

UNIVERSITY OF CALIFORNIA,  
IRVINE

A User-Tailored Approach to Privacy Decision Support

DISSERTATION

submitted in partial satisfaction of the requirements  
for the degree of

DOCTOR OF PHILOSOPHY

in Information and Computer Sciences

By

Bart Piet Knijnenburg

Dissertation Committee:  
Professor Alfred Kobsa, Chair  
Professor Donald J. Patterson  
Professor L. Robin Keller

2015



## **DEDICATION**

To

my wife Tuğçe,  
my brother Willem,  
my parents,  
and my friends

in recognition of their undying support.

To the reader of this dissertation:

“It is the mark of a truly intelligent person  
to be moved by statistics.”

George Bernard Shaw

# TABLE OF CONTENTS

	Page
TABLE OF CONTENTS	iii
LIST OF FIGURES	v
LIST OF TABLES	vii
ACKNOWLEDGMENTS	viii
CURRICULUM VITAE	ix
ABSTRACT OF THE DISSERTATION	x
CHAPTER 1: Motivation	1
CHAPTER 2: Problems with transparency and control	4
2.1 Privacy calculus	4
2.2 Transparency and control	6
2.3 Ironical effects of transparency	7
2.4 Bounded rationality and decision biases	8
2.5 Context effects: an example of decision biases (original work)	12
2.6 Control and Transparency Paradox	15
2.7 Form auto-completion tools and the control paradox (original work)	16
2.8 Conclusion: Transparency and control do not work	21
CHAPTER 3: Problems with privacy nudges	23
3.1 Privacy nudges	23
3.2 Failed nudges: Justifications and request order (original work)	26
3.3 Nudges: a one-size-fits-all approach	32
CHAPTER 4: Contextualization of privacy decisions	34
4.1 Introduction	34
4.2 Types of users and types of information	34
4.3 Multi-dimensionality and privacy profiles (original work)	37
4.4 Types of recipients	45
4.5 Recipient groups (original work)	48
CHAPTER 5: Personalized Privacy	63
5.1 Introduction	63
5.2 Adaptive justifications for app recommenders (original work)	67
5.3 Smart defaults for social network sites (original work)	71
5.4 Privacy recommendations for location-sharing (original work)	74
CHAPTER 6: Adaptive request order for demographics-based recommender systems (original work)	84
6.1 Motivation	84
6.2 Background	85
6.3 Pre-study: linking demographics to attributes	93
6.4 Study 1: testing demographics vs. attribute-based recommenders	98
6.5 Towards a better item request order	113
6.6 Study 2: testing different request orders	125
6.7 Conclusion regarding the adaptive request order studies	146



CHAPTER 7: General conclusion	148
7.1 Summary	148
7.2 Practical considerations	150
7.3 Future directions	156
REFERENCES	158

## LIST OF FIGURES

	Page
Figure 1: Experimental model.	18
Figure 2: Perceived Risk and Perceived Relevance per Website and Item type.	20
Figure 3: Disclosure rate per Website and Item type, for each of the three tools.	21
Figure 4: The Structural Equation Model (SEM) for the data of the experiment.	29
Figure 5: The total effects of the justifications.	30
Figure 6: The steps involved in our analysis of the dimensionality of information disclosure behaviors.	38
Figure 7: User profiles in the app recommender study.	40
Figure 8: Differences between profiles.	40
Figure 9: User profiles in the Facebook profile study.	41
Figure 10: Differences between profiles.	42
Figure 11: User profiles in the online retailer study.	44
Figure 12: Differences between profiles.	44
Figure 13: Analysis of the Family-related labels.	53
Figure 14: Analysis of the Friends-related labels.	54
Figure 15: Analysis of the Colleagues-related labels.	55
Figure 16: Analysis of the Classmates-related labels.	55
Figure 17: Analysis of the Acquaintances-related labels.	57
Figure 18: Mundo profile settings interface.	60
Figure 19: The effects of justification type and request order on disclosure.	68
Figure 20: The estimated effects of justification type and request order on the subjective constructs.	69
Figure 21: Observed disclosure levels in each default condition.	73
Figure 22: The effect of the recommenders on subjective measures.	82
Figure 23: Example of the values of energy-saving attributes for different age groups.	97
Figure 24: Example rules of the energy recommender when the user answers a question.	97

Figure 25: The healthy-living recommender with the attribute-based PE-method.	101
Figure 26: The energy-saving recommender with the demographics-based PE-method.	101
Figure 27: Structural Equation Model for the energy-saving recommender system.	108
Figure 28: Structural Equation Model for the healthy-living recommender system.	109
Figure 29: Total effects of domain knowledge and PE-method on trust in the provider.	111
Figure 30: Ordered demographic item disclosure rates for the energy- and health recommender.	115
Figure 31: Structural Equation Model for study 2.	132
Figure 32: Total effects of the manipulation on perceived control.	134
Figure 33: Total effects of manipulation x privacy concerns on perceived recommendation quality.	134
Figure 34: Total effects of manipulation x privacy concerns on perceived privacy threat.	135
Figure 35: Total effects of manipulation x domain knowledge on perceived privacy threat.	135
Figure 36: Total effects of manipulation x privacy concerns on trust in the provider.	136
Figure 37: Total effects of manipulation x domain knowledge on trust in the provider.	136
Figure 38: Total effects of manipulation x privacy concerns on system satisfaction.	137
Figure 39: Total effects of manipulation x domain knowledge on system satisfaction.	137
Figure 40: Total effects of manipulation x privacy concerns on choice satisfaction.	137
Figure 41: Total effects of manipulation x domain knowledge on choice satisfaction.	137
Figure 42: Total effects of manipulation x privacy concerns on total calories burned/avoided.	138
Figure 43: Total effects of manipulation x domain knowledge on total calories burned/avoided.	138
Figure 44: Average percentage of items seen and disclosed per request order.	139
Figure 45: Structural Equation Model for study 2, including disclosure behavior.	141

## LIST OF TABLES

	Page
Table 1: Subjective position and choice proportion of each option, for sharing with friends/colleagues.	14
Table 2: Subjective position and choice proportion of each option, for sharing with coupons/apps.	14
Table 3: Experimental manipulations: strategies to influence information disclosure.	28
Table 4: Items used in the app recommender study.	39
Table 5: Items used in the Facebook study.	41
Table 6: Items from the online retailer study.	43
Table 7: Items used in the SNS example.	50
Table 8: Labels used in the SNS example.	50
Table 9: Heuristics to find the best strategy.	70
Table 10: The co-occurrence of actions and evaluations.	77
Table 11: Recall in the 5 recommender conditions.	80
Table 12: Demographics questions asked.	93
Table 13: Attributes of the energy-saving and healthy-living measures.	95
Table 14: AVEs and loadings for the questionnaire items in the two conditions of Study 1.	105
Table 15: Simulations of model user behaviors under the two adaptive thresholds.	124
Table 16: Questionnaire used in study 2.	128

## ACKNOWLEDGMENTS

I would like to express immense gratitude to my committee chair, Professor Alfred Kobsa, who has supported me by not only advising me on my research, but to prepare me for my career as an academic. His “hands off approach with strategic interventions” has proven to be a perfect match to my preference for working independently. His keen insights have helped me improve my writing as well as my general approach to doing research.

I would like to thank my committee members, Professor Don Patterson and Professor Robin Keller, whose helpful feedback on my proposal has guided me in finishing the final studies of my thesis.

In addition, I would like to thank all my collaborators (there are too many to list them all by name), who have shaped my views on privacy and/or recommender systems through our interactions while writing. Among my co-authors, my “old advisor”, Professor Martijn Willemsen, deserves special recognition. He inspired my interest in human decision making and recommender systems, and taught me how to carefully analyze a problem through controlled experimentation.

I am also grateful for all the work that has been done by the various students who contributed to my research, either as remote collaborators under the supervision of Martijn, or as interns at UC Irvine.

I thank my supporters Professor Heng Xu, Professor Joe Konstan, and Dr. Burcu Bulgurcu for their advice and support in my search for an academic job. I also thank the people at Clemson University for offering me a tenure track position in the Human-Centered Computing division; I cannot wait to start a new academic adventure in South Carolina!

Last but certainly not least, I am greatly indebted to Steven Langerwerf, who has provided continuous technical support in terms of programming advice, server maintenance, and troubleshooting. Steven is a great friend, and a better personal tech support simply does not exist!

Financial support was provided by the University of California, Irvine, NSF Grant SES-1423629, a Google U.S./Canada Fellowship in Privacy, and gifts by Samsung, Intel, and Ericsson.

## **CURRICULUM VITAE**

### **Bart Piet Knijnenburg**

2002-05	B.S. in Innovation Sciences, Eindhoven University of technology
2003-06	Teaching Assistant, Innovation Sciences, Eindhoven University of Technology
2005-06	Education Officer of Study Association “Intermate”, Eindhoven University of Technology
2006-07	M.A. in Human-Computer Interaction, Carnegie Mellon University
2007-09	M.S. in Human-Technology Interaction, Eindhoven University of Technology
2007-09	Interaction Designer, Aduna/Vound
2009-10	Researcher and Lecturer, Eindhoven University of Technology
2011-12	Teaching Assistant, Information and Computer Science, University of California, Irvine
2012-14	Research Intern, Samsung Information Systems America
2014	Teaching Associate, Information and Computer Science, University of California, Irvine
2010-15	Ph.D in Information and Computer Science, University of California, Irvine

# ABSTRACT OF THE DISSERTATION

A User-Tailored Approach to Privacy Decision Support

By

Bart Piet Knijnenburg

Doctor of Philosophy in Information and Computer Sciences

University of California, Irvine, 2015

Professor Alfred Kobsa, Chair

As an increasingly important part of our social, professional and financial lives happens online, the frequency with which we have to deal with privacy problems is ever on the rise. In this dissertation I answer the question: **How can we help users to balance the benefits and risks of information disclosure in a user-friendly manner, so that they can make good privacy decisions?**

After briefly motivating this question in Chapter 1, I will first discuss problems with existing answers to this question in Chapters 2 and 3. In Chapter 2, I explain how providing transparency and control does not sufficiently help users in making better privacy decisions. Specifically, I demonstrate that people's privacy decisions fall prey to all sorts of decision biases, and that most privacy decisions are too complex for people to fathom. In effect, many people refrain altogether from exploiting provided transparency and control.

In Chapter 3, I explain how "privacy nudging" is also not sufficient in its presently studied form. Specifically, I demonstrate that although nudges relieve some of the burden

of privacy decision making, they tend to overlook the inherent diversity of users' privacy preferences and the context-dependency of their decisions.

The central thesis of this dissertation is that because of these shortcomings of transparency-and-control and privacy nudges, privacy scholars need to move beyond the "one-size-fits-all" approach that is embodied in both nudges and transparency and control. I argue that because of the high variability and context-dependency of people's privacy decisions, nudges need to be *tailored* to the user and her context.

In several studies, I contextualize users' privacy decisions by showing how disclosure depends on the person's privacy profile, the type of information, and the recipient of the information (Chapter 4). Then, I present the idea of a "Privacy Adaptation Procedure" and demonstrate its merit in Chapter 5. Finally, I test a complete implementation of the Privacy Adaptation Procedure in Chapter 6. The results of this final study give rise to reserved optimism regarding the feasibility of user-tailored privacy decision support.



## **CHAPTER 1: Motivation**

Privacy issues are an undying obstacle to the adoption of social and mobile technologies. Privacy concerns have been identified as an important barrier to the growth of social networks (Stutzman and Kramer-Duffield 2010; Tufekci 2008), e-commerce (Acquisti 2004), ubiquitous computing (Spiekermann and Berthold 2005), and location sharing services (Li and Chen 2010; Page et al. 2012). From the user's perspective, online privacy can be defined as the decision to share certain information with another person, a company, or with the general public. In such decisions, users face a dilemma: they want to enjoy the benefits that may result from sharing or disclosing information, but they also want to reduce the risk that this data may be used in unintended and unwanted ways (Mabley 2000; Taylor et al. 2009). For example, people are willing to provide personal information for monetary rewards or discounts (Hann et al. 2002; Xu et al. 2009), social benefits (Keith et al. 2011; Krasnova et al. 2010; Lee et al. 2008; Thambusamy et al. 2010; Youn 2009), or a personalized experience (Accenture 2012; Awad and Krishnan 2006; Chellappa and Sin 2005; Sheng et al. 2008; Xu et al. 2011), but they also fear security breaches (Fjermestad and Nicholas Romano 2009; Jarupunphol and Mitchell 2002), fraud (Andersen 2000; Cranor et al. 1999; Fox and Lewis 2001), identity theft (Fogel and Nehmad 2009; Nosko et al. 2010) and leaving a negative impression on others (Binder et al. 2009; Page et al. 2013; Tufekci 2008).

Most Internet users take a pragmatic stance on information disclosure (Harris 2000; Harris et al. 2003a; Westin et al. 1981; Westin and Maurici 1998). They trade off the anticipated benefits with the risks of disclosure (Mabley 2000; Taylor et al. 2009). This

decision process has been dubbed *privacy calculus* (Culnan 1993; Laufer and Wolfe 1977). In making this trade-off, these users typically decide to disclose *some* but not *all* information that is requested from them. Indeed, numerous privacy surveys demonstrate that Internet users want to limit the collection and dissemination of their personal data (Ackerman et al. 1999; Phelps et al. 2000; Schrammel et al. 2009; Sheehan and Hoy 2000; Staddon et al. 2012; Zhao et al. 2012).

As past research has shown (Acquisti 2004; Acquisti and Grossklags 2008), privacy-decisions are inherently difficult, because they have delayed and uncertain repercussions that are difficult to trade off with the possible immediate gratification of disclosure. The main research question I propose to answer in my dissertation is thus:

**How can we help users to balance the benefits and risks of information disclosure in a user-friendly manner, so that they can make good privacy decisions?**

Prior research has explored two approaches to this problem, and neither of them provides a satisfying solution. Providing transparency and control (Chapter 2) puts users in charge of their own privacy decisions, but privacy decision-making is often not rational, and transparency and control may only increase its difficulty. Privacy nudges (Chapter 3) relieve some of the burden by providing succinct justifications and sensible defaults, but they tend to overlook the inherent diversity of users' privacy preferences and the context-dependency of their decisions.

To overcome these problems, we need a more fundamental understanding of the reasoning behind privacy decisions in various contexts. I therefore conducted a number of studies to contextualize the privacy calculus (Chapter 4). Based on this deeper understanding of users' privacy decision behavior, **the core contribution of this dissertation is a Privacy Adaptation Procedure that offers *tailored* privacy decision support.** This procedure gives users personalized nudges and personalized justifications based on a context-aware prediction of their privacy preferences. Examples of the Privacy Adaptation Procedure are discussed in Chapter 5, and a complete implementation of it is tested in Chapter 6. Finally, I discuss some practical implications and future work in Chapter 7.

