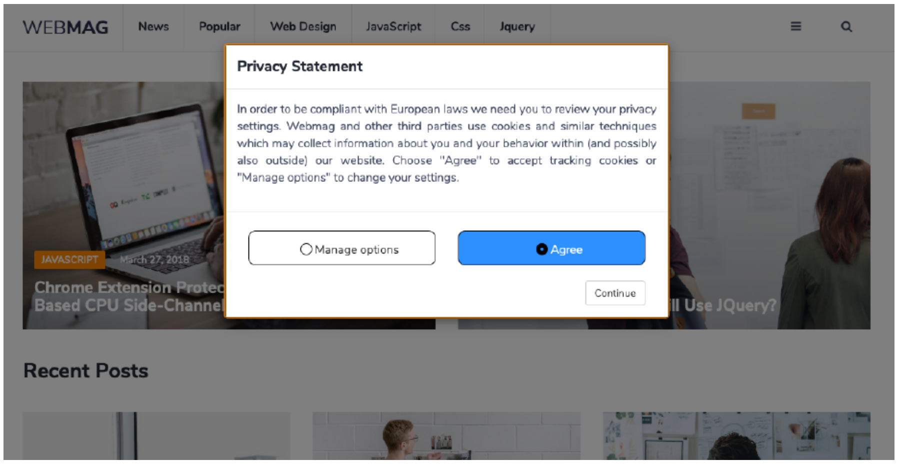

Figure 1: Example consent statement featuring all three dark patterns: default, aesthetic manipulation and obstruction. Website: Webmag (Graßl, 2019).

Typically, these patterns are used as Dark patterns, but it is currently being researched whether these patterns can be used for both Bright and Dark purposes.[2] This could be possible because Bright and Dark patterns have the same purpose: to nudge the user into a certain direction. However, they both rely on the user making a quick, non-deliberate decision. The user is urged to choose one of either options, without putting too much thought behind the choice. The process of deliberation, and therefore purposefully choosing one option over the other, has already been finished by the designer. Patterns such as Obstruction, Default or Aesthetic Manipulation will speed up the choosing process, in order to obstruct the user in the process of getting informed and to reflect upon privacy issues (Terpstra, Schouten, Rooij, & Leenes, 2019). Thus, when applied to privacy statements, they fail to acquire informed consent (Graßl, 2019).
Judging from existing research, a lot is known about how to design for a non-deliberate decision, even though the effects of these patterns on behavior are still studied. Not much is known, however, about design patterns that can foster a different type of behavior; that can cause a more deliberate decision which could support informed choices. In order to find out how to design such an interface, we need to understand the basics of meaningful control and informed consent.

1.2 Meaningful control and informed consent

The need for a pattern that can aid in acquiring informed consent becomes clear in our everyday lives. People face consent statements about privacy almost ev-

[2]In the as of yet unpublished work “Dark and Bright Patterns in Consent Requests” by Paul Graßl

ery day. Take cookie statements for example: banners asking for consent for the placement of cookies are shown on almost every website. These cookie statements are instances of privacy statements. Since May 2018, changes have been made in European law about these privacy statements, which can be found in a new law called the GDPR (*General Data Protection Regulation GDPR*, 2018). One aim of the GDPR is to protect the user by making sure user give informed consent. This informed consent will then give the user meaningful control over their personal data, so we must first investigate what meaningful control and informed consent mean, and why commonly used patterns are insufficient in acquiring it.

In order to define what it means to have meaningful control over personal data, a philosophical analysis of the privacy topic is needed. Control in the context of consent decisions revolves primarily around the topics access, notice, and consent (Solove, 2013). What is commonly understood as access and notice is that a user has explicit access to which specific information is shared regarding their own privacy. When applied to digital applications and services, having access and acquiring notice means that a user must not only know what the application they are using does with the data it gathers about its user, but also which data is gathered, and how it is maintained. This must be done in a way that is comprehensible to a user: they must understand what it means that this data is gathered. Moreover, the user must be given notice of any changes (*General Data Protection Regulation GDPR*, 2018, article 7, 12, 19).
Also, important in this discussion is the meaning of consent. We base our research on the notion that consent can only be meaningful under certain circumstances (Wertheimer, 2000). Firstly: consent should be understood as a performative or action; not as a subjective phenomenon (Wertheimer, 2000). This means that consent requires an action: it needs to be given in order for it to exist. Secondly: when a user is not informed enough about the issues they are giving consent to, the consent can not be considered meaningful. This means that the consent is only valid when the aforementioned access and notice are present. Thirdly: Consent relies on the giver to be in control of their decision, meaning that the consent is given freely, willingly, and at any time reversible. This includes that the option to make a decision must be provided: there has to be an option not to proceed (Soeder, 2019).

1.3 Boosting

A truly virtuous pattern should aim to give the user meaningful control and informed consent in the manner as described in the previous section. Although consent statements, such as cookie statements, are often not designed with a virtuous purpose, there are situations in which it is the intention of a party to ask for an informed and deliberate decision of a user. For example, when consent needs to be given by a participant of a research experiment, it is of great importance that the participant gives consent in that is performative and informed, and that the participant is in full control, as mentioned in section

1.2. A nudging policy fails in these aspects, because the decision made with a nudging policy requires little deliberation.
Instead of using a nudging policy, Grüne-Yanoff and Hertwig suggest the use of a boosting policy. They write: "The common denominator behind boost policies is the goal of empowering people by expanding (boosting) their competences and thus helping them to reach their objectives (without making undue assumptions about what those objectives are)". To put this policy into practice, it may be useful to incorporate it into the design of the consent statement itself. Boosting is a policy that can be implemented in different ways. Terpstra et al. (2019) describe that any individual using an IT system needs to:

- Understand what is at stake when they interact with digital technology
- Have the ability to reflect on the consequences of the decisions they make about their own privacy
- Have the ability to meaningfully express their privacy preferences.

Terpstra et al. (2019) introduces 'friction', a central design concept that can help to uphold these three values, and thereby gives an alternative approach to consent, by letting the user deliberate their choice more than with a nudging policy.

1.4 Friction

Friction in user experience design is anything that inhibits the user from intuitively and painlessly achieving their goals within a digital interface (Young, n.d.). It is often seen as a 'bad design' characteristic, as the goal of a good design often entails an easy and smooth process from start to goal completion (Cox, Gould, Cecchinato, Iacovides, & Renfree, 2016). An example of a bad use of friction in privacy consent statements is when a user can only agree to the terms and conditions of software after scrolling all the way down to the end of the document. This will often not cause the user to read the agreement, rather it causes frustration (Terpstra et al., 2019). Friction can also be used in a user-friendly way, for example when an 'are you certain'-dialogue pops up when a user tries to close an unsaved document. Such warnings, called microboundaries, can have a positive effect on a users experience (Cox et al., 2016).

Friction has not received much scientific attention yet, although its practical use is discussed by UX designers in blogs on the internet. Rekhi mentions that friction designs can be divided into three different categories: emotional friction, cognitive friction and interaction friction. Interaction friction refers to friction a user experiences when interacting with a product's interface, covers all aspects of the digital interface that can hinder users from accomplishing their goal (Rekhi, 2017). An example of this is a web design where the most important information is difficult to find, such as clinking through multiple pages before finding the phone number of an emergency service. Cognitive friction refers to the amount of mental effort needed to accomplish a goal. When cognitive load is high when a user performs a task, it means that there is a high amount of cognitive friction (Rekhi, 2017). An example of this is when the expectation of how an interface works does not match with its actual outcome of an interaction. This is sometimes used in video games, if the left and right arrows to move do not match the respective directions of movement of a character. Furthermore, there is emotional friction, which refers to emotions a user feels that prevent them from accomplishing their goal. Rekhi mentions the dating app Tinder as an example that reduces this type of friction, because the application reduced the intimidation that users experienced during online dating.

