Satisfying and Efficient Privacy Settings[1]

Manuel Rudolph, Svenja Polst

Fraunhofer IESE, Germany

manuel.rudolph@iese.fraunhofer.de, svenja.polst@iese.fraunhofer.de

Abstract

Data protection is becoming increasingly important for users of digital services. Recent studies show that users are concerned about having too little control over their personal data. However, users also complain that current interfaces for specifying privacy and security settings are too time-consuming and complicated. Therefore, we first identified the existing ways to configure these settings. Then we experimentally examined which type of specification is suited best for certain user types in terms of satisfaction or efficiency. Regarding efficiency, the type of specification that has the smallest number of options, called security level, is suited best for all users. Regarding satisfaction, there is no single type of specification that fits all user types; rather, different user types prefer different types of specification.

1 Introduction

Many users of digital services demand better control over their personal data (European Commission, 2015). At the same time, they rarely use the available privacy settings because they criticize that these are too complicated and time-consuming (Rudolph et al., 2018). In practice and in the literature, several types of user interfaces for privacy settings, which we call specification paradigms, have been proposed, such as wizards and privacy policy templates. They differ in their look and feel, interaction concepts and support options. The question arises whether different specification paradigms are equally suitable for different users with respect to satisfaction and efficiency. Users differ greatly in terms of their abilities (e.g. knowledge, available time, cognitive capacity) and preferences (e.g., the specification paradigms they like). Thus, there is probably no single specification paradigm that provides the best satisfaction and efficiency for all users. We assume that the right choice of specification paradigm by a specific user can positively influence the satisfaction with the tool itself and increase its

[1] This work is supported by the German Ministry of Education and Research (BMBF) under grant number 16KIS0328. The sole responsibility lies with the authors.

Veröffentlicht durch die Gesellschaft für Informatik e. V. 2018 in
R. Dachselt, G. Weber (Hrsg.):
Mensch und Computer 2018 – Tagungsband, 02.–05. September 2018, Dresden.
 https://doi.org/10.18420/muc2018-mci-0192

efficiency. Our contributions in this article are satisfaction- and efficiency-related recommendations of appropriate specification paradigms for specific user types based on experimental results. We also examined the specification effectiveness of the privacy policies (objective and perceived correctness), but omit this for space reasons. We present the selection of a user type model in Section 2 and the selected specification paradigms in Section 3, explain our experiment setting in Section 4, discuss the results in Section 5, and conclude the paper in Section 6.

2 Selection of a User Type Model

Users can be clustered into user types, for instance according to their character traits and resources. In the following, we will describe different user type models that we have found in the literature and argue for our selection of the model by Dupree (Dupree et al., 2016).

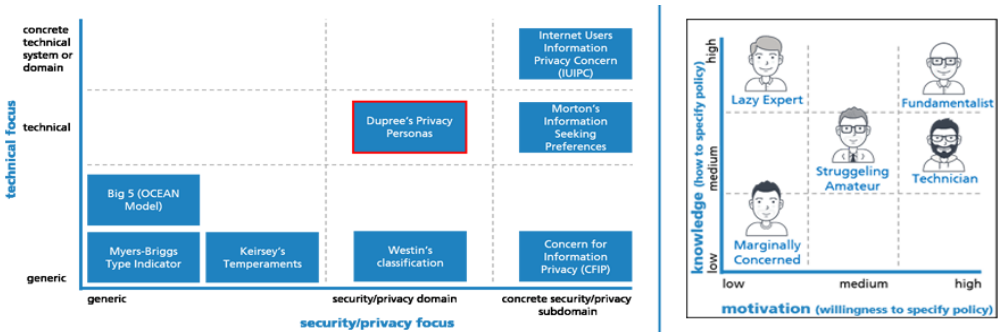

Figure 1: Left side: Classification of different user type models | Right side: Dupree's Persona Matrix