## **CHAPTER 2: Problems with transparency and control**

### **2.1 Privacy calculus**

Laufer and Wolfe (Laufer et al. 1973; 1977) coined the term “calculus of behavior” to refer to users’ conscious process behind their information disclosure decisions. Several researchers have since used the term “privacy calculus” to investigate antecedents of information disclosure (Culnan and Armstrong 1999; Culnan and Bies 2003; Dinev and Hart 2006; Hann et al. 2007; Keith et al. 2011; Li et al. 2010; Milne and Gordon 1993; Petronio 2002; Wilson and Valacich 2012; Xu et al. 2009, 2011), and it has become a well-established concept in privacy research (Li 2012; Pavlou 2011; Smith et al. 2011).

Li (2012) argues that the privacy calculus can be seen as a privacy-specific instance of decision-making theories like the utility maximization or expectancy-value theory (Awad and Krishnan 2006; Rust et al. 2002; Stone and Stone 1990). The expectancy-value theory states that people gather information about various aspects of each choice option, and assign a value to each of these aspects (Fishbein and Ajzen 1975). Utility maximization, in turn, states that people will trade off the different aspects and then choose the option that maximizes their utility (Bettman et al. 1998; Simon 1959).

What are the aspects that people trade off in privacy decisions? Two aspects are mentioned repeatedly in existing work: perceived risk and perceived relevance.

Featherman and Pavlou define privacy risk as the “potential loss of control over personal information, such as when information about you is used without your knowledge or permission” (Featherman and Pavlou 2003, p. 1036). This loss of control can lead to unintended uses and distribution of the information (Olivero and Lunt 2004; Sheehan and

Hoy 2000; Van Slyke et al. 2006). The *perception* of risk is the fear that these unintended consequences will happen (Jacoby and Kaplan 1972; Li 2012). In this sense, perceived risk can be seen as *contextualized* privacy concerns: concerns about the possible consequences of disclosing a *specific* piece of information to a *specific* recipient (Culnan and Bies 2003; Malhotra et al. 2004; Phelps et al. 2000; Smith et al. 1996).

Risk perceptions lead us to restrict access to our personal information (Li and Santhanam 2008; Petronio 2002). In fact, surveys have found that between 58.2% (Metzger 2007) and 72% (Hoffman et al. 1999) of all respondents cite risk as a reason not to disclose their personal information. Comparing effect sizes between studies, Dinev & Hart (2006) note that privacy risk may even be more likely to dissuade people from making an e-commerce transaction than the economic risk of the transaction (see also Bhatnagar et al. 2000). White therefore argues that “Marketers’ efforts may be wisely directed at attempts to mitigate any perceived “downside risks” associated with disclosure.” (White 2004, p. 43).

Several studies found a direct effect of perceived risk on disclosure intentions (Li et al. 2010, 2011; Norberg et al. 2007), while others believe this effect to be (partially) mediated by privacy concerns (Youn 2009) or trust (Dinev et al. 2006; Dinev and Hart 2006). Still others reverse the relationship between risk and concerns/trust, and find that risk is a mediator between concerns or trust and disclosure intentions (Malhotra et al. 2004; Van Slyke et al. 2006; Xu et al. 2005; Zhou 2012). The relationship between perceived risk, concerns, and trust is thus not entirely clear, but the relationship between perceived risk and disclosure—mediated or not—is strong and consistent. Perceived privacy risk may even have longer-term effects beyond disclosure; it may influence users’

intention to transact in a web shop (Kim et al. 2008; Pavlou 2003), or their intention to adopt an online service (Featherman and Pavlou 2003).

Whereas perceived risk describes the negative side of the privacy calculus, the positive side appears to be governed by the *perceived relevance* of the request. Just like perceived risk can be seen as contextualized privacy concerns, perceived relevance can be seen as *contextualized* benefit: the perceived benefit of disclosing a *specific* piece of information to a *specific* recipient (Li et al. 2011). Stone (1981) was the first to consider the effect of the perceived relevance of information requests on privacy-related behaviors, and this effect has since been demonstrated empirically (Li et al. 2011). Phelps et al. (2000) note that people's purchase intentions go down when a service requests information that does not serve the purpose of the request. They therefore argue that "marketers need to resist asking for such information in situations in which the relevance is not readily apparent" (Phelps et al. 2000, p. 38).

## **2.2 Transparency and control**

To help users with their privacy calculus, privacy experts recommend the practice of "transparency and control" or "informed consent": giving users comprehensive control over what data they wish to share, while at the same time providing them with more information about the implications of their decisions (Acquisti and Gross 2006; Benisch et al. 2011; Brodie et al. 2004; Egelman et al. 2009; Hui et al. 2007; Kolter and Pernul 2009; Metzger 2006; Rifon et al. 2005; Tang et al. 2010, 2012; Toch et al. 2010; Wenning and Schunter 2006; Xu 2007; Xu et al. 2009). Transparency and control are also at the heart of existing or planned regulatory schemes. These state that people should be educated about

the rationale and impact of the privacy decisions they are supposed to make, so that they can make these decisions in an unambiguous manner (EU 2012; White House 2012).

At least some minimum level of control over one's disclosure is necessary to engage in a privacy calculus: without control, the user does not have any influence on the risk/benefit tradeoff. Moreover, people can only make an informed tradeoff between benefits and risks if they are given adequate information. Information enables them to make an accurate assessment of the possible risks and benefits of disclosure. Based on this reasoning, advocates of transparency and control argue that it *empowers* users to regulate their privacy at the desired level (Bulgurcu 2012; Cavusoglu et al. 2013; Lederer et al. 2004; Sadeh et al. 2009; Taylor et al. 2009; Xu et al. 2012).

However, research in the past few years has unveiled a fair number of “privacy paradoxes”: situations or conditions in which transparency and control do *not* increase people's privacy, or even *decrease* it. Studies have also shown that people are not very rational decision makers with regard to privacy, as the privacy calculus theory tacitly takes for granted. I will describe these paradoxes in more detail below.

### **2.3 Ironical effects of transparency**

Although most researchers claim that users should be informed about the rationale behind information requests and the possible risks and benefits of disclosure (Egelman et al. 2009; Hui et al. 2007; Metzger 2006; Rifon et al. 2005; Xu et al. 2009), the reality is that doing so often makes users simply more fearful or unwilling to come to a decision. This results in an ironic effect of transparency: When sites inform users about the practices they employ to reduce the risks of disclosure, this information can have the opposite effect, and

make users *more* rather than less wary about their privacy. For example, marketers (Aagaard 2013; Bustos 2012; Gardner 2012) have discovered that displaying a privacy label on an e-commerce website—a supposed vote of confidence in the site’s privacy practices—may *decrease* instead of increase purchases. Similarly, privacy policies have been shown to incite privacy concerns rather than easing them (Pollach 2007).

Moreover, John et al. (2011) demonstrate that even subtle privacy-minded designs and information may trigger users’ privacy fears and thereby reduce disclosure and participation rather than increasing it. They found that a professional looking site garners higher privacy concerns than an unofficial and unprofessional looking site, because the former design reminds users of privacy. In other words: transparency does not make people more discerning about their privacy decisions, but merely makes them worry about privacy in general.

Finally, Adjerid et al. (2013) show that the impact of privacy notices depends on their specific presentation. Particularly, they find that the framing (see also Section 2.4) of privacy notices as either increasing or decreasing compared to previous levels of privacy (even though the current level is the same) increases or decreases disclosure behavior, respectively. Moreover, distractions and misdirections (which often happen in real-world online settings) can easily nullify any effect of privacy notices. Adjerid et al. conclude that transparency thus offers nothing more than a “sleight of privacy”.

## **2.4 Bounded rationality and decision biases**

Transparency and control only work when people are rational decision makers who will use the provided information and controls to their best advantage. The earliest work



on privacy indeed took this approach (Posner 1981; Stigler 1980), but like many decisions, users' privacy decisions turn out not to be particularly rational (Acquisti and Grossklags 2005, 2008).

Users' bounded rationality is demonstrated in the numerous decision fallacies that have been identified through empirical privacy research. One of these decision fallacies is the "herding effect" uncovered by Acquisti et al. (2012): people sometimes blindly follow others in their privacy decisions. They also demonstrate an "order effect": asking privacy-sensitive questions in a decreasing order of intrusiveness could increase overall levels of disclosure, because subsequent requests compare favorably to the previous more intrusive requests, and users will therefore be more likely to answer them positively (this is called the "door in the face" technique, cf. Cialdini et al. (1975)).

John et al. (2011) find a "default effect": people disclose more information when it is disclosed-by-default (opt-out) rather than withheld-by-default (opt-in). This default effect has been explained, behaviorally, cognitively, and socially. Behaviorally, disclosure takes less effort in a disclosed-by-default interface, and consumers consequently end up disclosing more information in this setting (Johnson and Goldstein 2003; Samuelson and Zeckhauser 1988). Cognitively, a disclosed-by-default interface puts users in a "reject frame" (i.e., they have to think of reasons for not sharing the information), whereas a private-by-default interface puts users in an "accept frame" (i.e., they have to think of reasons to share the information). Consequently, disclosure is lower in the private-by-default case, because decision-makers need to feel more committed to make an "accept" decision than to forego a "reject" decision (Ganzach 1995; Meloy and Russo 2004; Wedell

1997). Socially, defaults act as an implied endorsement of the default value by the system (McKenzie et al. 2006).

In a study on mailing list signup requests Johnson et al. (2002) explore this “default effect” in more detail, and disambiguate it from a “framing effect”: this effect states that people disclose more information when the request is framed positively rather than negatively. So whereas the default effect is about the *action* users have to take to attain a certain outcome (opt-in: action is required to disclose; opt-out: action is required to keep private), the framing effect is about the *phrasing* of the request. Johnson et al. confirm both effects, and find that twice as many people sign up in the positively phrased opt-out condition (a pre-checked checkbox labeled “[x] Notify me about more health surveys”; 89.2%) versus a negatively phrased opt-in condition (a pre-checked checkbox labeled “[x] Do not notify me about more health surveys”; 44.2%). Lai and Hui (2006) find similar effects (52.6% and 0.0% signups, respectively), but also note that these differences are bigger for people with low privacy concerns (69% vs. 0.0%) than for people with high privacy concerns (18% vs. 0.0%).

In a broader sense, the default effect may be due to the “endowment effect”: people are usually less willing to give up something they already have than they are willing to pay for acquiring something they do not have (Kahneman et al. 1990; Thaler 1980). Both Acquisti et al. (2009) and Tsai et al. (2010) show that people are indeed less willing to pay for gaining privacy than what they would demand to give it up. This may be the main reason why explicit monetary rewards seem to have varying effects on disclosure. Hui et al. (2007) find that participants are proportionally more willing to fill out a marketing survey with increasing monetary rewards ranging from \$0.60 to \$5.40. In a study on a location-

based coupon service, Xu et al. (2009) find that a rebate of \$0.20 on the monthly phone bill increases disclosure only when the system pushes the coupons to the user. However, when studying information disclosure in an online fax service, Li et al. (2010) find an “undermining effect of rewards” (p. 21) when users do not perceive the requested information to be relevant to the purpose of the e-commerce transaction. It has no effect when the information is perceived as relevant to begin with.

Self-efficacy is a determinant of privacy concerns (Cho 2010; Cho et al. 2009; Larose and Rifon 2007; Mohamed and Ahmad 2012; Yao et al. 2007; Youn 2009), but the aforementioned results suggest that people may lack the level of self-efficacy required to take control over their disclosure decisions. Ironically, Brandimarte et al. (2013) show that increased control can result in “misplaced confidence”, increasing disclosure even when the control provides no additional protection. They asked people for their permission to display the information they just disclosed on a new social network. The network would either *certainly* include their data (high perceived control) or only include the data of a randomly selected subset of participants (low perceived control). Users were more likely to agree with the former, even though chances of disclosure were obviously higher. Users may thus end up unwittingly disclosing more information merely because they perceive more control. This leaves them more vulnerable as a result of measures ostensibly meant to protect them. These findings are in line with existing work that shows that confidence often promotes complacency, rather than effective goal-relevant behavior such as self-protection behavior (Weinstein 1989).

## 2.5 Context effects: an example of decision biases (original work<sup>1</sup>)

We conducted a study in the field of location sharing to demonstrate some prevalent decision biases in people's privacy decision-making behavior. This paper investigated the effects of providing users more versus fewer options to set the settings of an app that can share their location with their friends, colleagues, a coupon system and certain apps. Specifically, we looked at *what* users could share (options: nothing, city, city block, exact location), and manipulated the presence of an option in the middle of the sequence (city) and an option at the endpoint of the sequence (exact location).

Regarding the presence of the middle option, some researchers suggest that users who would normally go for this option will “err on the safe side” when their preferred option is omitted, and rather not share their location at all (Benisch et al. 2011). Several researchers use this point to argue that finer-grained options (i.e., more control) are necessary to reduce their privacy concerns (Consolvo et al. 2005) and increase their level of location-sharing (Bokhove et al. 2012; Goncalves et al. 2012; Prasad 2012). In a similar vein, the absence or presence of the endpoint option (exact location) should only affect users who would normally choose the next option down (city block). Any user choosing one of the other options (none or city) should not be influenced by the availability of the exact location option, because if they wanted to share more they could have chosen the city block option.

---

<sup>1</sup> Sections titled “original work” describe studies conducted by the author as part of this dissertation. The work is summarized in sufficient detail to support the argument of the dissertation. A note at the end of such sections refers to any publications that discuss the results in more detail. As this research is conducted with the help of co-authors, these sections are written in the “we” form. The author is the lead researcher on all original work though, and was responsible for designing, conducting, analyzing and reporting the presented studies.

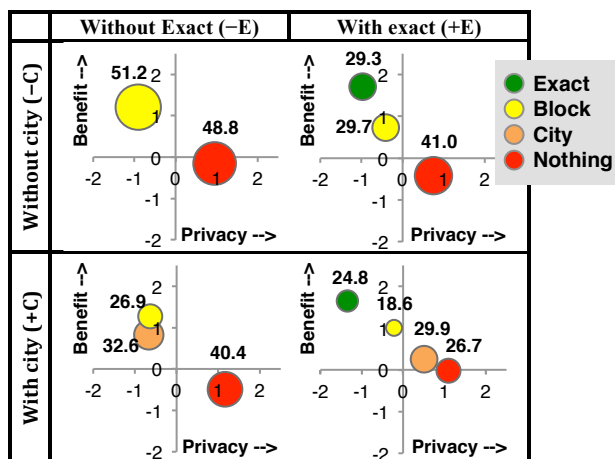
The results of our study (N=291) showed that users did not behave accordingly, though. Instead, users' behavior fell prey to two well-known decision biases:

- Tversky's "substitution effect" (Huber and Puto 1983; Tversky 1972): if option A is subjectively close to option B, then A serves as a *substitute* for B. In such a case, the absence of A will mainly increase the share of option B, and not so much the share of another option C. In other words: what users choose instead depends on their perception of the remaining options.
- Simonson's "compromise effect" (Simonson 1989): Given options B and C, if option A is on the other side of B than C, and far enough removed from B (i.e., a "distant competitor") then the presence of A will turn B into a *compromise* between A and C, and in turn, increase the share of B to the expense of C. Combining this with the substitution effect, the presence of an outlying option can increase sharing across the board.