A real world example of the usefulness of these types of friction which is frequently observed is the design and user experience of video games (Malone, 1981). In a good video game, the user is challenged by the game just enough that there is a good ratio between failing, causing a breakdown, and succeeding, causing a breakthrough (Iacovides, Aczel, Scanlon, & Woods, 2011). A user has a breakdown when they face a task that, at that particular moment, is too challenging for them to complete. The user is then stuck in the game, and has to try different approaches, which causes the gamer to learn new strategies (Iacovides, Cox, & Knoll, 2014). A breakthrough happens when the gamer has a moment of insight, which helps to increase involvement (Iacovides et al., 2014). Important to note here is that there is a difference between breakdowns caused by poor interface design, and breakdowns that are an essential part of the user experience, creating user involvement and immersion (Iacovides et al., 2011), (Cox et al., 2016). These findings can help us understand how to design for deliberate decision making. Similar to video games, we can try to slow down the user in their process of going through the design. The aim here is that the user is, through friction, slowed down enough such that they will gain new insights into the consent they are giving, hopefully resulting in a breakthrough: a rational, deliberate decision.

Cox et al. show that friction in the form of microboundaries can also be used for what we call decision boosting (empowering the user to make a more deliberate decision). A microboundary is an intervention that provides a small obstacle prior to an interaction that prevents us rushing from one context to another, which like the breakdowns and breakthroughs seen in video games may cause the user to deliberate the decision they make. An example from real life is that keeping a food diary is scientifically proven to aid in weight loss, because the extra step before eating something causes a more mindful decision about the food (Cox et al., 2016). However, Terpstra et al. (2019) mentions that it is only beneficial to a user when the purpose of the friction is clear. Applied to our cause, to support reflective thinking, an application should provide the means to learn about the topics the user is deciding upon. Furthermore, this information should be provided at the right time. Think for example about an app that tracks data from your navigation device. After a few trips, the application can make a solid guess where the user lives, and should then provide the user with a message to inform them about this, and allow them to change their privacy settings (Terpstra et al., 2019). When presented at the correct time, friction can not only cause reflective thinking, but it may also introduce alternate views and interpretations (Sengers, Boehner, David, & Kaye, 2005). In the context of privacy decisions, such considerations can cause frustration, so a means of change should be presented along with the information (Sengers et al., 2005). For example: a link to the privacy settings can be provided along with the consent statement. The user should be allowed to have meaningful control over their privacy. Not being able to act upon the newly found beliefs depreciates the decision that has been made (Sengers et al., 2005).

1.5 Deliberative Decision Making

Adding small moments of friction to a design can help in the decision making process (Cox et al., 2016). Friction can aid our cause of more deliberative decision making. Research suggests that there are two systems for decision making: the fast type, known as System 1, and the slow type, known as System 2 (Kahneman, 2015). System 1 is used for mindless, fast decision making, which is what we want to avoid when important decisions about privacy need to be made. Friction design can be used to disrupt thoughtless interactions, and prompt mindful and conscious decision making (Cox et al., 2016). These mindful and conscious decisions can be categorized as decisions made through System 2. "Various factors determine whether a System 1 or a System 2 process is used, such as the task at hand or the context. People generally complete day-to-day tasks on 'autopilot' (driven by System 1), as they automatically react to their environment. However, System 2 processes are used when a situation requires focused attention. One way to influence behaviour is to facilitate the transition from System 1 to System 2 to leverage more deliberate conscious processes" (Cox et al., 2016). If we want to cause users to do more conscious, mindful decision-making about their privacy, our task is thus to use friction in such a way to facilitate the use of System 2 when decisions about privacy are made.

1.6 Experiment

Based on the discussed literature, we expect that friction design is a suitable method to achieve deliberative thinking regarding privacy decisions. However, empirical evidence for this is largely missing, and the question of how to best include friction in a user interface remains. The current study addresses the matters by running an online experiment to investigate how a design that includes friction influences a user's deliberative process, and whether it indeed causes a more privacy friendly decision. Rational choice theory would predict that thinking with System 2 causes a privacy friendly decision (Smith, Dinev, & Xu, 2011) if this decision is of greater utility than the alternative. "Choosing a privacy-friendly option over a privacy-unfriendly option can be considered a rational choice if the privacy-friendly option bears greater utility in terms of costs and benefits than the privacy unfriendly option" (Graßl, 2019). Therefore, choosing a privacy-unfriendly option over a privacy-friendly option suggests that the decision is made through System 1 rather than System 2, and may thus not be deliberate but automatic. As part of this overall goal, we explore various approaches that introduce friction in privacy consent decisions.

There are however other factors to consider which may influence a person's decision on privacy matters. Previous research suggests that a person's individual privacy concerns influence a user's privacy decisions in general, and that people with higher privacy concerns are less likely to be default setting in a consent statement than people with low privacy concerns (Graßl, 2019; Awad & Krishnan, 2006; Lai & Hui, 2006; Malhotra, Kim, & Agarwal, 2004). Familiarity with design patterns, and knowledge about how or why one would make a certain design choice might affect the outcome of the experiment as well. It is, after all, the effects of a design pattern that we are investigating. Lastly, we need to consider the effects of the friction design on the user experience. Since friction is mostly seen as something to be avoided when designing an application, we need to address the effects of friction on the user's experience.

We hypothesized that in a consent statement situation with two choice options (privacy-friendly vs. privacy unfriendly), where the privacy-friendly option is rationally superior,

> Hypothesis 1: participants will be more likely to opt out of sharing personal data when an inappropriate request is made, if the participant is presented with friction in the design pattern of the consent statement (compared to a design that does not include friction).
>
> Hypothesis 2: participants will be more likely to make a deliberate choice about sharing their personal data when an inappropriate request is made, if the participant is presented with friction in the design pattern of the consent statement (compared to a design that does not include friction).
>
> Hypothesis 3: participants will have a less pleasant user experience when presented with friction in the design pattern of a consent statement (compared to a design that does not include friction).

Because little is known about the effects of friction in requests for sharing personal data, this pioneer study focused on the main effects of friction in a particular situation, rather than possible (and more speculative) interaction or moderation effects, in order to create a solid basis for further investigation.

2 Method

The following sections describe the experiment we performed to adress the hypotheses we made. Before starting the experiment the study was registered at

DCC, received for ethical approval.

2.1 Procedure and design

The experiment followed a between-subject design, where subjects were asked to interact with one of two different mock versions of an consent form: version A in which friction is integrated in the design, and version B, the neutral design. We wanted to test for general effects of friction on user's deliberation and behavior, since not much is known about the effects of friction in consent statements. Therefore, the friction in the neutral design B is minimal. The users were instructed to open a mock application and report on their first impressions of the visual design of the app (this cover story was used to create a realistic setting for the presentation of the consent statement). The application requires the user to confirm that their age is over 18 years, and it requires them to share their BSN (social security number). For privacy reasons only mock personal data was used. All participants were divided randomly in either the control group or the group that interacted with the friction design. First, the participants were asked to sign a consent form to participate in the experiment. The participant was then shown the mock application for a brief moment. Lastly, they were asked to fill in a questionnaire about their experience. In this questionnaire the participants were first debriefed about the cover story and the true purpose of the experiment, after which they were asked questions about their decisions, deliberative process and their user experience. Additionally, they were asked some control questions about their experience with Human Computer interaction and their individual privacy concerns.
Our dependent variables in this experiment are the privacy decision that is made, the amount of deliberation upon this decision, and the user experience. Our independent variable is the friction that is incorporated into the design.

2.2 Materials

To run the experiment online, we set up a mock application with Figma, a tool in which interactive mock applications can be created and tested (Figma, 2020). The consent statement before the experiment, and the questionnaire after were made with LimeSurvey, hosted on the Radboud University: Social Sciences server (LimeSurvey Development Team, 2012). The application will be discussed in detail in section 2.4.

2.3 Measures

For each participant we assessed their deliberative process, user experience, and several control factors with a questionnaire.

The questionnaire was presented to the participants after they had completed the interaction with the mock application. The questionnaire included a short introduction, explanation about the experiment, instructions for completion, and a question about previous experiences with Human Computer Interaction.

2.3.1 Privacy decision

The first questions that were asked were about whether the participant had chosen to share their personal information with the mock liquor store application. After this, a series of questions was asked about their deliberation of that decision, and their user experience of the application.