We characterized the relevant work according to two properties: their focus on IT security and privacy, and their focus on technical systems (see Figure 1, left side). Some very generic user type models describe human traits and behaviors without direct reference to a particular situation or domain. Examples are Keirsey's Temperaments (Keirsey, 1998), the Myers-Briggs Type Indicators (Myers et al., 1985), and the Big Five personality traits (Digman, 1990). Other approaches are related more closely to specific domains, e.g., computer usage, or include character traits that are relevant for security and privacy decisions. Westin's classification is based on users' privacy concerns and clusters users into three personas: Fundamentalist (high concern), Pragmatist (medium concern), and Unconcerned (low concern). Westin's work has been criticized (Urban & Hoofnagle, 2014) because he neglects the importance of knowledge or available information about privacy practices and domain-specific business processes. Considering this criticism, Dupree proposed her own five privacy personas (Dupree et al., 2016). They differ regarding the level of security and privacy knowledge and motivation for specifying one's own privacy and security settings. Morton's Information Seeking Preferences (Morton & Sasse, 2014) cluster users into five types: information controllers, security concerned, benefit seekers, crowd followers, and organizational assurance seekers. This approach is based on the ranking of 40 privacy-related statements and meant to support companies in providing

services that users prefer to adopt with respect to privacy behavior. Smith's approach (Smith et al., 1996) "Concern for Information Privacy (CFIP)" is quite generic and does not directly relate to technical systems. It measures the privacy concern of a person as a numerical value based on a calculation of fifteen privacy statements. In their approach called Internet Users' Information Privacy Concerns (IUIPC), Malhotra et al. (Malhotra et al., 2004) reflect on the concerns of Internet users about information privacy with a special focus on the individuals' perception of fairness in the context of data privacy by extending previous work (e.g., CFIP).

As we focus on the specification of privacy settings for technical systems, we identified the Dupree model as a suitable middle way and chose it for our experiment. She derived the five personas by performing personal interviews (see Figure 1, right side). The personas are called: *marginally concerned*, *amateur*, *technician*, *lazy expert*, and *fundamentalist*.

3 Different Types of Privacy Settings Interfaces

From the scientific literature and a state-of-the-practice analysis, we derived four prominent specification paradigms with which users can configure their privacy settings. They differ in terms of their look and feel, interaction concepts, and support options. All paradigms allow users to specify their privacy demands on the level of natural language. We define a privacy policy as the concrete privacy demand specified by a user via a privacy settings interface. We identified the following four specification paradigms:

- *Template instantiation*: This specification paradigm provides the user with a list of privacy policy templates. A privacy policy template is a kind of cloze where several options can be selected. The user can specify privacy policies by adjusting selection options, which can be represented by text fields, dropdown menus, radio buttons, or checkboxes. The users can themselves choose the order in which they want to specify different privacy policies. We found this paradigm in the Windows Group Policy Editor, the KAoS Policy Administration Tool (Uszok et al., 2003), and, to some extent, in the Facebook Privacy Settings.
- *Wizard*: The user can instantiate privacy policies in several small, consecutive configuration steps, each representing a kind of small template for privacy policies. Each step is well explained to the user. The user follows a concrete specification order. After executing all configuration steps, a complete set of privacy policies is specified. We found this paradigm in the Google Privacy Check Wizard, the PERMIS Policy Wizard (Permis, 2018), and in an active learning wizard by Fang (Fang & LeFevre, 2010).
- *Default policies*: In comparison to the template instantiation paradigm, the default polices paradigm contains predefined instantiations of privacy policies. Thus, the user has fewer configuration options and can only select from a limited number of predefined privacy policies per topic. This paradigm is used in the privacy settings of Twitter.
- *Security levels*: The user can select a single predefined set of privacy policies from a small list of sets. Thus, the user has to make one decision to configure the privacy settings. This paradigm is implemented by the Internet Explorer (Version 11).