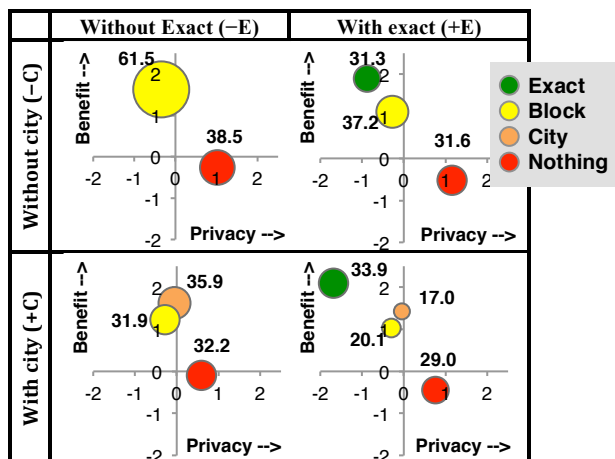
To measure the subjective distances between sharing options, we asked participants for a subset of their decisions how they perceived the benefits and privacy of the chosen option (on a 7-point scale), averaged these ratings as *part-worth utilities* (Huber 1974; Srinivasan 1988), and then mapped these averages onto a two-dimensional plane in Table 1 (for sharing with friends/colleagues) and Table 2 (for coupons/ apps). The different colored "bubbles" represent the different sharing options. The size of the bubbles, as well as their labels, represents the relative number of times each option was chosen. The position reflects the average perceived Privacy (horizontal axis) and Benefit (vertical axis) of each option. The top two panels in the tables represent the Without City (-C) conditions, while the bottom two panels represent the With City (+C) conditions. Similarly, the left two

panels represent the Without Exact (−E) conditions, while the right two panels represent the With Exact (+E) conditions.

**Table 1: Subjective position and choice proportion (size and label) of each option, for sharing with friends/colleagues.**



**Table 2: Subjective position and choice proportion (size and label) of each option, for sharing with coupons/apps.**



These graphs show that when a finer-grained sharing option is removed, users do not just “err on the safe side”, but instead deliberately choose the *subjectively closest* remaining option, as suggested by Tversky’s substitution effect.

Moreover, if an “extreme” option is introduced that is sufficiently distinct from the existing options, this not only causes some users to switch from the previously most

extreme option to this new option, but it also causes some users to switch from a less extreme option to the previously most extreme option, as suggested by Simonson's compromise effect. In other words, such an extreme option may increase sharing across the board.

Most importantly, like in all decisions, users seem to have no *fixed preference* for location sharing settings (Bettman et al. 1998; Coupey et al. 1998). Instead, their decisions depend on the other available options, and specifically on their *perception* of the differences between these options. These perceptions are easily influenced, and this is what causes most of the numerous decision biases discussed in Section 2.4.

*Refer to Knijnenburg et al. (2013a) for expanded analysis and discussion.*

## **2.6 Control and Transparency Paradox**

The Control Paradox states that while users claim to *want* full control over their data (Acquisti and Gross 2006; Benisch et al. 2011; Brodie et al. 2004; Kolter and Pernul 2009; Tang et al. 2010, 2012; Toch et al. 2010; Wenning and Schunter 2006; Xu 2007), they avoid the hassle of actually *exploiting* this control (Compañó and Lusoli 2010). In combination with overly permissive defaults (Bonneau and Preibusch 2010; Gross and Acquisti 2005), this leads to a predominance of over-sharing.

Recent work backs up the argument that people may simply not be motivated to make the effort to take control over their data. For example, Larose and Rifon (2007) find that privacy seals influence disclosure tendencies, but only for participants that are either motivated or have a high self-efficacy. Similarly, Besmer et al. (2010) employ social navigation cues to influence users' information disclosure on Facebook, and they too find

that participants were only influenced if they already had a tendency to change their settings. Finally, in a study of users' Facebook privacy settings, Gross and Acquisti (2005) conclude that "only a small number of members change the default privacy preferences, which are set to maximize the visibility of users profiles" (p. 79).

Similarly, Nissenbaum (2011) postulates the Transparency Paradox: privacy notices that are sufficiently detailed to have an impact on people's privacy decisions are often too long, detailed and complex for people to read. Indeed, while many people claim to read online privacy policies (Internet Society 2012; Milne and Culnan 2004), many do not actually read them (Adkinson et al. 2002; Berendt et al. 2005; Bergmann 2009; Harris 2001; Jensen et al. 2005; Kelley et al. 2010; Singleton and Harper 2002; Turner and Varghese 2002), or do not read closely enough to understand them (Pan and Zinkhan 2006).

## **2.7 Form auto-completion tools and the control paradox (original work)**

We investigated a prominent example of a technology that may inadvertently introduce the control paradox: web form auto-completion tools. Most modern browsers have an auto-completion feature that reduces the burden of information disclosure by reducing the required amount of typing, and by helping users to recall the correct information (Bicakci et al. 2011; Trewin 2006). Usability experts therefore recommend to web developers to build their forms in a way that enables them to be recognized by these auto-completion aids (Garrido et al. 2011; Wroblewski 2008). There even exist a number of third-party tools such as RoboForm, LastPass, and Dashlane, which provide more comprehensive (e.g., cross-device) auto-completion features.



Despite the apparent benefits of auto-completion, Preibusch et al. (2012) warn that since these tools typically fill *all* the fields on a form (even optional fields), they could increase the risk of over-disclosure due to a strong default effect. Moreover, auto-completion tools could counteract the privacy calculus: they make it so easy to submit a fully completed form that users may skip weighing the benefits and risks of disclosing a certain piece of information in a specific situation. Consequently, their information disclosure may no longer be *purpose-specific*.

To test this conjecture, we conducted an online user experiment (N=460) testing a mock-up of a traditional auto-completion tool against mock-ups of two alternative tools designed to reinstate users' privacy calculus and purpose-specific disclosure behavior. The three tools that we compared were:

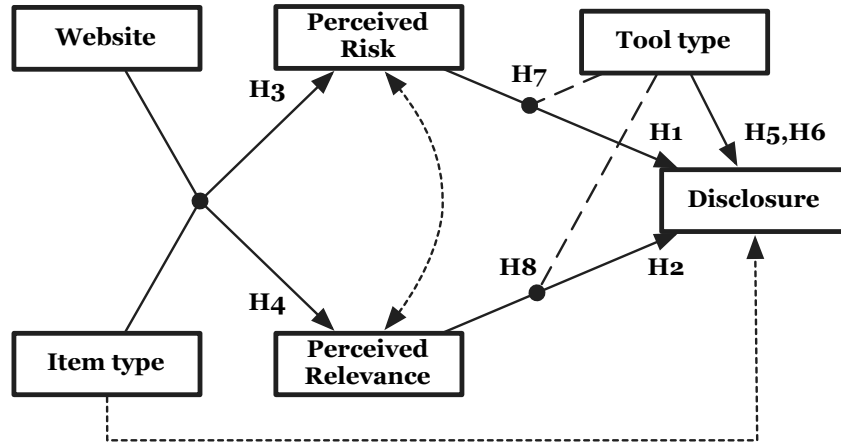
1. A traditional auto-completion tool that automatically fills all fields (Auto);
2. A tool that fills out the entire form like the traditional tool but features "remove" buttons next to each field so that users can easily remove individual entries (Remove);
3. A tool that leaves the form empty by default, but has "add" buttons to easily fill individual fields (Add).

The hypotheses of this study are outlined in Figure 1<sup>2</sup>. Hypotheses 1-4 concern the purpose-specificity of information disclosure, mediated by perceived risk and relevance.

---

<sup>2</sup> Somewhat surprisingly, H5 (arguing that disclosure would be higher for users of the Remove tool than for users of the Add tool) did not reach significance. Moreover, the dashed lines were added post hoc as effects that significantly improved model fit: a residual correlation between Perceived Risk and Perceived Relevance (accounting for the fact that items perceived as irrelevant are often also perceived as risky), and a main effect of Type of information on Disclosure (accounting for an inherent difference in disclosure tendency between these types of information that goes beyond perceived risk and relevance).

These hypotheses argue that perceived risk is lower (H3) and perceived relevance is higher (H4) when the type of requested information matches the purpose of the website, and that this in turn increases disclosure (H1-H2). Hypotheses 6-8 argue that users of the traditional auto-completion tool disclose more items (H6) and are less likely to consider perceived risk (H7) and relevance (H8) in their disclosure decision.



**Figure 1: Experimental model. H7-H8 describe a moderating effect on H1-H2.**

After being randomly assigned to one of the three tools, participants provided the tool with a wide range of personal information (general contact information, personal interests, job skills, and health record), and then “tested” the tool on a randomly selected external website (also mock-ups) which requested some of the information participants had provided to the auto-completion tool. Each of the three websites presented some kind of personalized service, and each was chosen to correspond to a particular subset of the personal information requested by the auto-completion tool: A blogging community matched personal interest items, a job search website matched job skills items, and a health insurer matched health record items. However, in our experiment these websites did not

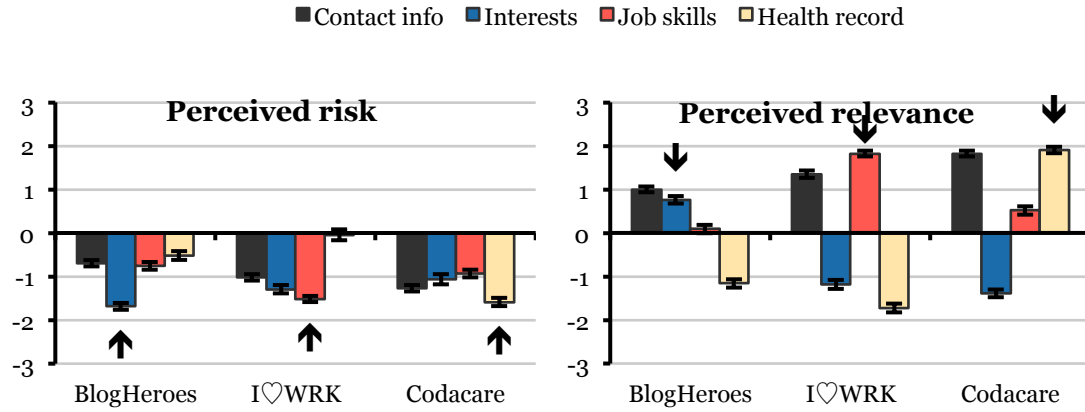
just ask for these “matching” items, but also for the items that did not clearly match the purpose of the website (e.g., health record items requested by the blogging community). Requesting both matching and non-matching items on the website provided a within-subjects manipulation that allowed us to measure the purpose specificity of disclosure in each of the three auto-completion tools.

Finally, participants were transferred back to the FormFiller website, where they would evaluate their satisfaction with FormFiller (a questionnaire adapted from (Knijnenburg et al. 2012c)). They also indicated their reasons for disclosing or not disclosing each requested item, specifically in terms of perceived risk (i.e., the statement “Providing [item] to [website] is:” rated on a 7-point scale from “very safe” to “very risky”) and perceived relevance (i.e., the statement “The fact that [website] asked for [item] was:” rated on a 7-point scale from “very inappropriate” to “very appropriate”).

In general (i.e., across all three conditions), we found significant effects of perceived risk (odds ratio: 0.818,  $p < .001$ ) and perceived relevance (odds ratio: 1.079,  $p < .001$ ) on disclosure, which is in line with hypotheses 1 and 2. Figure 2 shows the results for Hypotheses 3 and 4; risk is indeed lowest—and relevance is indeed highest<sup>3</sup>—when the type of information requested matches the purpose of the website. Finally, we find that the effect of perceived risk and relevance on disclosure is lower in the Auto condition (odds ratios: 0.863 and 0.989) than in the Remove condition (odds ratios: 0.754 and 1.133) and the Add condition (odds ratios: 0.811 and 1.114).

---

<sup>3</sup> We here ignore the “Contact Info” information type, because it does not match particularly stronger with any of the individual websites.




**Figure 2: Perceived Risk and Perceived Relevance per Website and Item type. The arrows point to the matching item types. Error bars are  $\pm 1$  Standard Error.**

This means that users of traditional auto-completion tools are less likely to consider perceived risk and relevance in their decision to disclose, and they thus neglect the purpose specificity of the requested information. Arguably, they suffer from the control paradox: while they have the opportunity to control their disclosure (by manually changing or removing field entries), they avoid the hassle of doing so. Our results show that the alternative tools overcome the control paradox: they indeed help users consider the perceived risk and relevance of each item in their decision making process. Consequently, users of these tools make decisions that are more purpose-specific. This can be confirmed by comparing the disclosure decisions of users in the Auto tool and the Remove tool in Figure 3. Note also that the average level of disclosure is about the same between the Add and Remove tool. This means that there is no default effect (see Section 2.4); another indication that users in these conditions more carefully decide what to disclose.

In a follow-up study with 290 participants we examined the underlying mechanisms that govern how the different tools influence users to make more heuristic versus more deliberate privacy decisions. In this study we were able to demonstrate that the Add and

Remove tools increase users' perceived elaboration self-efficacy, an important precondition for users to take deliberate control over their privacy decisions (Angst and Agarwal 2009; Lowry et al. 2012).



**Figure 3: Disclosure rate per Website and Item type, for each of the three tools. The arrows point to the matching item types. Error bars are  $\pm 1$  Standard Error.**

*Refer to Knijnenburg et al. (2013b) and Knijnenburg and Bulgurcu (2015) for expanded analysis and discussion.*

## 2.8 Conclusion: Transparency and control do not work

Transparency and control do not work well in practice, especially for systems that process large amounts of personal data, which is increasingly the case online (Zickuhr 2012). Providing transparency and control in such systems is simply unwieldy. Indeed, the complexity of online privacy policies is ever-increasing (Milne et al. 2006): they are often written in a legalistic and confusing manner, and require a college reading level to understand them (Antón et al. 2004; Cate 2006; Kelley et al. 2010; McDonald et al. 2009; Turow et al. 2005). Moreover, systems like Facebook that manage large amounts of personal user data have to resort to “labyrinthian” privacy controls (Consumer Reports

2012). As a result most Facebook users do not seem to know the implications of their own privacy settings (Liu et al. 2011; Strater and Lipford 2008), and share postings in a manner that is often inconsistent with their own disclosure intentions (Madejski et al. 2012).