2.3.2 Deliberation

The questions that were asked are from the DelibeRATE tool (Gillies, Elwyn, & Cook, 2014). The DelibeRATE tool is a short, ten-item tool which requires each question to be rated on a three-level Likert scale with range Yes/Unsure/No. The questions were changed to fit the experiment. An example of a change is the question: "I am aware of the advantages of participating in the trial or not", was changed to: "I was aware of the advantages of sharing my personal data with Drankstore b.v. or not". The questions that were asked can be found in appendix A.

2.3.3 User experience

The following questions were asked about the user experience of the application. They followed guidelines from the Short Version User Experience Questionnaire (UEQ), as described by (Rauschenberger, Schrepp, Pérez Cota, Olschner, & Thomaschewski, 2013; Schrepp, Hinderks, & Thomaschewski, 2017), and (Laugwitz, Held, & Schrepp, 2008). The UEQ-S consists of 8 questions, which test for the overall attractiveness of the design. The questions fall in the categories Pragmatic, with subcategories Perspicuity, Efficiency and Dependability or Hedonic, with subcategories Stimulation and Novelty. Participants were asked to rate their experience on a scale of 1-7 again. The questions that were asked can be found in appendix B.

2.4 The IRMA application

To test whether the user's deliberation about sharing their personal information is boosted by using friction and how this affected their user experience, we applied friction to the design of a real application. In many contexts, boosting when privacy is concerned is not desired by the designer or publisher of the app, as they may want to cause the user to share as much data as possible, because this may be of interest to them. For example, with companies like Facebook, getting the user to share their information is part of the revenue model. Even companies who have another way of making revenue can profit from a user's privacy unfriendly decision. Placing cookies to track interest, with the goal to make personalised advertisements, is common practice.

However, not everyone is interested in a user's data. One app that is specifically designed to help people to share as little information as possible is the IRMA application (Privacy by Design Foundation, 2020). This application stores personal data locally on the user's phone. The personal data is represented in the form of attributes: relevant characteristics of the user, like their age being over 18, or their BSN. These attributes can be shown to verify for example a user's age, marital status or BSN to other parties, without revealing more personal data than strictly necessary. An example is the use of IRMA to confirm that a user is above the age of 18 in an online store that sells alcoholic beverages. To test if friction actually increases deliberation, IRMA is very well suited, because it is specifically used to maintain personal, privacy sensitive data. Furthermore, IRMA is developing a functionality that can be used to safely reveal attributes to third parties, but does yet induce the user to think about whether the request from the third party is fair. Applying friction to the IRMA application when a user confirms the sharing of their personal information may help to improve this.

A mock version of the IRMA app was used for the purpose of this experiment, in combination with a mock version of an online non-existent liquor store named 'Drankstore b.v.'. We did not follow the exact same design of the IRMA application, but followed a similar but simplified concept. The liquor store was chosen to emulate an environment that had a clear purpose for an identification means such as IRMA. The age request for the user to be at least 18 years old made by the application is fair, as the legal age of purchasing alcoholic beverages in the Netherlands is 18 years. However, the web store has no purpose for asking a BSN. Rational behaviour would therefore mean that the user does not to share their BSN attribute. To see the effects of the friction design, we can apply it on the moment the user has to decide upon sharing their age and BSN attributes, hopefully supporting the user to make a rational, deliberate choice. There are however many different ways to implement friction in a positive way to this moment.

2.5 Pilot study

To find out which type of friction shows the most promising effects for the privacy decisions in the IRMA application, we performed a pilot study in which we applied five different forms of friction to a mock version of the IRMA app. The aim of this pilot was to find out which version of friction showed the most deliberate decision, and which version had the best user experience. The best design for our purposes is the one that shows both improved deliberation when compared to a design without friction, without causing an unpleasant user experience. The following five designs were tested in the pilot study:

1. The active warning. Egelman, Cranor, and Hong (2008) show in their study that providing an active warning that forces the user to notice the warning by interrupting their task showed an increase in the detection of phishing. For our study, the active warning is implemented by showing

a popup screen that warns the user about the possible danger of sharing their personal data.

Figure 2: Pilot version 1: Active warning

2. Colors as indicators. In order to inform a user about how sensitive a certain type of personal data is, we can use color coding. This type of design is often seen when a user has to choose a password, where a red color indicates a weak password, and a green color indicates a strong password. The effectiveness of color coding when a user has to choose a safe password has been researched, but the effects on decision making have not (Mathur, 2011). Even though this design does not stall the user as may be more obvious for a friction pattern such as shown in Version 1, it could have a similar effect on the user.

Figure 3: Pilot version 2: Visual friction in the form of colors as indicators.

3. Task Lockout. Gould, Cox, Brumby, and Wickersham (2016) and Back, Brumby, and Cox (2010) show that a brief pause between entering data and having the possibility to confirm and proceed made it more likely that users checked the information they were submitting. They also found that the longer the lockout lasted, the more participants switched tasks. The optimal lockout time depends on the task at hand, and the amount of information to be considered (O'Hara & Payne, 1999). We tested the task with lockout friction with a duration of 2 seconds. Only after the waiting time was over, as indicated by a bar on the screen, the user could click to accept or decline.

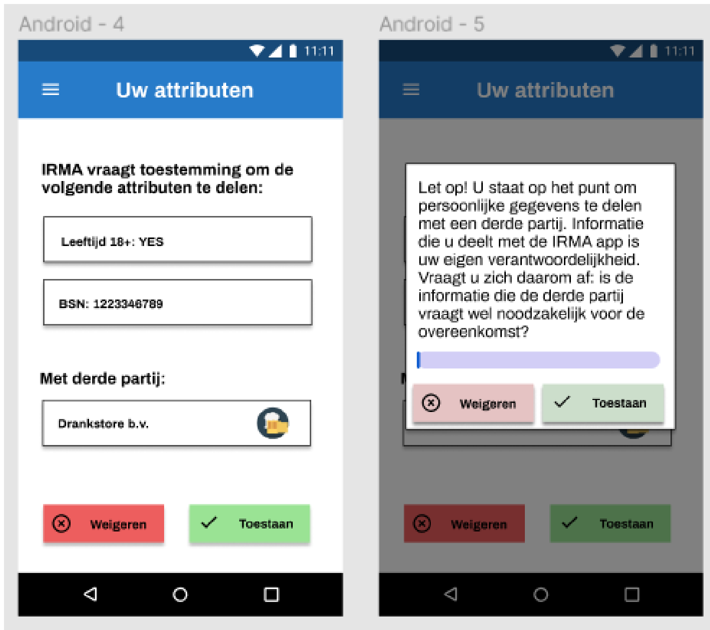

Figure 4: Pilot version 3: Task lockout.

4. Enter PIN. This design pattern is commonly used, for example in the banking applications. In order to see the current balance via their phones, a user has to enter a PIN. This does not only serve as a security measure, but it may also indicate to a user that they are dealing with a serious action.

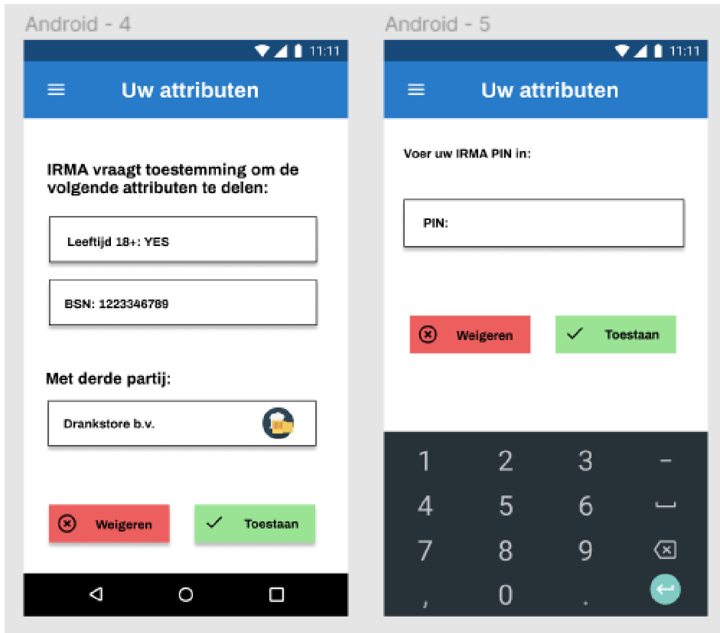

Figure 5: Pilot version 4: Entering IRMA Pin.