4 Experiment Setting

The aim of the experiment was to find out which paradigms are suitable for certain users in terms of satisfaction and efficiency. We define satisfaction as the indicator of how much the users like to use the specification paradigm. We define efficiency as the time needed to specify privacy policies with the given specification paradigm. Therefore, we defined the following research questions:

- RQ1: Which specification paradigm best fits a particular type of person represented by a persona in terms of satisfaction?
- RQ2: Which specification paradigm is suited best for a particular type of person represented by a persona in terms of efficiency?

4.1 Scenario & Tasks

The participants were asked to empathize with a concrete scenario in which they use three novel digital services for their village, which were explained to them by a text on the handout and a promotion video. The scenario description explained that these services have potential effects on citizens' privacy as personal data are processed. Concrete privacy demands were formulated by the authors as tasks, which the participants were asked to solve during the experiment. There were six tasks, for example: "When I place an order in the 'BestellBar' app, I do not under any circumstances want to receive advertising from other providers that refers to the ordered product. They may not use my data." The scenario description and the tasks were provided on a digital handout, which we recommended should be printed. The participants watched the video, then read the scenario description and afterwards performed the same six tasks on each specification paradigm. Four specification interfaces were created according to the selected specification paradigms shown in Section 3.

4.2 Procedures, Instruments & Execution

Our experiment was executed as a publicly accessible online experiment and was offered in German and in English. We decided on an online experiment in order to reach a large number of diverse people. To avoid misuse, participants could only start the experiment once with a unique eight-character participant ID. The participants were recruited from the authors' institution, friends and their friends. About 120 personal invitation emails with the handout attached were sent to interested parties. The handout contained instructions for starting the experiment, the individual participant ID, and the scenario description. The participants had 14 days to participate in the experiment. The expected duration of the experiment was 30-40 minutes, but there was no time limit. It was possible to interrupt the experiment and continue at the same point using the participant ID. However, it was not possible to repeat steps that had already been executed.

Our experiment was structured as follows: First, the participants had to agree to the informed consent and confirm that they are at least 18 years of age to avoid legal problems when analyzing data from minors. Afterwards, they were asked to answer demographic questions about age, gender, and educational level, as well as their relationship to the institution of the authors and its research topics. A self-assessment of their own expertise and motivation in the areas of IT security, their protection of their own privacy, and their experience in dealing with digital services was requested. The participants were then asked to choose the one out of five personas offered that suits them best. For this purpose, all five personas introduced by Dupree were described based on nine to twelve original characteristics and habits formulated in the first-person perspective ("I use public wireless networks without further protection measures"). The order of the displayed personas was determined randomly. Then the scenario with the concrete tasks was explained via video and handout. Next, the participants were instructed to complete all six tasks four times, each time using a different specification paradigm. The order of the specification paradigms had been determined randomly in order to minimize learning effects. Right after each specification paradigm, the participants rated on a five-point scale how they liked this type of specification. We also measured how long it took them to specify all tasks with each paradigm. Next, they were asked to rank the four specification paradigms according to their preference for using them in real life. They could comment on this ranking with free text. Finally, they were asked to determine how well they could identify with the scenario and the selected persona and were given the opportunity to submit comments in free text form.

5 Results & Discussion

Out of 120 invitations sent, 63 people started the experiment and 61 people completed it (53% being male). We did not find any evidence that would have caused us to regard complete records as invalid. The participants' age ranged from 18 to 82 years (M=40.5; SD=14.4). Eleven percent of the participants had a secondary school leaving certificate as their highest level of education, seven percent had an entrance qualification for higher education, 54 percent had a university degree, and 15 percent held a doctoral degree. 20 percent of the participants chose the persona *marginally concerned*, 34 percent the *amateur*, 18 percent the *lazy expert*, 23 percent the *technician*, and five percent the *fundamentalist*. We asked the participants how well the chosen persona corresponds to their personality, on a scale of 1 (not very good, but it best matched the five options) to 5 (I can identify very well with the persona). The average response was 3.8. Not a single person reported the value 1. In addition, the self-reported security knowledge was a good fit for the selected persona for most of the participants.