Due to the complexity of privacy decisions and users' bounded rationality, an increase in transparency and control often just aggravates the problem by introducing choice overload and information overload. Consequently, several scholars have recently questioned the effectiveness of the "transparency and control" paradigm. Specifically, Barocas and Nissenbaum (2009) argue that notice and control are a "red herring", because it is difficult for people to fathom all the information needed to make a conscious decision. Nissenbaum (2011) argues that "transparency-and-choice has failed" (p. 34) because detailed descriptions of privacy policies are impossible for people to understand, while simplified versions of these policies inevitably drain away details that are likely to make a difference in their decision. Solove (2013) argues that the paradigm does not "provide people with meaningful control over their data" (p. 1880) because of our cognitive limitations and the overwhelming number of entities collecting our personal information. Similarly, the U.S. President's Council of Advisors on Science and Technology argues that notice and consent fail because "it is simply too complicated for the individual to make fine-grained choices for every new situation or app" (PCAST 2014 p. 38)

## CHAPTER 3: Problems with privacy nudges

### 3.1 Privacy nudges

Instead of raising the burden on users by increasing transparency and control, solutions should *relieve* some of this burden by making it easier for users to process and execute the information disclosure decisions, without taking away users' control altogether. A recent approach to support privacy decisions that does this is *privacy nudging*. Nudges are subtle yet persuasive cues that makes people more likely to decide in one direction or the other (Thaler and Sunstein 2008). Carefully designed nudges make it easier for people to make the right choice, without limiting their ability to choose freely. Nudges ostensibly turn people's decision fallacies into mechanisms that help them (Acquisti 2009): they exploit these fallacies to create a *choice architecture* that encourages wanted behavior and inhibits unwanted behavior (Thaler and Sunstein 2008).

The type of nudge that is most extensively implemented in real systems is *justifications*. A justification is a succinct reason to disclose or not disclose a certain piece of information. Justifications are a nudge, because they make it easier to rationalize the decision (Bettman et al. 1998; Simonson 1989) and to minimize the regret associated with choosing the wrong option (Connolly and Zeelenberg 2002; Inman and Zeelenberg 2002). Justifications include providing a reason for requesting the information (Consolvo et al. 2005), highlighting the benefits of disclosure (Kobsa and Teltzrow 2005; Wang and Benbasat 2007), and appealing to the social norm (Acquisti et al. 2012; Besmer et al. 2010; Patil et al. 2011). The effect of such justifications seems to vary. In the study of Kobsa and Teltzrow (2005), users were about 8.3% more likely to disclose information when they

knew the benefits of disclosing the information. In an experiment by Acquisti et al. (2012), they were about 27% more likely to do this when they learned that many others decided to disclose the same information. However, Besmer et al. (2010) find that social cues have barely any effect on users' Facebook privacy settings: only the small subset of users who take the time to customize their settings may be influenced by strong negative social cues. Similarly, Patil et al. (2011) rate social navigation cues as a secondary effect.

Another justification strategy is to provide a privacy indicator or seal. This indicator or seal is a symbolic representation of a judgment about privacy, often made by an authoritative third party. Egelman et al. (2009) show that privacy indicators next to search results can entice users to pay a premium to vendors with higher privacy scores. In their study, participants paid about \$0.15 extra for a pack of batteries and about \$0.40 for a sex toy (on top of a \$15.50 average base price). Users of Xu et al.'s (2009) location-based coupon service were more likely to disclose information when the site displayed either a TRUSTe seal or a legal statement, with the seal working best. In Hui et al.'s (2007) marketing survey, however, a privacy seal did not significantly increase disclosure. Studying an online CD retailer, Metzger (2006) also found that their seal had no effect. Rifon and Larose (2005) show that warnings and seals at an online retailer website influence users in certain situations only.

In social media, justifications often relate to the real or potential audience of a shared piece of information. For example, in location sharing services, researchers have experimented with giving users real-time feedback on who is requesting or viewing their location (Jedrzejczyk et al. 2010; Tsai et al. 2009). The results of these experiments are mixed: users appreciate the information, but it can easily become excessive and annoying.



Similarly, Wang et al. (2013, 2014) implemented and tested a tool that provides users with detailed feedback about the potential audience when posting a Facebook message. They find that at least some users consider this tool helpful, but they find no significant differences in posting behavior.

Wang et al. (2013, 2014) consider two other types of nudges: sentiment feedback and a post timer. Sentiment feedback analyzes the tone of a message the user is about to post, and tells the user whether the message is likely to be perceived as positive or negative. The post timer delays Facebook posts by 10 seconds, which allows users to change their mind. While some of the participants in their study seemed to like these tools, others found them intrusive and annoying.

Another approach to nudging users' privacy decisions is to provide sensible defaults. Defaults (partially) relieve users from the burden of making information disclosure decisions by offering a path of least resistance: Correctly chosen defaults make it easier to choose the right action, or may not even require any action at all. Defaults also provide an implicit normative cue, e.g., a default order communicates what the system thinks is most important, and a default value communicates what the system thinks you should do. Finally, default values may work due to the 'endowment effect': people are less willing to pay for what they perceive to be a gain in privacy than what they would demand if the same decision were framed as a loss (Acquisti et al. 2009; Tsai et al. 2010).

Because of this, providing a certain default option nudges users in the direction of that default (Thaler and Sunstein 2008). Therefore, while most existing work on default effects in the privacy field (discussed in Section 2.4) regards them as a nuisance, several researchers have recently suggested that they can also be used as nudges (Acquisti 2009;

Adjerid et al. 2013; Balebako et al. 2011). Note, though, that such default-based nudges have not been tested in work other than my own (see Section 5.3). The same holds for the order in which information is requested, which is another variant of default-based nudges.

### **3.2 Failed nudges: Justifications and request order (original work)**

Our own study on nudges comprised a comprehensive evaluation of justifications and a first attempt to use request order as a nudge. The study (N=491) considered the fictitious mobile app recommender system “Applause”, developed by a fictitious company named “Appy”. The system was inspired by existing systems that have been developed both for research and commercial purposes (e.g. Böhmer et al. 2010; Davidsson and Moritz 2011; Girardello and Michahelles 2010, chomp.com). The system used in our study recommends apps for Android phones based on users’ context (e.g., location, app usage, credit card purchases) and demographics (e.g., age, hobbies, religion, household income).

Although context has recently attracted the most attention in mobile recommender systems research (Adomavicius and Tuzhilin 2011; Ricci 2011), several researchers have explored the combination of context and demographics in mobile recommenders (Lee and Park 2007; Lee and Lee 2007; Oh and Moon 2012; Zheng et al. 2012). In general, the users’ context provides a wealth of automatically accrued data that can be used to provide relevant recommendations tailored to the specific usage situation. Demographic information, on the other hand, can be used to overcome the “new user problem” (Lee and Park 2007), and is typically easier to interpret than context data.

In our study participants were first given a short introduction to the mobile app recommender, including two examples of how the system might use their data to provide

context-aware and personalized recommendations. They were then informed that they would be helping Appy to test the information disclosure part of the system. After randomly assigning them to one of 5×2 conditions (see below), participants were ostensibly “transferred” to the Appy website, where they would make 31 information disclosure decisions on 12 pieces of context data and 19 pieces of demographic data. Context requests asked users to indicate whether they would disclose the respective data, and could be answered with a simple ‘yes’ or ‘no’. For demographics requests, users were asked to provide the actual information, or to decline disclosure. All decisions were logged to our database. After 31 decisions, participants were transferred back to the experimenters’ website, where they filled out a questionnaire regarding their:

- Perceived value of disclosure help (3 items, e.g., “The system helped me to make a tradeoff between privacy and usefulness”)
- Perceived privacy threats (3 items, e.g., “The system has too much information about me”)
- Trust in the company (4 items, e.g., “I believe this company is honest when it comes to using the information I provide”)
- Satisfaction with the system (6 items, from (Knijnenburg et al. 2012c); e.g., “Overall, I’m satisfied with the system”)

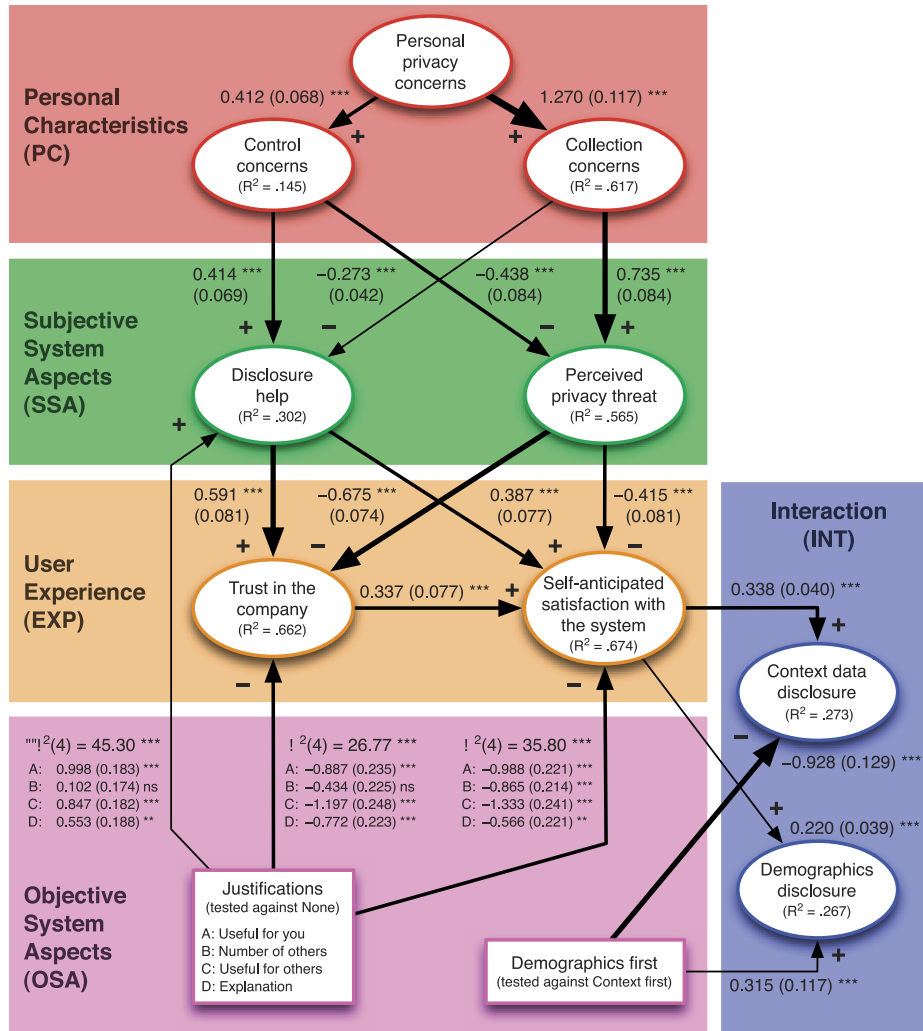
The experiment had two between-subjects manipulations. The first manipulation was the justification: four types of justification messages were tested against the baseline system with no justification messages, bringing the total to five conditions (see Table 3). The percentages in the messages ‘useful for you’, ‘number of others’ and ‘useful for others’

were randomly chosen from 5% to 95% (this percentage had barely any effect). The second manipulation was the request order, which was manipulated as demographical data first or context data first (see Table 3).

**Table 3: Experimental manipulations: strategies to influence information disclosure.**

Manipulation	Conditions	Description
Justification type	None	[Baseline condition with no justifications]
	Useful for you	"The recommendations will be about [XX]% better for you when you tell us/allow us to use..."
	Number of others	"[XX]% of our users told us/allowed us to use..."
	Useful for others	"[XX]% of our users received better recommendations when they told us/let us..."
	Explanation	"We can recommend apps that are [reason for request]"
Request order	Demographical data first	The system first requested the 19 pieces of demographical data, then the 12 pieces of context data.
	Context data first	The system first requested the 12 pieces of context data, then the 19 pieces of demographical data.


The experimental conditions, subjective evaluations, and disclosure behaviors were submitted to a Structural Equation Model that was built iteratively using a split-half method, guided by our validated framework for the user-centric evaluation of recommender systems (Knijnenburg et al. 2012c). The final model (Figure 4) has a good model fit:  $\chi^2(912) = 1540, p < .001$ ;  $RMSEA = 0.037$ , 90% CI: [0.034, 0.041],  $CFI = 0.977$ ,  $TLI = 0.976$ .



**Figure 4: The Structural Equation Model (SEM) for the data of the experiment.**  
**A: Useful for you, B: Number of others, C: Useful for others, D: Explanation.**  
**(Significance levels: \*\*\*  $p < .001$ , \*\*  $p < .01$ , 'ns'  $p > .05$ )**

The model shows that the justifications have a significant impact on perception of disclosure help, trust in the company, and self-anticipated satisfaction with the system. The 'useful for you', 'useful for others' and 'explanation' justifications each significantly increase the perceived value of disclosure help. However, this positive effect is canceled out by a negative effect on trust in the company and on self-anticipated satisfaction with the system. Figure 5 shows that the total (direct plus mediated) effects of the justifications on trust in

the company are essentially zero, and that the total effects on self-anticipated satisfaction with the system and disclosure behavior are negative. In other words, the justifications do not work; in fact they rather *decrease* users' satisfaction and disclosure. Moreover, separate analyses confirmed that justifications decreased disclosure regardless of the percentage in the justification message (except for the 'number of others' justification, but even for that message a high percentage merely reduces the negative effect, and never actually increases disclosure). It may thus be that these justifications fall prey to the ironic effect of transparency (see Section 2.3): by justifying the disclosure, we inadvertently signaled that the act of disclosure is not trivial and may involve risks: the justifications bring the disclosure decision to the foreground and demand users' attention. Consequently these justifications fail as a nudge.



**Figure 5: The total effects of the justifications (A: Useful for you, B: Number of others, C: Useful for others, D: Explanation) on the different outcomes, tested against the baseline condition (No justification). Vertical axes are in sample standard deviations of the measured. Error bars are  $\pm 1$  standard error of the measurement.**

Finally, the request order has a direct impact on disclosure behavior. Basically, disclosure of each type of data is higher when it is requested first. Requesting demographics first increases demographics disclosure but decreases context disclosure, and vice versa. Neither of the request orders thus increases disclosure across the board, but the effect of the request order could still be useful if one type of data is more valuable to the system than the other type. Interestingly, the request order has no effect on user satisfaction, indicating that users may not care whether context or demographics items are requested first.