5. Manually enter BSN. This pattern asks the user to manually enter the word 'BSN' before proceeding.[3] A similar pattern can be found in the real world, for example on the website Github. Github provides hosting for software development version control. When a user wants to delete a repository that stores their software, the user has to manually type the name of that which they are deleting.

[3] In preparation of the pilot study we investigated a version where the user had to manually enter their BSN, but since this showed to be annoying to the experimenters, we opted for a design where only the letters 'BSN' needed to be entered instead.

Figure 6: Pilot version 5: Manually entering the word 'BSN'.

To test which of the five versions would be best, we performed a pilot study with 17 participants which were found by sending out emails. The participants were distributed equally over the five different designs. The uneven distribution of completed experiments can be explained by participants not completing the questionnaire, or participants deciding not to partake in the experiment after reading the consent form. Five different implementations of friction were added to a mock version of the IRMA application. The pilot study followed a between-subject design, where the participants were shown one of the five versions of the mock IRMA application. The participants were asked to open an application for a web store where they had to share fake personal data via IRMA with the webstore. They were asked to report on their first impression of the visual design of the application. Before showing the web store, they either confirm or decline the request of sharing their age and BSN attribute (privacy-unfriendly). After making this decision, the participants that chose to share their attributes were shown the friction design, while others were shown the friction desing directly. Then, after the decision was made, the web store application was shown for three seconds, regardless of their decision. The participants were told about the cover story, after which they were asked to fill in a short questionnaire.The participants of the pilot study were asked the questions as formulated in the UEQ-S and the DelibeRATE tool, as described in section *3.5 Questionnaire*. Furthermore, they were asked about their experience with Human Computer Interaction Design, the initial choice they made about sharing their personal data with Drankstore b.v. (to allow or disallow sharing BSN and age attribute), and the final decision they made regarding this choice (after they encountered the pattern in the cases where these options were presented to them before and after friction). The results of this can be found in table 1.

We used the highest mean score acquired by the DelibeRATE tool, given that it does not impact user experience significantly more than any of the other designs.

Version	1	2	3	4	5
Amount of participants	3	2	3	6	3
Declined at first option	1	0	1	3	1
Declined at second option	2	-	1	-	1
Mean DelibeRATE score	0.65	0.35	0.720	0.4375	0.5333
Pragmatic Quality	-1.167	0.500	0.778	1.167	0.250
Hedonic Quality	0.194	0.000	0.667	0.889	0.000
Overall UX	-0.067	0.292	0.720	1.00	0.021

Table 1: Results of the pilot study

Table 1 shows for each version, how many participants have seen it. It also shows how many of these participants chose not to share their age and BSN attributes. For versions 2 and 4 there was one option, since the friction was applied in the first and only screen in which the participants could choose to accept or decline. For versions 1, 3 and 5 there were two seperate opportunities to accept or decline: once before the friction was applied, and once after. Table 1 shows how many people declined at the first option, and how many people declined finally. This means that one participant, who used version 1, chose a different option after interacting with friction. How friction affected the choice of the participants of versions 2 and 4 is can not be told from this experiment.

Table 1 shows the mean of the DelibeRATE scores for each version, where version 1 and 3 score highest. These versions happen to also be very similar in design: version 3 shows the same text as version 1, but adds the time-locking element.

Table 1 shows the Pragmatic quality, Hedonic quality and overall user experience of the design. Version 4 scores best, however, only one of the six participants that used this version filled in the UEQ-S questionnaire. It is therefore hard to say whether these results are representative. Second highest scores version 3, for which all participants filled in the UEQ-S.

Because there was a low response rate, and many participants chose not to answer non-mandatory questions, we cannot draw too many conclusions from these results. However, the results of this pilot study might give us a small clue about the effects of the friction designs. The participants that were shown version 1 chose the option to deny sharing data the most. From the questionnaire we see that version 3 made the user deliberate their choice most (mean = 0.767, calculated with DelibeRATE tool), and version 4 showed the best user experience (Overall UEQ-S score of 1,00). Version 3 had the second highest user experience, with a score of 0.720. We chose to use version 3 for our main experiment, because it showed a good combination of deliberation and user experience. We are aware however, that this decision was made based on extremely little data, and should be seen as an educated guess that needs to be explored more in the future.

Furthermore, because of the experience gained by this pilot study, we decided to make a few changes in the questionnaire for the main experiment. First, the questions of both the DelibeRATe tool and the UEQ-S were made mandatory, even though the guide to these tools states that questions should have the option to be left blank. This decision was made to prevent the participants from leaving all answers blank, which happened often during the pilot study. Second, an open question was added at the end, asking the participant if they had any specific thoughts about the data they shared in the experiment. One of the participants of the pilot study disclosed that they did choose to share their data, but was confused as to why a liquor store would need their BSN. These thoughts show that the participant may have had a mindful thought process, even if it did not show up in the end result of the questionnaire.

2.6 Participants

Of the 56 participants were recruited for this experiment, 37 participants are students of Artificial Intelligence who did an introductory course to Human Computer interaction, or had other previous experience with user experience design. The students currently partaking in the iHCI course, the course on Human Computer interaction at Radboud University taught in the first year of the Artificial Intelligence Bachelor's programme, received an email by their teacher about the experiment. The other participants had no education on user experience design, and received the information about the experiment via individual emails. We aimed for the highest number of participants to recruit. Inclusion criteria for study participation were an age above 18 years, and proficiency in the Dutch and English language. The nationality of the participants was not measured. Participants received no money or other reward for the successful completion of the study, which was estimated to take around 10 minutes. The total sample population consisted of 29 females (53 percent) and 23 males (41 percent). The rest of the participants preferred not to indicate their gender. 37 participants (66 percent) fell into the age category 18-25 years, the others were scattered over the other categories and can be seen in Figure 6.

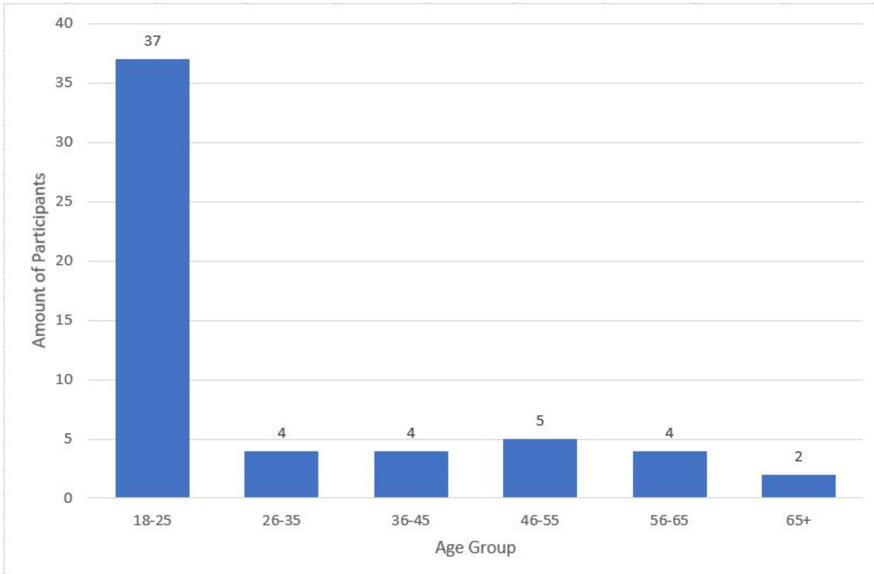

Figure 7: Plot of the amount of participants per age group

3 Results

3.1 Data Analysis

To analyse the results of the participants' decision about sharing their data we used summary statistics such as the mean and interquartile ranges.
The maximum DelibeRATE questionnaire overall score is 20, with each item being scored from 0 to 2 (0=No, 1=Unsure, 2=Yes). The overall DelibeRATE score and individual question scores were analysed and are presented using summary statistics such as the median and interquartile ranges or number per category and percentage as appropriate, as suggested by Gillies et al.. The higher the score, the more deliberation participants had when making a decision about sharing their personal data. To determine whether there was a difference in DelibeRATE scores between participants consenting or not, we performed a Mann–Whitney U test (two-sided, 5% significance level), as Gillies et al. suggest.
User experience design was tested using the Short Version of the User Experience Questionnaire, which has its own analysis tool available. The analysis of the different questions into the two categories (pragmatic and hedonic) is calculated with the accompanying Excel sheets. Furthermore, at the end of the experiment a comment section was provided for the participants to ask about their general thoughts of the experiment, and whether they had noticed something strange about sharing their personal data. Results of these comments were qualitatively analyzed by assessing if there was a trend or surprising thought.