5.1 Satisfaction

After each use of a specification paradigm, the participants indicated how much they liked it on a five-point scale ranging from 1 ("I really dislike this specification paradigm") to 5 ("I really like this specification paradigm"). Overall, they liked the *template instantiation* paradigm most (cf. Table 1 and Figure 2 – left side). In second and third place are *wizard* and *default policies*. The *security level* paradigm was considered least satisfying. The participants

also ranked the paradigms according to their preference. In the ranking, the *security level* paradigm was most often ranked last, regardless of the chosen persona.

Table 1: Satisfaction with specification paradigms for personas (M: mean, SD: standard deviation, Mdn: median)

				Ranking in percent							Ranking in percent			
	M	SD	Mdn	1st	2nd	3rd	4th	M	SD	Mdn	1st	2nd	3rd	4th
	Marginally Concerned							Technicians						
Template Inst.	3.9	0.9	4	17	58	25	0	4.1	1.1	4	43	7	29	21
Default Policies	3.3	1.4	3.5	17	8	42	33	3.4	1.3	4	14	21	43	21
Security Levels	3.0	1.2	3	17	25	8	50	3.2	1.5	3	29	14	7	50
Wizard	4.0	1.2	4	50	8	25	17	3.8	1.2	4	14	57	21	7
	Amateurs							Fundamentalists						
Template Inst.	3.8	0.9	4	43	48	10	0	4.3	1.2	5	67	0	33	0
Default Policies	3.3	1.2	4	24	10	48	19	4.3	0.6	4	0	33	33	33
Security Levels	2.1	1.2	2	0	14	19	67	3.3	2.1	4	0	0	33	67
Wizard	3.8	0.7	4	33	29	24	14	4.3	0.6	4	33	67	0	0
	Lazy Experts							All Participants						
Template Inst.	4.0	1.1	4	45	45	0	9	4.0	1.0	4	39	38	16	7
Default Policies	3.0	1.1	3	0	9	64	27	3.3	1.2	4	15	13	48	25
Security Levels	1.9	0.8	2	9	9	27	55	2.6	1.4	2	11	15	16	57
Wizard	3.8	1.3	4	45	36	9	9	3.9	1.0	4	34	34	20	11

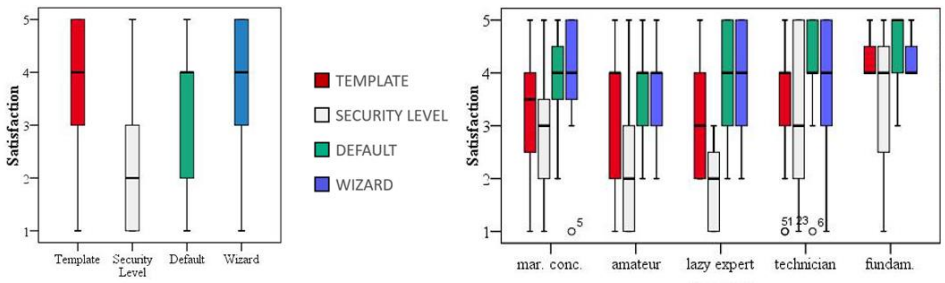

Figure 2. Left side: Satisfaction with paradigms | Right side: Satisfaction with specification paradigms per persona