There are two possible explanations for the effect of request order. One is that users become more wary of privacy threats as the data collected about them accumulates, the other is that users get tired of answering so many disclosure requests. There are reasons to believe that the former explanation holds more ground. If the latter explanation were correct, the effect should be most pronounced for demographics disclosure, because in the presented system it takes more effort to disclose demographic data than context data (since demographics disclosure requires the user to key in the data, whereas context disclosure merely requires users to click a 'yes' or 'no' button). In fact, though, the effect is stronger for context data disclosure than for demographics disclosure. Therefore, tiredness or boredom is likely not the reason for the request order effect.

*Refer to Knijnenburg et al. (2012b) and Knijnenburg and Kobsa (2013a) for expanded analysis and discussion.*

### **3.3 Nudges: a one-size-fits-all approach**

The nudges evaluated in existing work and our own work showed predominantly disappointing results: they usually only worked for some users, and left others unaffected or even dissatisfied with the applied nudge. The problem with the current approaches to privacy nudging is that they take an implicit stance on whether the purpose of the nudge should be to increase disclosure, or to decrease it. System designers may claim that it is always in users' best interests to provide more information, e.g., for personalization purposes (Accenture 2012). They use nudges to increase disclosure, but these nudges may cause the more privacy-minded users to feel 'tricked' into disclosing more information than they would like (Brown and Krishna 2004). Lawmakers may instead believe that privacy is an absolute right that needs to be defended at all costs. Their imposed protective practices may make it more difficult though to disclose information and enjoy the benefits that disclosure would provide. This puts the less privacy-minded individuals at a disadvantage.

Nudges are rarely good for everyone, and some researchers therefore argue that they may threaten consumer autonomy (Smith et al. 2013; Solove 2013). These researchers argue for "smart nudges", such as smart default settings that match the preferences of most users. To implement smart defaults, managers would have to analyze users' privacy settings and make the most common setting the default. Such smart defaults would arguably most closely match the average users' true privacy preferences.

But what if the "average users' privacy preferences" do not exist? To wit, there is ample evidence that people vary extensively in their information disclosure behavior (Harris 2000; Harris et al. 2003b; Westin et al. 1981; Westin and Maurici 1998), and that even for the same person this decision depends on the context in which it is made (Benisch



et al. 2011; Besmer et al. 2010; Consolvo et al. 2005; Hsu 2006; Johnson et al. 2012; Kairam et al. 2012; Lederer et al. 2003; Li et al. 2010; Nissenbaum 2009; Norberg et al. 2007; Olson et al. 2005; Patil and Kobsa 2005; Toch et al. 2010; Watson et al. 2012). Indeed, the variability and context-dependency of privacy preferences is at the core of many privacy theories such as Altman's *privacy regulation theory* (Altman 1975), Nissenbaum's *contextual integrity* (Nissenbaum 2004, 2009), and Petronio's *communication privacy management* (Petronio 1991, 2010).

Fundamentally, then, the current implementations of nudges take a “one-size-fits-all” approach to privacy (Spiekermann et al. 2001a): They assume that the “true cost” (John et al. 2011) of disclosure is roughly the same for every user, for every piece of information, in every situation. I argue that because of the high variability and context-dependency of people's privacy decisions, it would be better to *tailor* nudges to the user and her context (Kobsa 2001; Wang and Kobsa 2007). The core idea of this proposal is thus to adapt privacy decision support to the user (Chapter 5). But this means that I must first “map out” the variability and context-dependency: on what dimensions do people differ in their information disclosure behavior, and which contextual variables influence this decision? This is the topic of Chapter 4.

## CHAPTER 4: Contextualization of privacy decisions

### 4.1 Introduction

A wide variety of contextual variables have been considered in previous research, including demographics and characteristics of the users themselves, the type of information requested, characteristics of the requester, the time and place of the request, and potential interactions between these variables. As the time and place of requests specifically pertains to ubiquitous computing, we limit ourselves to the contextual variables that apply more generically, namely the characteristics of the users themselves, the type of information requested, and the characteristics of the requester. I discuss related work regarding these variables in more detail below.

### 4.2 Types of users and types of information

Existing studies on information disclosure typically either treat users' disclosure of each requested item as an independent observation (Acquisti et al. 2012; Joinson et al. 2008) or they treat them as summated composite scores, essentially considering them to be unidimensional (John et al. 2011; Joinson et al. 2010; Knapp and Kirk 2003; Metzger 2004, 2006, 2007).

A more contextualized approach would attempt to uncover several *dimensions* of personal information, and/or cluster people to create distinct *profiles* of disclosure behavior. Here I discuss a number of studies that have attempted this approach. Since most of these studies have some shortcomings, I subsequently present our own work

establishing the multi-dimensionality of information disclosure behaviors and uncovering user disclosure profiles.

In an experiment with a customer loyalty club at a grocery store, White (2004) finds differences between the disclosure of embarrassing information and contact data, depending on users' relationship with the vendor. The two dimensions are predetermined *ex ante*, and no measures of convergent and discriminant validity are reported.

Phelps, Nowak and Ferrell (2000) ask participants about their willingness to disclose items in five predefined categories (demographic, lifestyle-related, purchase-related, personal identifiers, and financial). They show that participants are more likely to disclose the former three categories than the latter two. In Phelps et al. (2001), the authors define three categories on the same data (lifestyle and shopping, personal financial, and demographic). They show that these groups have high convergent validity, but they do not test discriminant validity.

Buchanan et al. (2007) uncover two dimensions of privacy behavior among 12 items: a general caution dimension and a technical protection dimension. In an interesting investigation of the attitude-behavior link, their attitudinal measure of "Privacy Concern" is correlated with the general caution dimension, but not the technical protection dimension. The opposite was true for the Westin Privacy Score (Westin and Maurici 1998). The IUIPC scale (Malhotra et al. 2004) correlated with both general caution and technical protection.

Koshimizu et al. (2006) consider participants' feelings about a community-based video surveillance system. They find seven factors and three main user profiles. The authors do not provide the statistical fit of their factor model nor a statistical justification for selecting three clusters.

In a large-scale pan-European survey of privacy practices Lusoli et al. (2012) find four factors of personal data disclosed on eCommerce sites: social information, biographical information, sensitive information, and security information. They show that disclosure is quite uniform across European countries, with some differences between northern and southern Europe. They also find six dimensions of protection behaviors: reactive practices (e.g., spam- and spyware filters), proactive practices (e.g., contacting websites about their privacy practices), withholding information, minimizing disclosure, avoiding the use of technology, and lying.

Olson et al. (2005) study disclosure behavior in an interpersonal privacy context. They find six different dimensions among their 40 types of information. However, they find no clear theme among the items that would justify their relatedness and they do not provide statistical evidence for the particular number of dimensions.

De Souza and Dick (2009) measure disclosure behaviors as a single score, classify participants into two clusters (based on attitudes), and then show a difference in behavior between the two clusters. Similarly, Ackerman et al. (1999) contend that participants' levels of disclosure in a generic e-commerce setting differs per uncovered cluster, but that the relative sensitivity of each item is consistent. This suggests that the measured intentions are unidimensional, and that users' disclosures differ in degree but not in kind. These claims are however not directly tested by either De Souza and Dick (2009) or Ackerman et al. (1999).

Spiekermann et al. (2001a) study disclosure behavior in an ecommerce system with an anthropomorphic online shopping bot. They perform similar clustering as in Ackerman et al. (1999), but their resulting four clusters fall onto two attitudinal dimensions: identity

disclosure and profile disclosure. Participants' tendency to disclose their address (identity-related behavior) and their tendency to answer shopping bot questions (profile-related behavior) differ per cluster, and are in line with the attitudinal profile of the cluster.

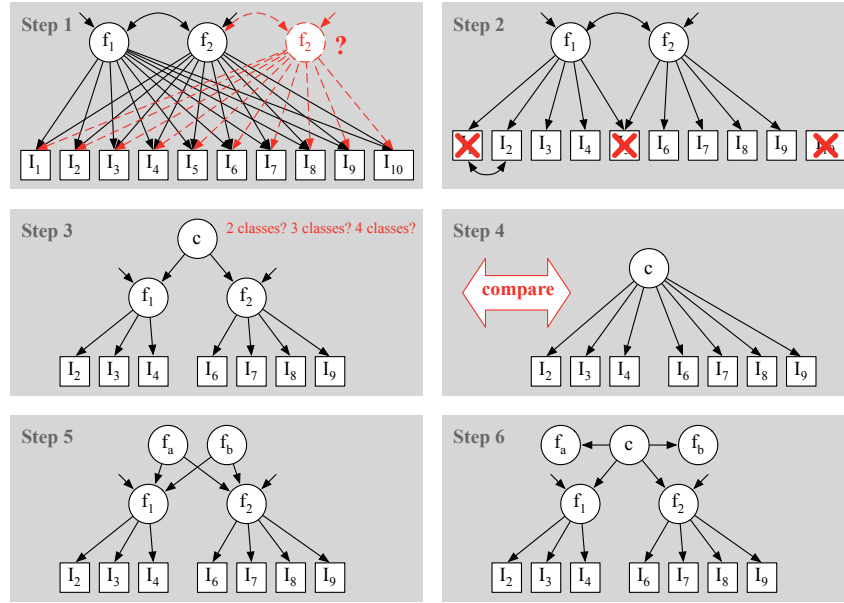
### **4.3 Multi-dimensionality and privacy profiles (original work)**

In our own work on types of users and types of information we evaluated three information disclosure datasets using a six-step statistical analysis (Figure 6). This six-step analysis improves upon existing analyses in the following ways:

- It derives the dimensionality of the behavior from the behavioral data;
- It provides statistical justifications for the chosen number of dimensions;
- It classifies participants into profiles based on their behavior;
- It provides statistical justifications for the selected number of classes.

The first step is to submit each dataset to a series of Exploratory Factor Analyses (EFAs) to discover the inherent dimensionality of the data (i.e., the optimal number of factors). In step 2 we conduct a Confirmatory Factor Analysis (CFA) to create a “clean” factor model. This model describes the dimensionality of the disclosure behavior, and is tested for overall model fit, convergent validity, and discriminant validity. In step 3 we perform a series Mixture Factor Analyses (MFAs; Muthén 2007) to classify participants on these factors. In step 4 the final MFA solution is compared to a simple Latent Class Analysis (LCA) with the same number of classes, where the items themselves (rather than the factors) are used for the classification. This step is taken to validate that the same grouping occurs when classification is performed without the restrictions imposed by the factor analysis. In step 5, we test for an attitude-behavior link by measuring the effect of

attitudinal factors on the behavioral factors using Structural Equation Modeling (SEM). Finally, in step 6 we test whether there are significant differences between classes in terms of these attitudes, as well as participants' demographics and related behaviors using a Multiple Indicators and Multiple Causes (MIMIC) model and simple linear regressions.



**Figure 6: The steps involved in our analysis of the dimensionality of information disclosure behaviors**

Below I report the dimensionality (step 2), the classification solutions of the MFAs (step 3) and the differences between classes (step 6) in the three datasets; the reader can refer to the original paper (Knijnenburg et al. 2013c) for the results of other steps.

The first dataset originated from the mobile app recommender study discussed in Section 3.2 (Knijnenburg and Kobsa 2013a). Table 4 shows that our analysis confirmed the existence of two dimensions: context and demographics. Furthermore, we found 4 user profiles showing distinctly different behaviors along these two dimensions: users with a low disclosure tendency on both dimensions (LowD, 67 participants), users with a low

context data disclosure tendency but a high demographic data disclosure tendency (DemoD, 176 participants), users with a medium level of disclosure tendency on both dimensions (MedD, 71 participants), and users with a high disclosure tendency on both dimensions (HiD, 179 participants; Figure 7). We also showed that users' general privacy concerns, collection concerns, and mobile Internet usage could potentially be used to distinguish between the different classes (Figure 8), although differences in these variables between classes are not overwhelming.

**Table 4: The items used in the app recommender study, along with their average rate of disclosure and the factor loadings of the CFA. The dashed lines delineate the four categories of demographics items, the order of which was randomized (these categories did not produce different factors).**

Type of data	ID	Items	Level of disclosure	Factor loading
<b>Context</b>  Alpha: 0.79 AVE: 0.652  Factor correlation: 0.432	1	Recommendation browsing	87.0%	0.767
	2	Location	84.8%	
	3	App usage	82.2%	
	4	App usage location	67.1%	0.749
	5	App usage time	73.2%	
	6	Web browsing	48.3%	
	7	Calendar data	62.9%	0.874
	8	E-mail messages	36.7%	0.835
	9	Phone model	84.6%	0.940
	10	Accelerometer data	65.3%	0.659
	11	Microphone	50.9%	0.796
	12	Credit card purchases	20.1%	
<b>Demographics</b>  Alpha: 0.86 AVE: 0.784  Factor correlation: 0.432	13	Favorite sports (fan)	86.8%	0.718
	14	News interests	92.7%	0.905
	15	Amount of TV watching	92.3%	
	16	Amount of reading	93.5%	
	17	Phone data plan	87.6%	
	18	Gender	94.9%	0.915
	19	Age	93.3%	
	20	Education	92.7%	
	21	Field of work	83.6%	
	22	Housing situation	87.4%	0.911
	23	Population density of area	90.7%	
	24	Relationship status	88.6%	
	25	Children	89.3%	
	26	Household income	74.2%	0.964
	27	Household savings	66.3%	0.957
	28	Household debt	64.5%	0.802
	29	Race	89.1%	
	30	Political preferences	86.4%	
	31	Workout routine	90.1%	




Figure 7: User profiles in the app recommender study (4-class MFA).

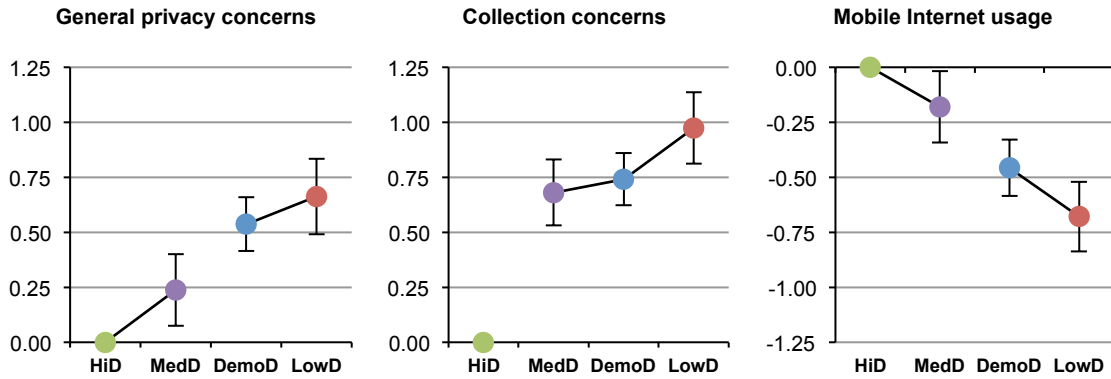


Figure 8: Differences between profiles in terms of general privacy concerns, collection concerns, and mobile Internet usage. Points that are not connected are significantly different from one another. The vertical axes are scaled in sample standard deviations of the measured factor. HiD is fixed to zero, and error bars are  $\pm 1$  standard error of the difference with HiD.