3.2 Main Analysis

To test the three hypotheses, the analysis of the results was made up of three categories: the consent that was given or not, the deliberation of the decision, and user experience design. Lastly, we analyzed the comments to see if this tells us something about the participant's thought process regarding sharing their personal data in the experiment. All tests that were performed were deemed significant when results show a p-value < 0.05. To correct the p-values obtained by the statistical tests we used the Bonferroni-Holm correction(Holm, 1979). This method reduces the family-wise error rates for multiple hypothesis-tests, which could result in a Type 1 error[4] (Abdi, 2010). Unless stated otherwise, the reported p-values have all been corrected.

3.3 Consent

Because of the peculiar way our friction design functioned, the participants that could interact with this design had two separate moments in which they could agree or disagree to share their age and BSN attribute. Only when the participant chose to agree to share their attributes, they were shown the friction design, while the participants that immediately chose not to share their attributes did not interact with the friction design. This leaves us with results that can be interpreted in multiple ways. First, we can measure the difference in amounts of people that chose to share their personal information for Version 1, the control group, and Version 2, the friction design group. These are the between-subject results. Second, we can also analyse how many participants of the friction design group chose to change their mind about sharing information after interacting with the friction design. These are the within-subject results.

3.3.1 Between-subject results

Regarding consent behavior we observed that the majority of participants for either version chose to agree to share their age and BSN attributes with Drankstore b.v at the first chance: (26/29 for the control version 1 and 22/27 for version 2 with friction). Consent behavior did not change significantly between females and males, nor as a function of age, or previous education in Human Computer interaction. The participants that used version 2 had a second option to agree or disagree, of which 3/22 choose, after agreeing first, to deny on the second chance. To convert the data to numerical values we transformed the decision to reject sharing data to 1 (all results that show 'Weigeren' on first or second attempt), and the decision to accept to 0 (all results that show 'Toestaan' on first or second attempt).

[4]This method ensures that the family-wise error rate is smaller than the significance level (Holm, 1979)

Version	Version 1	Version 2
Total amount of participants	29	27
Rejected sharing their information	3	8
% of rejections	0.103	0.296

Table 2: Summary statistics for the DelibeRATE tool

Because of the small sample size we chose to perform a Fisher's exact test on the data to check if the difference between the two versions was significant. The p-values 0.536 before correction, and 1.000 after Holm's correction suggest that there is no significant difference between the results of version 1 and version 2. Consequently, we did not find support for Hypothesis 1.

3.3.2 Within-subject results

Out of the 27 participants in the friction design group, 22 chose to agree to share their personal information before interacting with the friction design. These 22 participants were then shown the time lockout design, after which they could choose to share their attributes once again. 3 Participants changed their mind after interacting with the time lockout friction design. To see whether a significant amount of people changed their mind, meaning that there is a significant difference in the dependent means, we performed a t-test for dependent means. With this test we check whether the change of 22/22 participants to 19/22 participants is a significant change. The p value of 0.1602 indicates that this result is not significant.

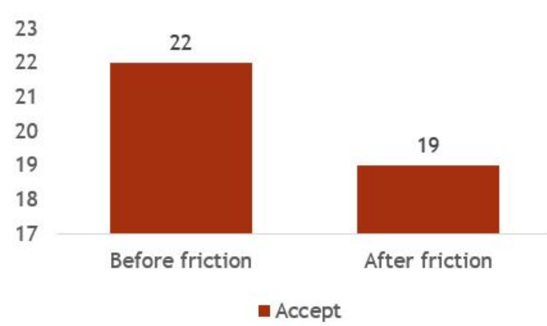

Figure 8: Bar plot for the within subject results: acceptance of sharing the age and BSN attributes before and after interacting with the friction design.

3.4 Deliberation

To analyze the results we use the recommended analysis methods provided by the DelibeRATE tool. Before the analysis can start, all answers need to be converted to numerical values (Yes=2, Uncertain=1, No=0), and the scores for each question need to be added up for each participant. This results in scores between 0 (minimal deliberation), and 20 (maximal deliberation). We also removed the data for the participants in the friction design group that chose not to accept sharing their personal information on the first attempt, because they never interacted with the friction design. Next, we can apply summary statistics to this data.

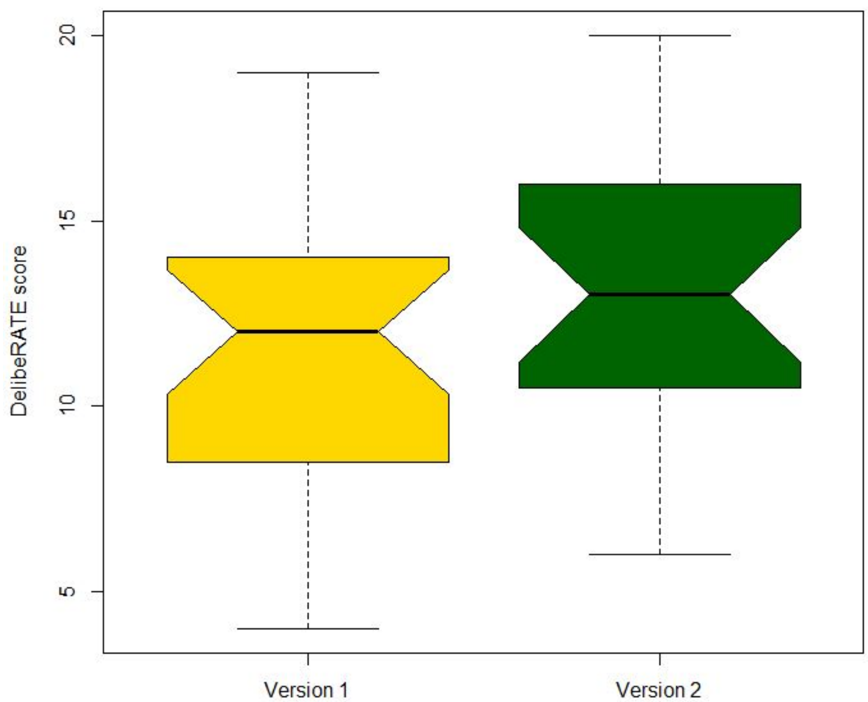

Figure 9: Box plot of the DelibeRATE scores for Version 1 (neutral design) and Version 2 (design with friction)

Figure 7 shows the Box plots that summarize the results of the DelibeRATE tool. The results show a difference between the two versions:

Version	Version 1	Version 2
Mean	11.63	13.43
Median	12	13
Variance	16.93	15.89
Standard Deviation	4.12	3.99

Table 3: Summary statistics for the DelibeRATE tool

After assuring the data was normally distributed we performed a t-test to see if the difference between the means was significant. A p-value of 0.032 shows us that this result is significant, however after using Holm's correction we obtained a p-value of 0.160, which means that the result is not to be considered significant. Hence, we did not find support for Hypothesis 2.

3.5 User Experience

To analyze the results of the UEQ-S, which will indicate how good or bad the user experience was for each of the versions, we used the analysis tool that comes with the UEQ-S. First, all scores are transformed from the original 1-7 scale to a scale from -3 to +3, with -3 as the most negative score and 3 as the most positive score. The tool then calculates a score for each individual question, after which these scores are categorized as Hedonic or Pragmatic qualities. Then the scores for the Hedonic and Pragmatic scales are calculated, which together form the overall score of the user experience. We also removed the data by participants in the friction design group that chose not to accept sharing their personal information on the first attempt, because they never interacted with the friction design.