The satisfaction results separated by persona are displayed in Table 1 and Figure 2 (right side). They show that *security levels* are the least satisfying paradigm within all persona groups and that the *default policy* paradigm is in third place for all personas except the *fundamentalists*. The *marginally concerned* preferred the *wizard* and *template instantiation* paradigms. In the final ranking, 50 percent voted for *wizard* in first place and a majority voted for *template instantiation* in second place. Thus, we will recommend the *wizard* paradigm for the *marginally concerned* if satisfaction is the focus. The *amateurs* almost equally liked the *template instantiation* and *wizard* paradigms after the individual specifications with the paradigms. However, in the final ranking, they clearly voted for *template instantiation* in first place, followed by *wizard*. 67 percent ranked *security levels* last. Thus, we recommend the *template instantiation* paradigm for *amateurs* if satisfaction is the focus. The *lazy experts* voted similarly as the *amateurs*. Overall, the *template instantiation* paradigm was liked most, followed by the *wizard* paradigm. As the differences for the *template instantiation* and *wizard* paradigms are marginal,

we recommend both for *lazy experts* if satisfaction is the focus. The *technicians* liked the *template instantiation* paradigm the most, followed by the *wizard* paradigm. However, the other two paradigms also got quite high ratings. This more even distribution of satisfaction over the paradigms is also evident in the rankings, where each paradigm was ranked first by at least 14 percent of the *technicians*. The *template instantiation* paradigm was ranked top by 43 percent of the participants. Although satisfaction with the paradigms is more equally distributed among the *technicians*, we can recommend *template instantiation* as the most satisfying paradigm. We cannot give any recommendations for the *fundamentalists* as the sample group was too small (n=3). Nevertheless, we can tell that directly after use, the *template instantiation*, *default policies*, and *wizard* paradigms were equally satisfying. In the final ranking, two participants voted for *template instantiation* in the first place; *wizard* was ranked second.

In summary, with respect to satisfaction (RQ1) we can recommend the *wizard* specification paradigm for the persona *marginally concerned* and the *template instantiation* paradigm for the personas *amateur* and *technician*. Both paradigms are equally suitable for the persona *lazy expert*. Due to the small number of *fundamentalists*, we cannot give any recommendations.

5.2 Efficiency

We measured the time each participant needed to complete each of the specifications in the four paradigms. We excluded the data sets of two participants from the analysis, as each had an extreme outlier in one paradigm. This can only be explained by a longer pause during the experiment. The other time data are reasonable regarding the minimum time needed to fulfill a task properly. In terms of the total population, the *security levels* paradigm proved to be the most efficient (M=1.8 minutes) method for specifying privacy settings. There are smaller differences in the average time of the other paradigms, ranging from 3.1 to 3.8 minutes (cf. Figure 3, left side). The second most efficient paradigm was *template instantiation*. The participants needed the longest time for *wizard*. The *lazy experts* needed less time for solving all tasks in all four paradigms (cf. Table 2 and Figure 3, right side) than the other personas (about 2 minutes less than the average time of the remaining population). This may confirm the low motivation of the *lazy experts* to spend time on privacy specification according to the Dupree model. The *marginally concerned* needed an average of about 2.5 minutes longer than the remaining population, which can be explained by their low level of knowledge. The other three personas needed an average of 11.8 to 12.9 minutes for all paradigms. *Lazy experts*, *fundamentalists*, and *amateurs* needed the longest for *wizard* and performed the fastest with *security levels*. For *technicians* and *fundamentalists*, the time needed for *template instantiation* and *default policies* was almost equal. The *technicians* were clearly also the fastest when using the *security levels* paradigm; however, they took an equal amount of time with all others.

In general, when recommending the most efficient paradigm, one has to consider the desired degrees of freedom in the specification. We measured the degrees of freedom in the experiment as the concrete decisions a participant has to take when specifying all six tasks in a paradigm. These were significantly higher for *template instantiation* and *wizard* (both 18 decisions) than for *default policies* (6 decisions) or *security levels* (1 decision).