The second dataset comprises the 359 US participants of Wang et al.'s (2011b) cross-cultural Facebook study (222 female; median age: 28, range: 18 to 75). Participants indicated on a seven-point scale their level of comfort with disclosing 16 different Facebook profile items to “everyone on the Internet”. Table 5 shows that these items formed 4 dimensions: Facebook activity (Act), Location (Loc), Contact info (Con), and Life and interests (Int). We also determined 5 distinct behavioral profiles regarding these four dimensions: users with a low disclosure tendency on all four dimensions (LowD); users with a high tendency to disclose location and interests, but not activity and contact info



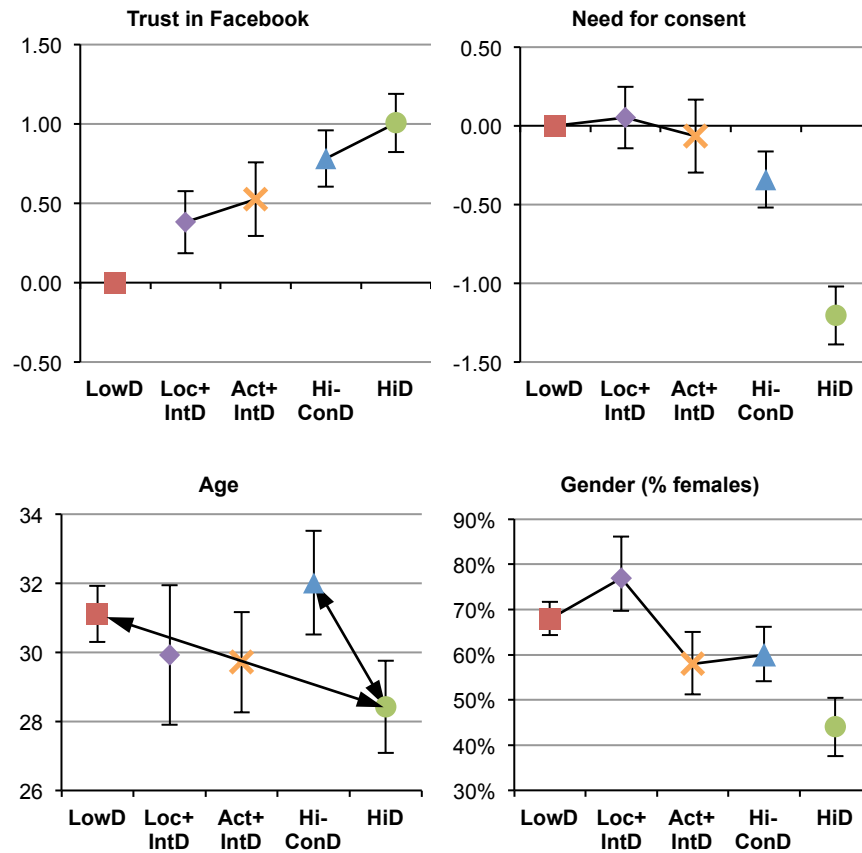
(Loc+IntD); users with a high tendency to disclosure activity and interests, but not location and contact info (Act+IntD); users with high disclosure tendencies on all dimensions except contact info (Hi-ConD); and users with a high disclosure tendency on all dimensions (Figure 9). Finally, we showed that users' trust in Facebook, need for consent, age, and gender could be used to distinguish between the different classes (Figure 10); again, the found differences are not overwhelming.

Type of data	ID	Items	Level of comfort							Factor loading
			1	2	3	4	5	6	7	
<b>Facebook activity</b>  Alpha: 0.93 AVE: 0.790	1	Wall	83	44	32	47	52	53	48	0.820
	2	Status updates	80	56	43	39	55	44	42	0.953
	3	Shared links	64	45	36	68	54	45	47	0.885
	4	Notes	102	55	44	55	37	32	34	0.907
	5	Photos	132	55	38	34	37	31	32	0.874
	16	Friend list	60	34	50	73	51	43	48	
<b>Location</b>  Alpha: 0.95 AVE: 0.919	6	Hometown	62	50	32	63	61	44	47	0.924
	7	Location (your current city)	72	62	41	56	46	37	45	0.960
	8	Location (your current state/province)	60	54	42	58	53	40	52	0.958
<b>Contact info</b>  Alpha: 0.85 AVE: 0.792	9	Residence (your street address)	240	33	23	23	19	10	11	0.884
	11	Phone number	262	29	15	22	19	4	8	0.933
	12	Email address	159	58	35	41	28	22	16	0.849
<b>Life and interests</b>  Alpha: 0.88 AVE: 0.756	13	Religious views	45	30	27	109	50	33	65	0.740
	14	Interests (favorite movies, books, etc.)	32	26	36	79	72	51	63	0.913
	15	Facebook groups that you are a member of	35	33	38	79	65	51	58	0.942
	10	Employer	110	53	36	73	43	24	20	

**Table 5: The items used in the Facebook study, along with the frequencies at each level of comfort (ranging from “not at all comfortable” to “very comfortable”), and the factor loadings of the CFA.**



**Figure 9: User profiles in the Facebook profile study (5-class MFA).**



**Figure 10: Differences between profiles in terms of trust in Facebook, need for consent, age, and gender. For age, arrows indicate a significant difference between HiD compared to LowD or Hi-ConD. For the other variables, points that are not connected are significantly different from one another.**


The third and final dataset (N=154) was gathered specifically for this study. Participants were first asked to enter the answer to 24 demographic questions into a text field, with the option to not disclose it instead. We then asked participants for each item how likely they were to provide the answer to an online retailer.

We constructed our 24 items so that 6 of them were related to health (Hlth), 6 to interests (Int), 6 to work (Wrk), and 6 to more general issues including contact information (Con). Table 6 shows that these 4 dimensions were confirmed in the dataset: all items fell on the hypothesized factors, except for the “computer software” item. We predicted that


this item would load on the work factor, but instead it loaded on the interest factor. In hindsight, this makes perfect sense. We also determined 4 behavioral profiles regarding these dimensions: users with a high tendency to disclosure contact information, but none of the other dimensions (ConD); users with a medium disclosure tendency on all dimensions except contact information (for which disclosure is low; Med-ConD), users with a high tendency to disclose interests and contact information, but not health and work-related information (Int+ConD); and users with a high disclosure tendency on all dimensions except contact information (Hi-ConD; Figure 11). Finally, we found age differences between these profiles (Figure 12), but no differences in privacy concerns were detected between the classes.

**Table 6: The items from the online retailer study, with frequencies at each comfort level and factor loadings of the CFA. ‘#’ indicates the request order (randomized, reversed in the second condition)**

Type of data	ID	#	Items	Level of comfort							Factor loading
				1	2	3	4	5	6	7	
<b>Health</b>  Alpha: 0.92 AVE: 0.782	A1	8	Physical health	21	20	13	19	31	31	19	0.919
	A2	23	Number of doctor visits in the past month	36	21	18	23	19	17	20	
	A3	20	Weight (lbs)	25	26	19	24	24	22	14	
	A4	22	Dietary restrictions	19	17	15	26	27	27	23	
	A5	12	Whether you use birth control	43	26	15	26	12	17	15	
	A6	14	Whether you have diabetes	22	27	13	20	25	20	27	
<b>Interests</b>  Alpha: 0.91 AVE: 0.850	B1	7	Favorite pastime	13	13	10	22	31	35	30	0.894
	B2	2	Favorite musical band/artist	6	12	11	23	27	44	31	
	B3	4	Favorite food	5	11	5	24	26	49	34	
	B4	21	Favorite movie	7	10	9	20	33	40	35	
	B5	24	Last holiday location	23	16	11	26	25	24	29	0.929
	B6	6	Relationship status	16	23	9	19	31	33	23	
	B7	3	Computer software you are familiar with	8	9	11	28	31	41	26	
<b>Work</b>  Alpha: 0.93 AVE: 0.823	C1	5	Highest completed degree	15	16	14	24	25	33	27	0.943
	C2	9	Work experience (years)	18	24	7	26	27	29	23	
	C3	18	Current/previous occupation	27	27	18	19	27	19	17	
	C4	13	Current/previous field of work	15	25	16	25	31	24	18	
	C5	19	Current/previous income level	40	31	20	23	17	14	9	
<b>Contact info</b>  Alpha: 0.87 AVE: 0.761	D1	1	Name	16	21	15	17	25	36	24	0.912
	D2	16	Gender	5	9	4	22	29	43	42	
	D3	15	Age	8	11	7	27	37	37	27	
	D4	17	Address	43	32	4	18	24	17	16	
	D5	10	E-mail address	15	30	13	19	35	23	19	
	D6	11	Phone number	50	37	14	21	19	7	6	



**Figure 11: User profiles in the online retailer study (4-class MFA).**



**Figure 12: Age differences between profiles. Arrows indicate significant differences.**

Using three datasets of online information disclosure intentions and behaviors, this work demonstrates that information disclosure behaviors are not unidimensional but instead consist of multiple related dimensions. Furthermore, we show that people can be classified into distinct groups that show very different behaviors along these dimensions. It is important to make such distinctions, since they may reveal that groups of people with the same amount of overall disclosure can show very different “disclosure profiles” if one looks at more than one dimension. Our datasets contain a number of examples for this. In dataset 1, two groups of participants exhibit the same medium level of disclosure, but one group (MedD) discloses both context items and demographics items at a medium rate while the other group (DemoD) discloses almost all demographics items but almost no context items. In dataset 2, one group (Loc+IntD) has high intentions to disclose location-

related items but low intentions to disclose activity-related items, and another group (Act+IntD) has the opposite intentions. Similarly, in dataset 3 all participants seem to have different disclosure tendencies for contact information compared with all other types of information, but in different extent and direction.

Note that we found differences between classes in terms of demographic variables, traits and attitudes, but these differences were not very large. This makes it possible to use these variables as predictors to assign privacy profiles to users, but accurate prediction will not be possible. It seems especially difficult for these variables to distinguish between different dimensions of disclosure behavior.

In sum, this study shows the importance of contextualizing information disclosure behavior in terms of the *type of information* that is being requested (dimensionality) and the *disclosure profile* of the user (classification). It also presents a generic approach to perform this contextualization, an approach that I have since applied to other datasets (Wisniewski et al. 2014).

*Refer to Knijnenburg et al. (2013c) and Knijnenburg (2013, 2014) for expanded analysis and discussion.*

#### **4.4 Types of recipients**

Many existing studies on information disclosure have identified the recipient of the information as an important contextual variable that influences people's disclosure decisions (Benisch et al. 2011; Consolvo et al. 2005; Hsu 2006; Johnson et al. 2012; Kairam et al. 2012; Lederer et al. 2003; Norberg et al. 2007; Olson et al. 2005; Patil and Kobsa 2005; Toch et al. 2010; Watson et al. 2012). This makes managing one's privacy a rather

complex task, especially in social situations: users report difficulties managing their privacy on their mobile phones (Kelley et al. 2013) and their social networking services (SNS) (Lipford et al. 2008; Madden 2012).

Suggestions have been made to simplify privacy settings interfaces, for example, by allowing users to *categorize* the recipients as a means to more efficiently manage what they want to share with whom. Below I discuss related work that serves as a rationalization for this approach, as well as work that investigates existing categorizations. In many cases, however, existing categorizations have not been formally tested. Because of this shortcoming, I subsequently present our own work establishing the optimal categorization of recipients for an SNS privacy settings interface.

Early research on social sharing has shown that people want to share their personal information with their social connections selectively (Deuker 2012; Krasnova et al. 2009; Lederer et al. 2003; Olson et al. 2005; Patil and Lai 2005). Indeed, SNS users tend to restrict access to their profiles by sharing certain information with certain people only (Kairam et al. 2012; Madden 2012; Young and Quan-Haase 2009). Facebook and Google+ both allow users to share specific posts and profile items with specific categories of recipients (Kairam et al. 2012; Watson et al. 2012). Facebook has three categories by default (Close Friends, Friends, and Acquaintances), while Google+ has four (Friends, Family, Following and Acquaintances).

On both networks users can create their own categories, but this feature is rarely used to make disclosure more selective (Deuker 2012; Strater and Lipford 2008; Watson et al. 2012; Wisniewski et al. 2014), because when people are prompted to categorize their friends into semantically meaningful categories, they often create categories that are

inadequate for making privacy decisions (Kelley et al. 2011). System-defined categories based on an analysis of users' disclosure behavior may thus be better suited for the purpose of selectively disclosing personal information, but the optimal management of such categories is not straightforward.

Social networks and (mobile) operating systems serve as platforms for third-party applications (apps). Apps often request access to some of the user's personal information. Most platforms require users to grant apps all these permissions with a single click, but research has shown that users want more granular control over their app permissions (Xu et al. 2012). This would make it more difficult to set each specific permission, though. Given that users tend to have different preferences for different types of apps (Kelley et al. 2013), the platform could help users by storing default settings based on the type of app that is requesting the permissions. This would require a manageable yet sufficiently detailed categorization of app types (e.g., Apple has a long list of app categories). It is unclear though, how to give this categorization the optimal level of granularity for expressing one's default permission preferences.

Online we could also benefit from a contextualized representation of the recipients of information, especially when it comes to information disclosed via web forms. I already discussed that users typically display purpose-specific disclosure behaviors on web forms, except when using a traditional form auto-completion tool (see Section 2.7). More recent versions of such tools allow users to specify the types of websites that may receive certain personal data (Ardagna et al. 2011; Bokhove et al. 2012; Cranor et al. 2002; Kolter and Pernul 2009). Again, it is unclear whether the provided or user-specified categories in such

systems are adequate, and what categories have the optimal classification and level of granularity for expressing one's privacy preferences.

Summarizing, in managing default settings for interactions with SNS contacts, application permissions, and web form disclosures, *categorizations based on the recipient* may provide a significant simplification of the management task. Crucially, though, in all of these situations the optimal categorization strategy is as of yet unknown. In the next section we therefore investigate the optimal categorization for SNS contacts using a novel categorization methodology that can be applied more generically to all sorts of recipient categorization problems.

#### **4.5 Recipient groups (original work)**

Although the idea of group-based contextualization of recipients has been proposed in the literature, no work to date has considered studying the optimal segmentation of recipients into groups (Olson et al. (2005) is a notable but limited exception). We therefore developed a practical methodology for creating a privacy-relevant segmentation of the recipients of personal information, and applied this methodology to develop a concise categorization scheme for the specification of privacy preferences in social networks. Finally, we tested these categorizations at different levels of granularity to see which categorization users found most satisfying to use.