Item	Mean	Variance	Std. Dev.	No.	Negative	Positive	Scale
1	2,3	1,3	1,1	27	obstructive	supportive	Pragmatic Quality
2	-0,4	1,2	1,1	27	complicated	easy	Pragmatic Quality
3	1,1	1,5	1,2	27	inefficient	efficient	Pragmatic Quality
4	2,0	1,7	1,3	27	confusing	clear	Pragmatic Quality
5	1,5	2,6	1,6	27	boring	exciting	Hedonic Quality
6	-0,3	1,1	1,0	27	not interesting	interesting	Hedonic Quality
7	-0,2	2,0	1,4	27	conventional	inventive	Hedonic Quality
8	-0,5	1,7	1,3	27	usual	leading edge	Hedonic Quality

(a) Version 1

Item	Mean	Variance	Std. Dev.	No.	Negative	Positive	Scale
1	2,4	0,6	0,8	28	obstructive	supportive	Pragmatic Quality
2	-0,6	2,1	1,4	28	complicated	easy	Pragmatic Quality
3	0,3	2,0	1,4	28	inefficient	efficient	Pragmatic Quality
4	1,3	1,9	1,4	28	confusing	clear	Pragmatic Quality
5	0,9	3,8	2,0	28	boring	exciting	Hedonic Quality
6	-0,2	2,7	1,7	28	not interesting	interesting	Hedonic Quality
7	0,0	3,4	1,9	28	conventional	inventive	Hedonic Quality
8	-0,2	3,1	1,8	28	usual	leading edge	Hedonic Quality

(b) Version 2

Figure 10: Scores per item for (a) Version 1 the neutral design, and (b) Version 2 the design with friction

Looking at the scores per item gives us an indication of which qualities differ per version. Both designs have a mediocre (around 0) score for the complicated-easy, not interesting-interesting, conventional-inventive and usual-leading edge qualities, and a good (>1) score for the obstructive-supportive, confusing-clear, and boring-exciting qualities. Both designs score very well on the obstructive-supportive quality, but version 2 has a higher mean and smaller variance and standard deviation. The biggest difference between the versions is in the efficient-inefficient and confusing-clear qualities, where version 1 scores much higher than version 2.

Short UEQ Scales	
Pragmatic Quality	1,231
Hedonic Quality	0,120
Overall	0,676

(a) Version 1

Short UEQ Scales	
Pragmatic Quality	0,857
Hedonic Quality	0,125
Overall	0,491

(b) Version 2

Figure 11: The mean scores of the UEQ-S for (a) Version 1, the neutral design, and (b) Version 2, the design with friction

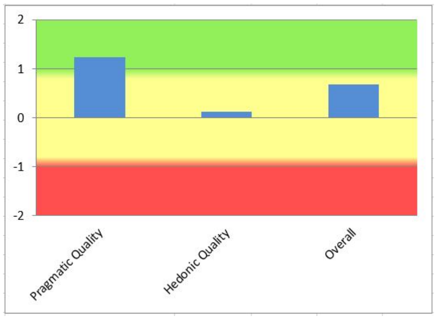

(a) Version 1

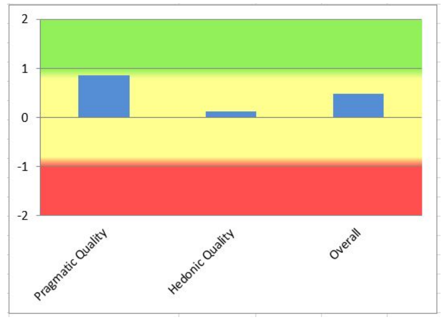

(b) Version 2

Figure 12: The scores of the UEQ-S for (a) Version 1,and (b) Version 2. Green scores indicate a good user experience, and red scores indicate a bad user experience

Figure 7 and 8 summarize the user experience for the Pragmatic and Hedonic Qualities, and the overall user experience as calculated with the UEQ-S. We see that the Pragmatic quality, consisting of scores for the obstructive-supportive, easy-complicated, efficient-inefficient and confusing-clear qualities is higher for Version 1 (mean = 1.231) than for Version 2 (mean = 0.857). The Hedonic scores for both versions are almost equal (mean Version 1 = 0.120 and mean Version 2 = 0.125). The overall score is highest for Version 1 (mean = 0.676), even though there is only a small difference compared to Version 2 (mean = 0.491). Figure 8 shows us that, even though there are differences between the scores, none of the scores should be interpreted as a negative or bad user experience. To see if the differences between the means are significant we performed a t-test on the transformed data, which showed us that none of the differences between the means were significant. The Pragmatic Quality had a p-value of 0.511 and the Hedonic Quality had a p-value of 1.000. The overall score had a p-value of 1.000. Therefore, we did not find support for Hypothesis 3.

3.6 Control Questions

To get a better understanding of the participants behavior during the study several control questions were asked. To the question whether they had any previous experience with Human Computer Interaction or User Experience design, 76% of the control group answered yes, and 61% of the group that interacted with the friction design answered yes. For neither of the groups there was any significant change in their consent behavior (p = 1.000). Equally, we did not measure any significant difference in their DelibeRATE score (p = 1.000) or their overall user experience (p = 1.000).
Furthermore, it was asked whether participants used a plugin in their browser to automatically remove cookies. Around 20% of the control group and 19% of the group that interacted with the friction design answered yes. Neither of the groups showed any significant change in the consent behavior of the people

that answered yes in comparison the people that answered no ($p = 1.000$), their DelibeRATE score ($p = 1.000$) and their overall user experience ($p = 1.000$).

In the comment section participants were asked if they noticed something strange about sharing their data, and if they had any general comments. The purpose of these questions is to find out if people are alarmed at all by sharing their BSN. To analyze the answers, which were given in an open text field, we use coding. Coding is not an exact measure of information, but it provides us with themes: general constructs within the data (Gibbs & Flick, 2007). Because of the small sample size it was possible to read and analyze all comments. Further analysis was done by finding three common categories among the comments. The first category consists of 5 participants commenting that they did not properly understand the experiment, for example this comment: "I Don't understand anything about this experiment and Survey". Further analysis on the relation of these comments to the age group of the participants showed that all these comments were given by participants in age categories above 45 years. The second category consists of comments around the theme of IRMA, where participants expressed that they did not like IRMA, did not want to use IRMA again, or did not understand the utility of IRMA. An example of this is the following comment: "I could not interact or find a reason why I should use the application other than a request to participate in this survey and I trusted the person who asked me". The third category consists of the theme BSN, where people expressed that they did not understand why their BSN was needed, or that they would never share their BSN if it was not part of the experiment. For example: "The presence of the BSN. Would not have given my actual BSN".

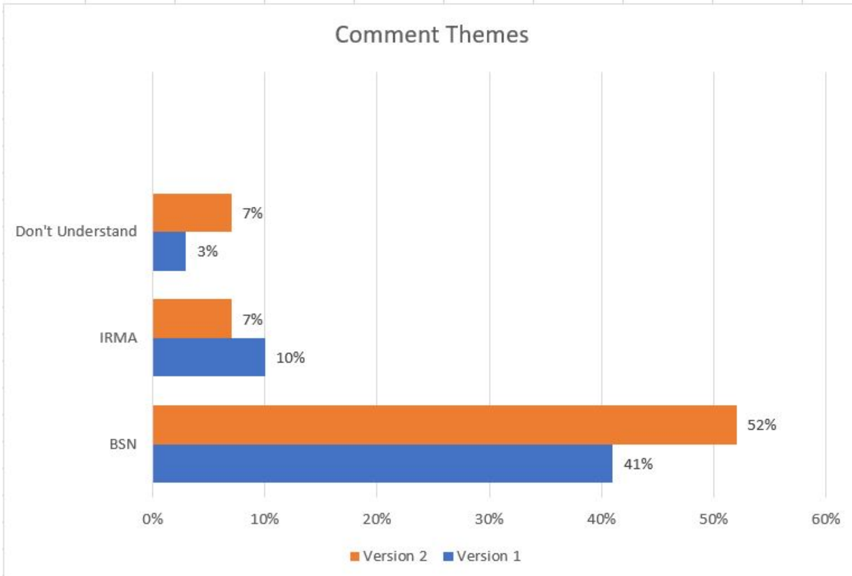

Figure 13: Bar Plot showing the percentage of participants whose comments could be categorized in one of the three categories. Version 1 is the control version, and version 2 is the design with friction.