Table 2: Mean time of paradigm use in minutes

Mean times in minutes	Template Instantiation	Security Levels	Default	Wizard	All paradigms
Marginally Concerned	3.4	2.6	4.3	4.0	14.3
Amateur	3.0	1.6	3.4	3.8	11.8
Lazy Expert	2.7	1.1	2.7	3.7	10.3
Technician	3.5	1.8	3.5	3.5	12.3
Fundamentalist	3.5	1.4	3.5	4.5	12.9
All Participants	3.1	1.8	3.5	3.8	12.2

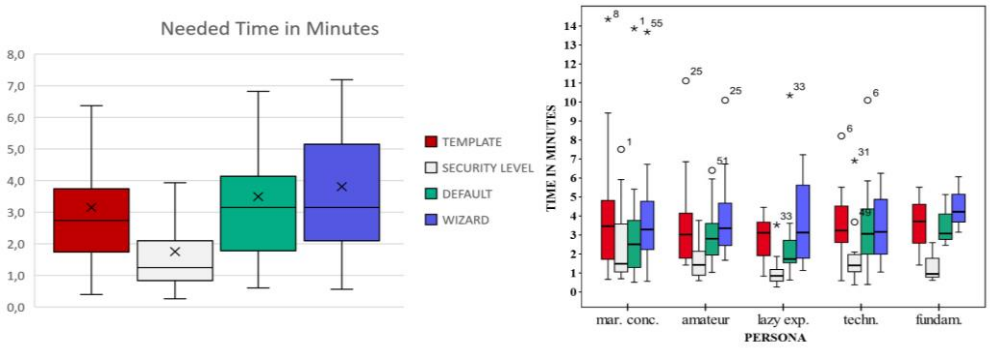

Figure 3: Time needed to complete all tasks using a specification paradigm in minutes of the whole population (left) and for specific personas (right)

In summary, if only efficiency is important (RQ2), the *security levels* paradigm can be recommended for all user types. On average, the *template instantiation* paradigm had the second fastest results. Thus, if more degrees of freedom are desired, we recommend for all participants *template instantiation* as the second most efficient method. Additionally, the *template instantiation* paradigm is very popular among all user types as it was most satisfying for all personas except the *marginally concerned*, where it resulted in the second highest level of satisfaction.

5.3 Comparison of Satisfaction, Efficiency & Correctness

Besides satisfaction and efficiency regarding privacy settings interfaces, we also evaluated the correctness of the tasks performed by the user types by measuring the objective correctness (by comparing policies configured according to tasks to the sample solution). Overall, we found that correctness correlates with the degrees of freedom of the specification paradigms. The fewest mistakes were made with the *security levels* paradigm. More freedom led to more mistakes. However, for the whole population, the average number of mistakes differed only marginally between the *default policies*, *template instantiation*, and *wizard* paradigms. There are deviations in the recommendations regarding the three qualities satisfaction, efficiency, and effectiveness. For instance, the *security levels* paradigm requires the shortest specification time and is related to the fewest mistakes, but the participants disliked it. We conclude that the

selection of the most suitable specification paradigm must be aligned with a ranking of the relevance of the qualities satisfaction, efficiency, and correctness.

5.4 Threats to validity

There are several threats to internal validity. We did not observe the participants and therefore we cannot exclude the possibility that they talked about the experiment with other participants before or during their participation. The participants might also not have been able to concentrate well enough in an appropriate environment to efficiently solve the tasks and evaluate the paradigms. We also cannot rule out that the participants took small breaks during the specification. We excluded two participants when analyzing efficiency due to obvious large breaks. Given the large number of participants (n=61 for satisfaction and n=59 for efficiency) and the instructions before the experiment, we consider these threats to be low. In addition, participants who did not identify well with the scenario may have had less motivation for dealing with the paradigms in the experiment. This may have had a negative effect on the results.

We also face threats to external validity. The *security levels* paradigm in combination with the given tasks most likely does not reflect reality since the preset tasks were perfectly matched by one of the *security levels*. This is rarely the case in real life. We decided to propose a perfect solution as the lack of a perfect match might have confused the participants and thus endangered internal validity. Another threat is that the experiment was conducted in a scenario representing a single use case for privacy settings (mono-operation bias). Also, we cannot guarantee that the participants are representative of the population regarding security knowledge and skills. Further experiments confirming our results in various scenarios and with other participants would increase the generalizability of our results and thus their external validity.