Our methodology for creating a privacy-relevant segmentation of recipients is based on the principles of *discriminant and convergent validity*; terms that originate from the field of *psychometrics* ( DeVellis 2011).



In psychometrics, the term discriminant validity is used to indicate that two scales measure two separate constructs (Bagozzi and Phillips 1991). We used the term to indicate that users' preferences for two categories of recipients are sufficiently different from each other. When discriminant validity is low (i.e., when users have essentially the same preferences for the two categories), this increases the effort for users, because then they have to make the same setting for both of these categories. One could then reduce the privacy-setting effort without sacrificing precision by merging these non-discriminant categories into a single category.

The term convergent validity is used in psychometrics to indicate that the items of a scale robustly measure a single construct (Bagozzi and Phillips 1991). We used the term to indicate that users will likely apply the same settings to all members of a category. When the convergent validity of a category is low (i.e., when users have diverging preferences for certain category members), users' preferences are not precisely captured by that category. The accuracy could then be increased by splitting the category into subcategories.

Establishing both discriminant and convergent validity creates a concise set of categories that allows users to express their sharing preferences with maximal precision (fewest errors) and minimal effort (fewest clicks).

#### *4.5.1 Methodology*

Our methodology consists of 5 steps. The first step is to select a set of information items that are a fair representation of the types of information that would be shared using the system (SNS example: Table 7), and step 2 is to create a set of 3-10 generic category

labels and a large number of specific category labels that break the generic label down across one or more dimensions (SNS example: Table 8).

**Table 7: Items used in the SNS example.**

Type of data	Index item
Facebook activity	Status updates Photos
Location	Hometown Location (your current city)
Contact info	Phone number Email address
Life and interests	Religious views Interests/Likes (favorite movies, books, etc.)

**Table 8: Labels used in the SNS example.**

Set of labels	Label (and designated generic label)	Set of labels	Label (and designated generic label)
Generic labels	Family (fa) Friends (fr) Colleagues (co) Classmates (cl) Acquaintances (ac)	Age	Older relatives (fa) Younger relatives (fa) People of my age (ac) People younger than me (ac) People older than me (ac)
Extent of kinship	Household members (fa) Next of kin (fa) Extended family (fa) Distant relatives (fa) In-law family (fa)	Work-related	Work friends (co/fr) Colleagues on my team (co) Colleagues not on my team (co) Subordinates (co) Superiors (co)
Friendship demogr.	Guy friends (fr) Girl friends (fr) Younger friends (fr) Friends my age (fr) Older friends (fr)	School-related	High school classmates (cl/ac) Best friends from high school (cl/fr) College classmates (cl/ac) Best friends from college (cl/fr) Best friends from elementary school (cl/fr)
Distance	Local family (fa) Family members who live far away (fa) Local friends (fr) Friends who live far away (fr) People in my neighborhood (ac)	Acquaintances	Friends of friends (ac) People I've only met online (ac) People I hardly know (ac) People I follow (but don't know me) (ac) People I've lost contact with (ac)
Disliked people	Family members I don't like (fa/ac) Colleagues I don't like (co/ac) Sketchy/questionable friends (fr/ac) People I don't like (ac) Ex-boyfriends/ex-girlfriends (fr/ac)	Commonality	People I've met only once or twice (ac) People I talk to only once in a while (ac) People I have nothing in common with (ac) People with the same hobbies as me (fr) People I talk to on a daily basis (fr)
Extent of friendship	Old friends (fr) New friends (fr) Family friends (fr/fa) Best friends (fr) Good friends (fr)		

Step 3 is to recruit participants and collect data from a number of participants (SNS example: N=449) regarding their disclosure and perceived risk of disclosing each of the selected items to different people that match each label (for each participant, test a subset of the labels). Step 4 is to analyze the data, which itself consists of the following steps:

- **Risk and disclosure level analysis:** compares the average level of disclosure or risk for a specific label with its generic label. If the difference is substantial, then the specific label should be treated as a separate category.
- **Risk and disclosure variance analysis:** compares the internal variance of disclosure or risk for a specific label with its generic label. If the variance of the specific label is substantially smaller than that of the generic label, then the specific label should be treated as a separate category.
- **Risk and disclosure difference analysis:** compares the within-subjects differences in disclosure or risk between two labels by measuring how much additional variance is introduced by combining the labels. This analysis is performed on all *within-subjects* label pairs: for each label pair, only participants that evaluated both labels are considered. For each of these participants, compute the mean risk per item for each label separately. Compute the squared difference between these risk levels, and average these values over all items and all participants. If this difference is substantially large, then that specific label should be treated as a separate category.

Finally, step 5 is to determine the optimal categorizations based on the results. Using different cut-off criteria for the analyses, one can create categorizations at different


levels of granularity, and these can be tested against each other in an online user experiment with the target system (or a mockup thereof).

#### 4.5.2 *Analysis of the SNS example*

I report the analysis of the SNS example below. For brevity I only report the risk analyses, split up by generic label.

Figure 13 shows the results for the Family label. The label has good discriminant validity; it is different from all other generic labels, except maybe the Friends label. The Family label also has good convergent validity. No Family-related labels have lower internal variance. Also, most labels do not differ much from the generic label and can be subsumed, with a few exceptions:

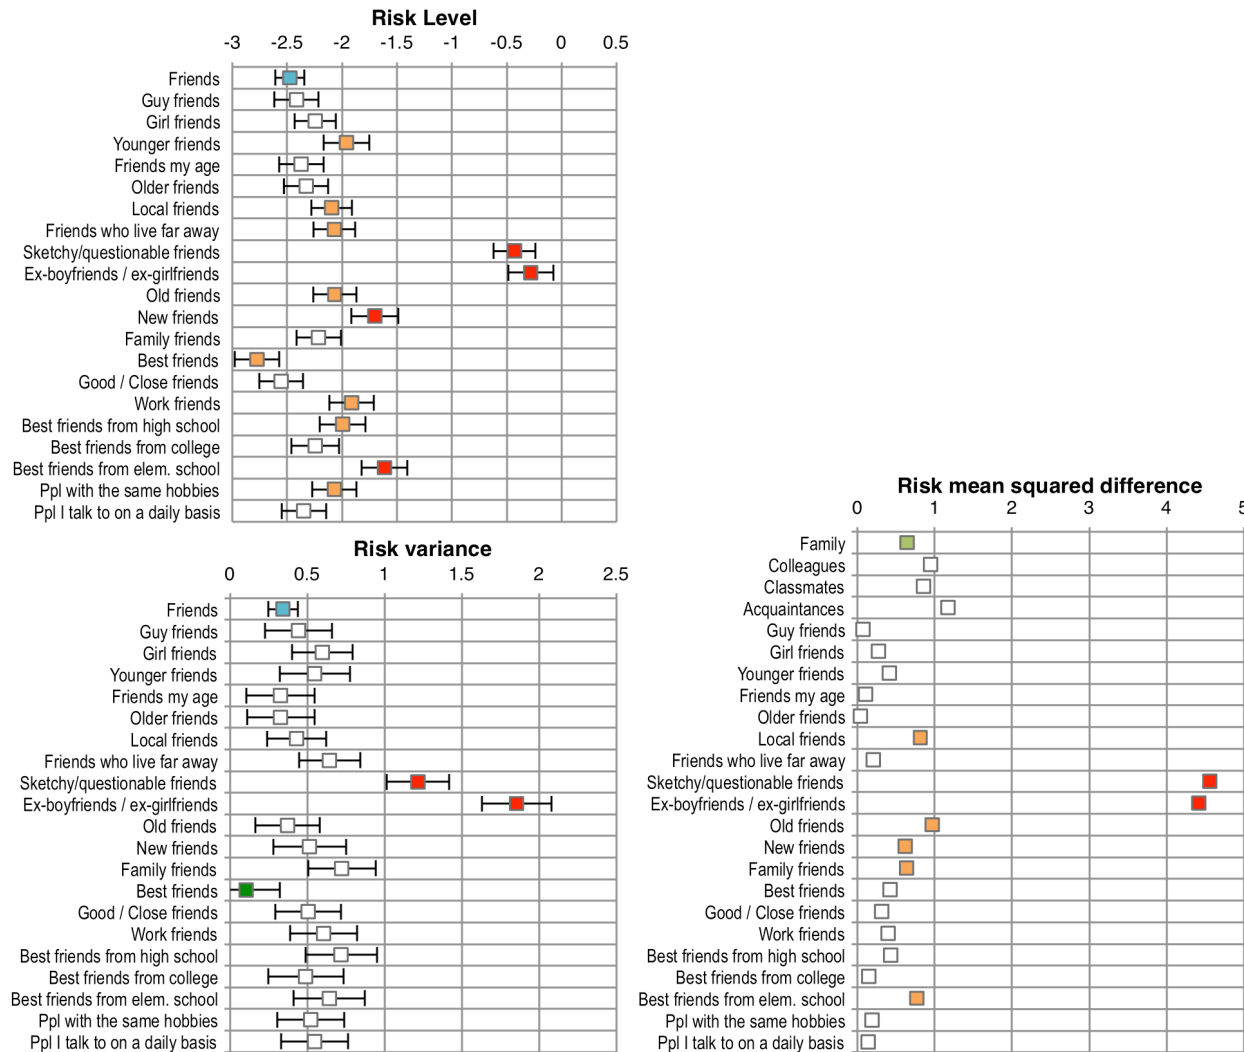
- Extended, Distant and In-law family have a higher risk level. The label Distant relatives has some risk difference, and Extended and In-law family show more risk variance than Family. Distant relatives and In-law family do not fit well, though.
- Older and Younger relatives have a higher risk level. These labels are similar to each other; they describe family members from other generations.
- The “Family members I don’t like” label does not fit the Family category at all.



**Figure 13: Analysis of the Family-related labels**

Figure 14 shows the results for the Friends label. It has good discriminant validity; it is different from all other generic labels, except maybe the Family label. It has good convergent validity, although the Good/close and Best friends labels have a lower risk variance. Most other labels do not differ much from the generic Family label and can thus be subsumed, with a few exceptions:

- New, Young, and Work friends have a higher risk level. New and Family friends also differ slightly in terms of disclosure from the generic label. These friends have in common that one may restrict information access to them to keep up good appearances.
- Elementary school best friends and Old friends have a higher risk level, and differ slightly from the generic label. These are friends one may have lost contact with.
- Sketchy friends and exes do not fit the Friends category.




**Figure 14: Analysis of the Friends-related labels**

Figure 15 shows the results for the Colleagues label. The label has reasonable discriminant validity, although it is somewhat similar to Classmates and Acquaintances. It has reasonable convergent validity, although the following split could be made:

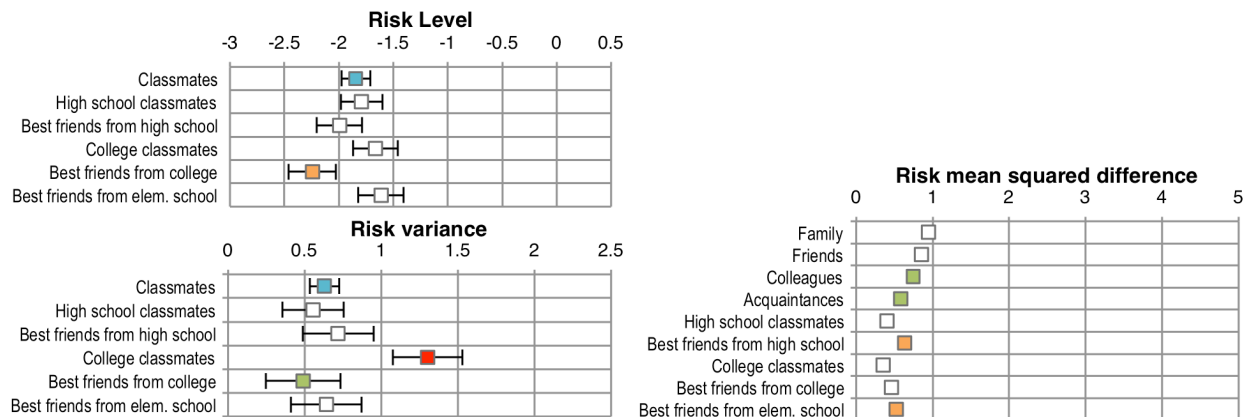
- The Colleagues on my team and Work friends labels show less variance than the generic label, and somewhat lower risk levels. These labels fit well together, although Work friends also fits under Friends.

- Superiors, Subordinates, and Colleagues on other teams have a higher risk level. The latter two do not fit well together though.
- Colleagues I don't like/trust do not fit the category.



**Figure 15: Analysis of the Colleagues-related labels**

Figure 16 shows the results for the Classmates label. It has reasonable discriminant validity, but it is a bit similar to Colleagues and Acquaintances. It has a reasonable convergent validity, but a split between school friends and classmates fits better. School friends fall well enough under the generic Friends label, though.