Due to the small number of participants commenting on the first two themes, it is hard to say what this information tells us. The third category however, shows that almost half of the participants, regardless of whether or not they had rejected sharing their BSN, did notice that sharing their BSN may not be in their best interest. We can also tell that there are more comments about this theme for the version with friction, compared to the version without friction. [5]

Overall, it can be noted that neither their prior knowledge of user experience design, nor their individual privacy concerns seemed to have influenced the participant's answers in the questionnaire. Furthermore, it can be noted that although most participants chose to share their BSN with Drankstore b.v., they did notice that this was unusual. How these findings relate to the results of the analysis will be discussed in the following section.

4 Conclusion and Discussion

The goal of this project was to investigate whether friction as a design pattern could improve on the deliberation of users concerning their privacy decisions, and how friction affected their decision on sharing personal information. In addition, we wanted to investigate how this affected their user experience. Although

[5] We analyzed the difference between the two versions with Fisher's exact test. With a p-value of 1.052e-05, we show that commenting on sharing BSN is significantly higher when participants are shown the design with friction.

we found that more participants reported that they were alarmed by sharing their BSN in the group that was shown a pattern with friction compared to the group that was shown a neutral pattern, we did not find significant differences between the groups when it came to their privacy decision, deliberation or user experience. There appeared to be a significant difference between the deliberation scores, meaning that the group that was shown the design with friction did indeed deliberate their decision more than the group that was shown the neutral design. However, after correcting the p-value for multiple experiments, the results were no longer significant.
Most participants that left a comment reported that they thought that sharing their BSN with an online store was unusual, even though the majority of these people still decided to consent to this. This could potentially indicate that the friction design does have an effect on the thought process of the participants, even if it does not give the rationally preferred outcome of not sharing data. With regards to the changes in the GDPR I believe that a design that includes friction in order to gain informed consent should be further explored.

4.1 Explanatory approaches

We argue that there are 5 possible explanations for the observed behavior of the participants, concerning their final decision about sharing their personal data and their deliberation of this decision. Firstly, the experiment setting may have worked as a pattern itself, meaning that the email that the participants received with instructions on how to perform the experiment at home may have nudged them to consent to sharing their personal data. In the email it was mentioned that participants could share their personal data via IRMA, and even though it was not mentioned that this would be necessary to complete the experiment, it may have been interpreted this way by some of the participants. This would explain why the choices they made did not show significantly more deliberation, since the participants might not have found that there was much to deliberate on. Second, despite using a cover story similar to the one used in Graßl (2019), which had been tested for credibility, we cannot rule out that participants may have consented to sharing their personal information because they perceived everything as part of a "trustworthy" research study. Third, participants may have simply consented because the button to agree was on the right side of the screen. Custom for many applications is that the option to continue is on the right side of the screen, which may mean that the participants were simply conditioned to click buttons on this side of the screen more than the buttons on the left side of the screen. This also explains the lack of deliberation, since this behavior is generally a thoughtless process. Fourth, participants may also have been conditioned by their everyday life to agree to these types of consent statements, because many websites and apps require you to accept these statements in order to make use of the services. It may be that this conditioned behavior overwrote the effects of the friction design that was implemented. This could have prevented the participants from doing much deliberation, since the process may have been mindless due to their conditioning. Lastly, even though

we asked participants to act as if the mock BSN we provided them with was their own, they may have acted differently if they had to use their own personal information.
To explain the lack of difference in user experience between the two groups, we argue that this could be due to the design being carefully picked among the five designs tested in the pilot study. We deliberately chose a design that had a good user experience, to avoid the participants being annoyed by the design, which could affect their deliberation in unforeseen ways.

4.2 Implications

Although friction as a pattern to boost decision making seems to be promising when looking at the results of this experiment, the real life application of it may be more complicated. Most services that would ask a user to share more information than strictly necessary, benefit from the user sharing this information. Therefore, these services would have no interest in implementing such a pattern.Using friction as intended in this experiment in privacy statements by such services thus seems unrealistic. However, this does not mean that there is no use at all for such patterns. If users can be conditioned not to deliberate on their privacy decisions by use of Dark Patterns, then perhaps they can also be conditioned to deliberate on their choices more by using patterns with friction, as long as they are exposed to these patterns. Services with truly virtuous goals, such as IRMA, could gain benefit by exposing a user to such patterns the user will be conditioned to make more privacy-friendly decisions on the long run.

4.3 Limitations

Due to the specific way the friction pattern that we decided to experiment with was designed, a few limitations were unavoidable. Our friction design had two separate moments of choice, one before and one after the friction. This meant that, if participants choose on the first option not to share their personal information, they never interacted with the friction design. Even though this type of friction was tested to be the most beneficial for this research with the pilot study, a further research must take into account that other ways of implementing friction can have a vastly different effect on the behavior of the participants, and also on the statistics. An experiment in which participants, regardless of their choice about personal information, get to see a friction design, can show different results.

The ecological validity of this experiment was compromised by choosing to use a browser mock application. This decision was made mainly for practical reasons: the experiment had to be done online, at home, and with minimal assistance. Although the browser mock application did solve these issues, it did cause limitations in other areas. For one, the experimental environment setting was less controlled than it would have been, were the experiment performed in a more neutral setting instead of in the participant's own homes. How factors like background noise, bystanders and other distractions impacted the experiment is hard to say. Furthermore, since the mock application did not have to be installed on the participants' own devices, the setting deviated from real-life consent statements in apps.

Furthermore, even though the choice to test the friction design on the IRMA application was not an arbitrary one, this does have implications for the generizability of the experiment. IRMA is an application that has virtuous goals, and actively tries to help the user to be more conscious of their privacy. Most consent statements however are not designed with this goal in mind, and may therefore provide different results when tested with a similar experiment. Furthermore, some of the participants may have heard of IRMA before, and being reminded of IRMA's purpose might have caused them to make a different choice about sharing their data than they would have, had they not heard of IRMA before.

Lastly, due to the small amount of participants, it is hard to say whether the results of this experiment are meaningful at all. Performing the experiment again, with a larger, and more diverse group of people, would give us more insight into many of the aforementioned issues, since the results for this small experiment do provide us substantial ground for further investigation. Furthermore, since the pool of participants was not diverse in the age category, the results do not tell us much about the effects of friction on age categories >25 years.

4.4 Future Work

Future research should investigate why people that are alarmed by sharing privacy sensitive data still choose to do so in particular situations. Specifically, to address the limitations of this work, future research should investigate what it is that brings people to make these irrational decisions. Furthermore, this future research should address the role of the cover story and means of instructions, and what role conditioning of a similar situation plays in these types of decision making.

5 Acknowledgements

I would like to thank my supervisors Hanna Schraffenberger and Franc Grootjen for their advice and support during this process. Furthermore, I would like to thank Paul Graßl, for raising interesting issues in his Master thesis, and providing support and help with my research, especially on statistical tests. I would also like to thank my peer reviewers, over at the faculty of Philosophy, for providing insights from a different perspective. Lastly, I would like to thank all participants who performed the experiment.

References

Abdi, H. (2010). Holm's Sequential Bonferroni Procedure. *Encyclopedia of Research Design*, *1*(8), 1–8.

Awad, & Krishnan. (2006). The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to Be Profiled Online for Personalization. *MIS Quarterly*, *30*(1), 13. doi: 10.2307/25148715

Back, J., Brumby, D. P., & Cox, A. L. (2010). Locked-out: Investigating the Effectiveness of System Lockouts to Reduce Errors in Routine Tasks. In *Proceedings of the CHI '10 Extended Abstracts on Human Factors in Computing Systems* (3775–3780). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/1753846.1754054` doi: 10.1145/1753846.1754054

Cox, A. L., Gould, S. J., Cecchinato, M. E., Iacovides, I., & Renfree, I. (2016). Design Frictions for Mindful Interactions: The Case for Microboundaries. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing systems* (1389–1397). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/2851581.2892410` doi: 10.1145/2851581.2892410

Egelman, S., Cranor, L. F., & Hong, J. (2008). You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proceedings of the '10 CHI Conference on Human Factors in Computing Systems* (1065–1074). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/1357054.1357219` doi: 10.1145/1357054.1357219

Fansher, M., Chivukula, S. S., & Gray, C. M. (2018). #darkpatterns: UX Practitioner Conversations About Ethical Design. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (1–6). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/3170427.3188553` doi: 10.1145/3170427.3188553

Figma. (2020). *Where teams design together.* Available at `https://www.figma.com/` (2020/04/27).