Finally, we see some threats to conclusion validity with regard to our recommendations of the most suitable specification paradigms. The number of participants per persona was quite small, especially the number of *fundamentalists* (three persons). In addition, we cannot rule out that our participants are not representative for the whole population. More participants are needed to draw conclusions that are representative for all users. The selection of specification paradigms for our experiment is based on our observations of the paradigms most commonly used in practice. We cannot exclude the existence of other paradigms better suited to the personas.

6 Conclusion & Future Work

In this paper, we analyzed whether the efficiency and satisfaction of configuring their own privacy settings with different specification paradigms differs among users. We grouped them according to Dupree's model based on knowledge of IT security and privacy as well as motivation to configure their own privacy settings. It turned out that there are differences between the user types, but overall these are marginal. Our recommendations regard either efficiency or satisfaction, as well as the desired granularity in specifying privacy settings using these paradigms. In the end, software developers must consider the importance of the two qualities satisfaction and efficiency when choosing the right specification paradigm for the end users.

We measured the objective correctness of the specifications (by comparing the results with the sample solution) and the perceived correctness (by comparing objective correctness with self-estimation). We will discuss the comparison of efficiency, satisfaction, and correctness in more detail on the level of individual personas in future work and will then make further empirically proven recommendations. We plan to conduct a similar experiment in a different scenario with different participants to underpin the generalizability of our results.

References

Digman, J. M. (1990). Personality structure: Emergence of the five-factor model. Annual review of psychology, 41(1), 417-440.

Dupree, J. L., Devries, R., Berry, D. M., & Lank, E. (2016). Privacy personas: Clustering users via attitudes and behaviors toward security practices. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (pp. 5228-5239). ACM.

European Commission (2015): Special Eurobarometer 431 - Data Protection. Available online at http://ec.europa.eu/commfrontoffice/publicopinion/archives/ebs/ebs_431_en.pdf

Fang, L., & LeFevre, K. (2010). Privacy wizards for social networking sites. In Proceedings of the 19th international conference on World wide web (pp. 351-360). ACM.

Keirsey, David (1998): Please understand me 2: Prometheus Nemesis Book Company.

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.

Morton, A., & Sasse, M. A. (2014). Desperately seeking assurances: Segmenting users by their information-seeking preferences. In Privacy, Security and Trust (PST), 2014 Twelfth Annual International Conference on (pp. 102-111). IEEE.

Myers, I. B., McCaulley, M. H., Quenk, N. L., & Hammer, A. L. (1998). MBTI manual: A guide to the development and use of the Myers-Briggs Type Indicator (Vol. 3). Palo Alto, CA: Consulting Psychologists Press.

Permis. (20.02.2018). Computing Laboratory, University of Kent. http://sec.cs.kent.ac.uk/permis.

Rudolph, M., Feth, D. and Polst, S. (2018). Why Users Ignore Privacy Policies: A Survey and Intention Model for Explaining User Privacy Behavior. 19th International Conference on Human-Computer Interaction (HCII), Las Vegas, USA.

Smith, H. J., Milberg, S. J., & Burke, S. J. (1996). Information privacy: measuring individuals' concerns about organizational practices. MIS quarterly, 167-196.

Urban, J., & Hoofnagle, C. (2014). The privacy pragmatic as privacy vulnerable. In: Workshop on Privacy Personas and Segmentation. Symposium on Usable Privacy and Security (SOUPS). Menlo Park, CA, July 9-11.

Uszok, A., Bradshaw, J., Jeffers, R., Suri, N., Hayes, P., Breedy, M., Bunch, L., Johnson, M, Kulkarni, S. & Lott, J. (2003, June). KAoS policy and domain services: Toward a description-logic approach to policy representation, deconfliction, and enforcement. In Policies for Distributed Systems and Networks, 2003. Proceedings. POLICY 2003. IEEE 4th International Workshop on (pp. 93-96).