Forbrukerrådet. (2018). *Deceived by Design, How Tech Companies use Dark Patterns to Discourage us from exercising our Rights to Privacy.* Retrieved from `https://www.forbrukerradet.no/undersokelse/no-undersokelsekategori/deceived-by-design/`

General Data Protection Regulation GDPR. (2018, May). European Council. Retrieved from `https://gdpr-info.eu/`

Gibbs, G. R., & Flick, U. (2007). Thematic Coding and Categorizing. In *Analyzing Qualitative Data* (38–56). Sage, Los Angeles.

Gillies, K., Elwyn, G., & Cook, J. (2014, 07). Making a Decision About

Trial Participation: The Feasibility of Measuring Deliberation during the Informed Consent Process for Clinical Trials. *Trials*, *15*, 307. doi: 10.1186/1745-6215-15-307

Gould, S. J., Cox, A. L., Brumby, D. P., & Wickersham, A. (2016). Now Check Your Input: Brief Task Lockouts Encourage Checking, Longer Lockouts Encourage Task Switching. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (3311–3323). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/2858036.2858067` doi: 10.1145/2858036.2858067

Graßl, P. (2019). *Dark Patterns in Consent Statements.* (unpublished Master thesis, Radboud University)

Grüne-Yanoff, T., & Hertwig, R. (2016). Nudge versus Boost: How Coherent are Policy and Theory? *Minds and Machines*, *26*(1-2), 149–183.

Holm, S. (1979). A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, *6*(2), 65–70. Retrieved from `http://www.jstor.org/stable/4615733`

Iacovides, I., Aczel, J. C., Scanlon, E., & Woods, W. I. (2011). What can Breakdowns and Breakthroughs tell us about Learning and Involvement experienced during Game-play. In *Proceedings of the 5th European Conference on Games Based Learning* (275–281). Athens, Greece.

Iacovides, I., Cox, A. L., & Knoll, T. (2014). Learning the Game: Breakdowns, Breakthroughs and Player Strategies. In *Proceedings of the CHI '14 Extended Abstracts on Human Factors in Computing Systems* (2215–2220). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/2559206.2581304` doi: 10.1145/2559206.2581304

Kahneman, D. (2015). *Thinking, Fast and Slow.* Farrar, Straus and Giroux. (ISBN: 9780374275631)

Lai, Y.-L., & Hui, K.-L. (2006). Internet Opt-in and Opt-out: Investigating the Roles of Frames, Defaults and Privacy Concerns. In *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research: Forty Four Years of Computer Personnel Research: Achievements, Challenges the Future* (253–263). Claremont, California, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/1125170.1125230` doi: 10.1145/1125170.1125230

Laugwitz, B., Held, T., & Schrepp, M. (2008). Construction and Evaluation of a User Experience Questionnaire. *USAB*, *5298*, 63-76. doi: 10.1007/978-3-540-89350-9-6

LimeSurvey Development Team. (2012). Limesurvey - The Free and Open Source Survey Software Tool! [Computer software manual]. Retrieved from `limesurvey.org`

Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, *15*(4), 336–355.

Malone, T. W. (1981). Toward a Theory of Intrinsically Motivating Instruction. *Cognitive Science*, *5*(4), 333 - 369. Retrieved from `http://www.sciencedirect.com/science/article/pii/S0364021381800171`. doi: https://doi.org/10.1016/S0364-0213(81)8 0017-1

Mathur, A. (2011). Improved Password Selection Method to prevent Data Thefts. *International Journal of Scientific and Engineering Research*, *2*(6), 118-119.

Nodder, C. (2013). *Evil by Design: Interaction Design to lead us into Temptation.* Wiley. (ISBN: 13: 978-1118422144)

Norman, D. A., & Draper, S. W. (1986). *User Centered System Design; New Perspectives on Human-Computer Interaction.* USA: L. Erlbaum Associates Inc. (ISBN: 0898597811)

O'Hara, K. P., & Payne, S. J. (1999). Planning and the User Interface: the Effects of Lockout Time and Error Recovery Cost. *International Journal of Human-Computer Studies*, *50*(1), 41 - 59. Retrieved from `http://www.sciencedirect.com/science/article/pii/S1071581998902348` doi: https://doi.org/10.1006/ijhc.1998.0234

Privacy by Design Foundation. (2020). *Over IRMA.* `https://privacybydesign.foundation/irma/`. (Online; accessed 13 July 2020)

Rauschenberger, M., Schrepp, M., Pérez Cota, M., Olschner, S., & Thomaschewski, J. (2013). Efficient Measurement of the User Experience of Interactive Products. How to use the User Experience Questionnaire (UEQ). Example: Spanish Language Version. *International Journal of Interactive Multimedia and Artificial Intelligence*, *2*, 39-45. doi: 10.9781/ijimai.2013.215

Rekhi, S. (2017). *The Hierarchy of User Friction.* Available at `https://medium.com/@sachinrekhi/the-hierarchy-of-user-friction-e99113b77d78` (2020/04/13).

Schrepp, M., Hinderks, A., & Thomaschewski, J. (2017). Design and Evaluation of a Short Version of the User Experience Questionnaire (UEQ-S). *IJIMAI*, *4*(6), 103–108.

Sengers, P., Boehner, K., David, S., & Kaye, J. APACrefauthors (2005). Reflective Design. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility* (49-–58). Aarhus, Denmark: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/1094562.1094569` doi: 10.1145/1094562.1094569

Smith, H. J., Dinev, T., & Xu, H. (2011). Information Privacy Research: an interdisciplinary Review. *MIS quarterly*, *35*(4), 989–1016.

Soeder, M. O. (2019). *Privacy Challenges and Approaches to the Consent Dilemma.* (Unpublished Master Thesis, Philipps-University Marburg) doi: 10.2139/ssrn.3442612

Solove, D. J. (2013). Introduction: Privacy Self-Management and the Consent Dilemma. *Harvard Law Review*, *126*(7), 1880-1903.

Terpstra, A., Schouten, A. P., Rooij, A. D., & Leenes, R. E. (2019). Improving Privacy Choice through Design: How Designing for Reflection could Support Privacy Self-management. *First Monday*, *24*(7). doi: 10.5210/fm.v24i7.9358

Wertheimer, A. (2000). What is Consent? And why is it Important? *Buffalo Criminal Law review*, *3*(1), 557–568.

Young, V. (n.d.). *Strategic UX: The Art of Reducing Friction.* Retrieved from `https://www.dtelepathy.com/blog/business/strategic-ux-the-art-of-reducing-friction`

6 Appendix

A

The questions were asked regarding deliberation:

1. I knew the option, of consenting or not, to sharing my personal data with Drankstore b.v. is available to me.
2. I understood that option, of consenting or not, to sharing my personal data with Drankstore b.v. was available to me.
3. I was aware of the advantages of sharing my personal data with Drankstore b.v. or not.
4. I was aware of the disadvantages of sharing my personal data with Drankstore b.v. or not.
5. I knew how I felt about the effects of sharing my personal data with Drankstore b.v.
6. I could imagine what the effects of sharing my personal data with Drankstore b.v. would be like.
7. I can judge whether sharing my personal data with Drankstore b.v. is best for me.
8. I can judge whether sharing my personal data with Drankstore b.v. is worst for me.
9. I had given the option, of sharing personal data with Drankstore b.v. enough thought.
10. I felt ready to make a decision about sharing my personal data with Drankstore b.v. or not.

B

The questions were asked regarding User Experience:

1. Complicated (1) or Easy (7)
2. Boring(1) or Exciting (7)
3. Obstructive (1) or Supportive (7)
4. Inefficient (1) or Efficient (7)
5. Confusing (1) or Clear (7)
6. Not interesting (1) or Interesting (7)
7. Conventional (1) or Inventive (7)
8. Usual (1) or Leading Edge (7)