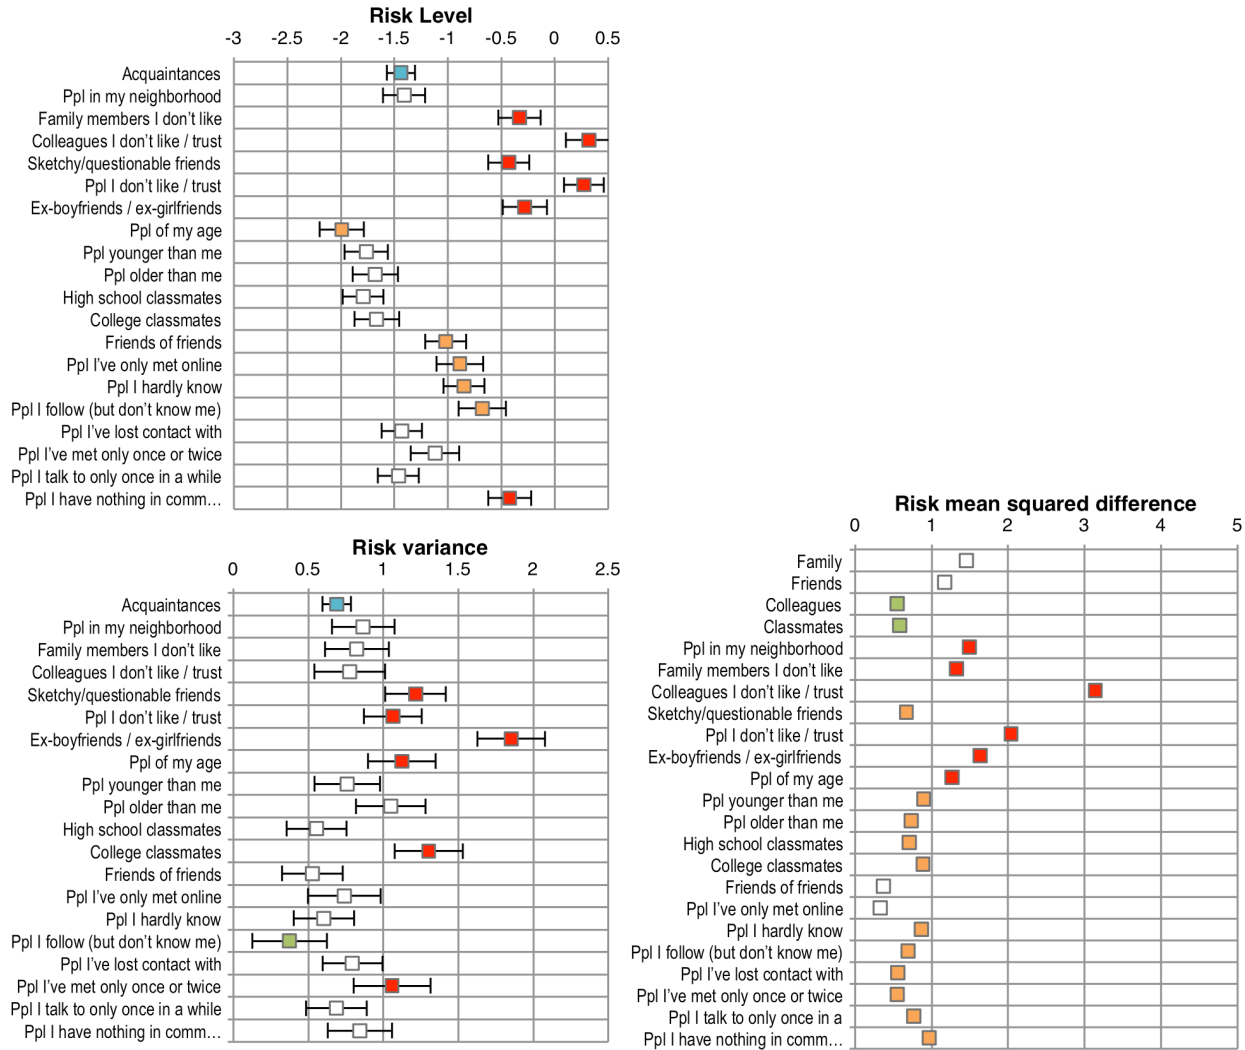


**Figure 16: Analysis of the Classmates-related labels**

Finally, Figure 17 shows the results for the Acquaintances label. It has reasonable discriminant validity, but it is a bit similar to Colleagues and Classmates. It has low convergent validity: several specific labels have less internal variance, or do not fit very well under the generic label. The risk of these labels is generally higher. Broadly speaking, there are three subcategories:

- Untrusted people, i.e., Friends and Family I don't trust, People I don't trust, and Colleagues I don't trust.
- Unknown people, i.e., Friends of friends, People I've met online, People I hardly know, People I follow, People I've met only once or twice, People I have nothing in common with, and Neighbors.
- Infrequent contacts, i.e., People I've lost contact with and People I only talk to once in a while.





**Figure 17: Analysis of the Acquaintances-related labels**

*Refer to Knijnenburg et al. (2014a) for expanded analysis and discussion.*

#### 4.5.3 Determining the optimal category granularity

Based on the results, we developed the following four categorizations, and tested them (minus the very coarse variant) against each other to determine the optimal level of granularity:

**Very coarse (2 groups):**

- Family/friends
- Classmates/Colleagues/Acquaintances

**Normal (5 groups):**

- Family
- Friends
- Classmates
- Colleagues
- Acquaintances

**Granular (10 groups):**

- Immediate family
- Extended family
- Close friends
- Regular friends
- Best-behavior friends
- Peer colleagues on my team
- Other colleagues
- Classmates
- Infrequent contacts
- People I hardly know or don't trust

**Very granular (14 groups):**

- Immediate family
- Relatives older or younger than me
- Extended family
- Close friends
- Regular friends
- Best-behavior friends
- Friends from my past
- Peer colleagues on my team
- Colleagues on other teams
- Superiors/Subordinates
- Classmates
- Infrequent contacts
- People I hardly know
- People I don't trust

Participants in this study (N=387) were recruited to test the settings interface of Mundo, a hypothetical new social network site. They were asked to enter the names of people fitting a list of 50 different “person descriptions”. Subsequently, they were told to imagine using a new social network on which they would categorize their contacts and set their sharing preferences. To categorize their contacts, participants were given the ~50 names they provided earlier, and asked to categorize them into one of the 5, 10, or 14 categories.

Figure 18 shows the settings interface used in the subsequent task of setting the sharing preferences. The default setting, presentation order of categories, and ability to make exceptions for specific contacts varied according to the experimental condition. Participants were asked to carefully consider what to share with whom, and to set their settings accordingly.

Finally, participants were asked to fill out a questionnaire that measured their perceptions of over-disclosure threat, the perceived ease of use of the settings interface, and their anticipated satisfaction with the Mundo system. The survey concluded by measuring participants' interpersonal privacy concerns.

**Mundo Profile Settings**  
Please set your sharing preferences below.

<b>Immediate family</b> Madeleine Kipling, Makenzie Trueman, Mellony Ralph, Brannon Garth, Val Jonette, Randolph Sally, Saranna Lillie. <a href="#">Change individual settings...</a>	Status updates	Photos	Home-town	City & state	Phone number	E-mail address	Religious views	Likes (movies, books...)
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<b>Extended family</b>	Status updates	Photos	Home-town	City & state	Phone number	E-mail address	Religious views	Likes (movies, books...)
Kennedy Talon	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Indiana Jed	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Sammy Effie	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Layton Bobbie	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Melantha Izzy	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Gussie Kamryn	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
<b>Close friends</b> Lenny Grasmie, Hester Ashleigh, Carlie Lake, Ebony Kathryn, Claudia Alton <a href="#">Change individual settings...</a>	Status updates	Photos	Home-town	City & state	Phone number	E-mail address	Religious views	Likes (movies, books...)
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 18: Mundo profile settings interface.

In Section 2.5 I mentioned that some researchers suggest that finer-grained control (e.g., by means of more granular categories) is a necessary precondition to foster information sharing in social networks services (SNS), because users will otherwise “err on

the safe side” (Benisch et al. 2011; Sadeh et al. 2009; Tang et al. 2012) Our own work discussed there argued against this, and demonstrated users may share *more* when given only coarse-grained controls. In the current study we found no differences in disclosure between the normal, granular and very granular conditions.

Regardless of whether more coarse controls entice over- or under-disclosure, researchers seem to agree that more granular controls allow users to set their privacy settings to a level that better reflects their sharing preferences (Toch et al. 2010; Tsai et al. 2009; Wang et al. 2011a). For example, in an empirical study of a location-sharing settings interface, Tang et al. (2012) found that users of the finer-grained settings interface were more comfortable with their privacy settings. We therefore hypothesized that more granular categories could reduce users’ over-disclosure threat. However, our results showed that compared to the normal granularity condition, participants perceive higher (rather than lower) over-disclosure threat in the more granular categorization conditions (yet interestingly not significantly in the very granular categorization condition). This is opposite to what we hypothesized.

Finally, several researchers have noted that users of fine-grained interfaces find it difficult and time-consuming to accurately set their privacy settings (Madejski et al. 2012; Sadeh et al. 2009; Strater and Richter 2007). Too much control may thus have a detrimental effect on users’ perceived ease of use of the privacy settings interface. We therefore hypothesized that more granular categories would be more difficult to use. In our study, the perceived ease of use was indeed significantly lower in the very granular categorization condition.

To summarize these findings, we found that category granularity has no effect on disclosure, but that the granular (10-category) categorization increases perceived over-disclosure threat, and the very granular (14-category) categorization is harder to use. SNS managers should thus employ the 5-category categorization. Different types of systems may require different categorizations, though. In these situations, managers can employ our methodology for developing and testing these categorizations.

*Refer to Knijnenburg and Kobsa (2014) for expanded analysis and discussion.*

## CHAPTER 5: Personalized Privacy

### 5.1 Introduction

The ultimate goal of my dissertation is to use the contextualized understanding of users' privacy calculus to develop a Privacy Adaptation Procedure to support people's information disclosure decisions. This procedure first predicts users' privacy preferences and behaviors based on their known characteristics. In systems with lots of privacy settings, it then provides automatic initial default settings in line with users' "disclosure profiles" (e.g., by default, it discloses Mary's location to her best friends on weekends, but it does not disclose John's location to his boss when he is on vacation). By alleviating some control, these "smart defaults" could overcome the Control Paradox, but at the same time they arguably respect users' inherent privacy preferences. In conversational systems, request orders can be adapted to users' privacy preferences in a similar fashion, by prioritizing requests for information that the user is predicted to consider non-sensitive, and therefore more likely to disclose.

Another aspect of the Privacy Adaptation Procedure is to show disclosure justifications in situations where they are needed, but to only show them to users who can be expected to react rationally to them, so that they will not cause privacy scares in the other users (thereby overcoming the ironic effect of transparency, see Section 2.3). For instance, the procedure would not explain to John why a dating site would like to know the brand, make and year of his car, since it predicts that John would not be swayed by such an explanation anyways but rather lose trust in the system. The system would show this information to John only if John explicitly asks for it. Likewise, if David is predicted to

disclose details of his car on the dating site anyway, the procedure would conclude that no justification is needed, and not display one without being asked. The procedure would however justify the disclosure request for Mary if it predicts that she is amenable to a justification and would otherwise not disclose this information.

As Solove (2013) put it, “there is no silver bullet, and so we must continue to engage in an elaborate dance with the tension between [privacy] self-management and paternalism.” The Privacy Adaptation Procedure aims to strike this balance between giving users no control over, or information about, their privacy at all (which will be insufficient in highly sensitive situations, and may deter privacy-minded individuals) and giving them full control and information (which, due to increased confusion, is arguably bad in all other situations and for all other people). Arguably, the procedure relieves some of the burden of the privacy decision from the user by providing the right privacy-related information and the right amount of privacy control that is useful, but not overwhelming or misleading. This way, it is designed to give users “realistic empowerment”: it should enable them to make privacy-related decisions within the limits of their bounded rationality.

Practically speaking, implementing a Privacy Adaptation Procedure consists of two parts: First, the system must determine the context of the current privacy decision. The contextual variables explored in Chapter 4 have a strong influence on people’s privacy decisions, and should therefore be used as inputs for this part. Then, it must present the default or justification that best fits this context.

There are three ways in which a system can determine the context of a privacy decision. The simplest solution is to ask the user to specify the context, such as when the system asks the user to categorize a recipient. Another option is to derive contextual



variables from other variables, such as when the system uses user characteristics (e.g., age, gender (Cho et al. 2009; Li and Chen 2010), cultural background (Lin et al. 2012; Wang et al. 2011b), or mobile Internet usage) or attitudes (e.g., privacy concerns) to determine users' disclosure profile. In Section 4.3 I mentioned that these variables seem to only be able to halfway reliably predict the amount of information disclosure, but that they often cannot distinguish between different dimensions of disclosure behavior. Finer-grained privacy attitude scales ought therefore to be developed (van de Garde-Perik et al. 2008). Finally, a system can determine the context "on the fly", by observing users' disclosure behaviors during the interaction. As a finer-grained, dynamic, and thus ultimately more precise method to establish context, most existing work in the field of privacy recommendation uses this "on the fly" approach.

For example, Ravichandran et al. (2009) apply *k*-means clustering to users' contextualized location sharing decisions to come up with a number of default policies. They show that a small number of default policies for the user to choose from could accurately capture a large part of their location sharing decisions.

Sadeh et al. (2009) apply a *k*-nearest neighbor (kNN) algorithm and a random forest algorithm to learn users' privacy preferences in a location-sharing system. They show that users had difficulties setting their privacy preferences, and that the applied machine learning techniques can help users in specifying more accurate disclosure preferences.

Pallapa et al. (2014) proposed context-aware approaches to privacy preservation in wireless and mobile pervasive environments. One of their solutions leverages the history of interaction between users to determine the level of privacy required in new situations. They demonstrate that this solution can deal with the rise of privacy concerns while at the

same time efficiently supporting users in a pervasive system full of dynamic and rich interactions.

Finally, in a social network context, Fang and LeFevre (2010) developed a privacy wizard that is able to configure users' privacy settings automatically and accurately with a machine learning model that they developed. The wizard removes the burden of setting privacy settings using tools that most users would otherwise find too difficult to understand and use.

Balebako et al. (2011) argue that the help provided via such machine learning systems can be seen as adaptive nudges. Indeed, Smith et al. (2013) argue that "Smart defaults can become even smarter by adapting to information provided by the consumer as part of the decision-making process." (p. 167) The presented studies cannot provide much of an argument in favor or against such adaptive nudges though, because like most studies in the field of recommender systems they focus on the *accuracy* of the privacy preference modeling techniques; the effect on the users' *experience* (e.g., satisfaction, perceived privacy threat) has not been investigated yet (Knijnenburg et al. 2012c).

We therefore conducted a series of studies that investigate the potential (Section 5.2) and actual (Sections 5.3 and 5.4) benefits of privacy adaptation to users' experience. Because they are pioneering work, the "context learning" behind these studies is very simplistic. In Chapter 6 however, I present a series of experiments with a recommender system that implements a more sophisticated Privacy Adaptation Procedure. This system is the capstone of my dissertation project.

## 5.2 Adaptive justifications for app recommenders (original work)

In Section 3.2 I discussed our study that demonstrates that justifications failed to increase users' disclosure or satisfaction. Upon further analysis of these surprising results, we noticed that the optimal justification depended on two contextual variables: the characteristics of the user and type of information requested. In this section I discuss the results of this contextual analysis.

To use disclosure tendency as a contextual variable, we split the users into two groups: one with a low disclosure tendency (up to 22 disclosed items), which comprises 33.3% of the participants, and one with high disclosure tendency (23-31 disclosed items). Gender further split each group into roughly half. We then determined the optimal combination of request order and justification *within each group*. We used three different definitions of "optimal", depending on the potential goal of the system: increasing the disclosure of context data, increasing the disclosure of demographic data, and increasing the users' subjective experience of the system (using the subjective variables perceived disclosure help, perceived privacy threat, trust in the company's privacy practices, and satisfaction with the system).

Figure 19 displays the combined effect of justification type and request order on disclosure behavior for the four user groups (males/females with low/high disclosure tendency). The 5-way interaction effect (justification type  $\times$  request order  $\times$  gender  $\times$  disclosure tendency  $\times$  information type) is significant. Figure 20 shows the effects of these variables on users' subjective experiences. Within each user group, we compare the best strategy for each data type (marked with an arrow) against all other strategies. Strategies that perform significantly worse than the best strategy are labeled with a  $p$ -value.

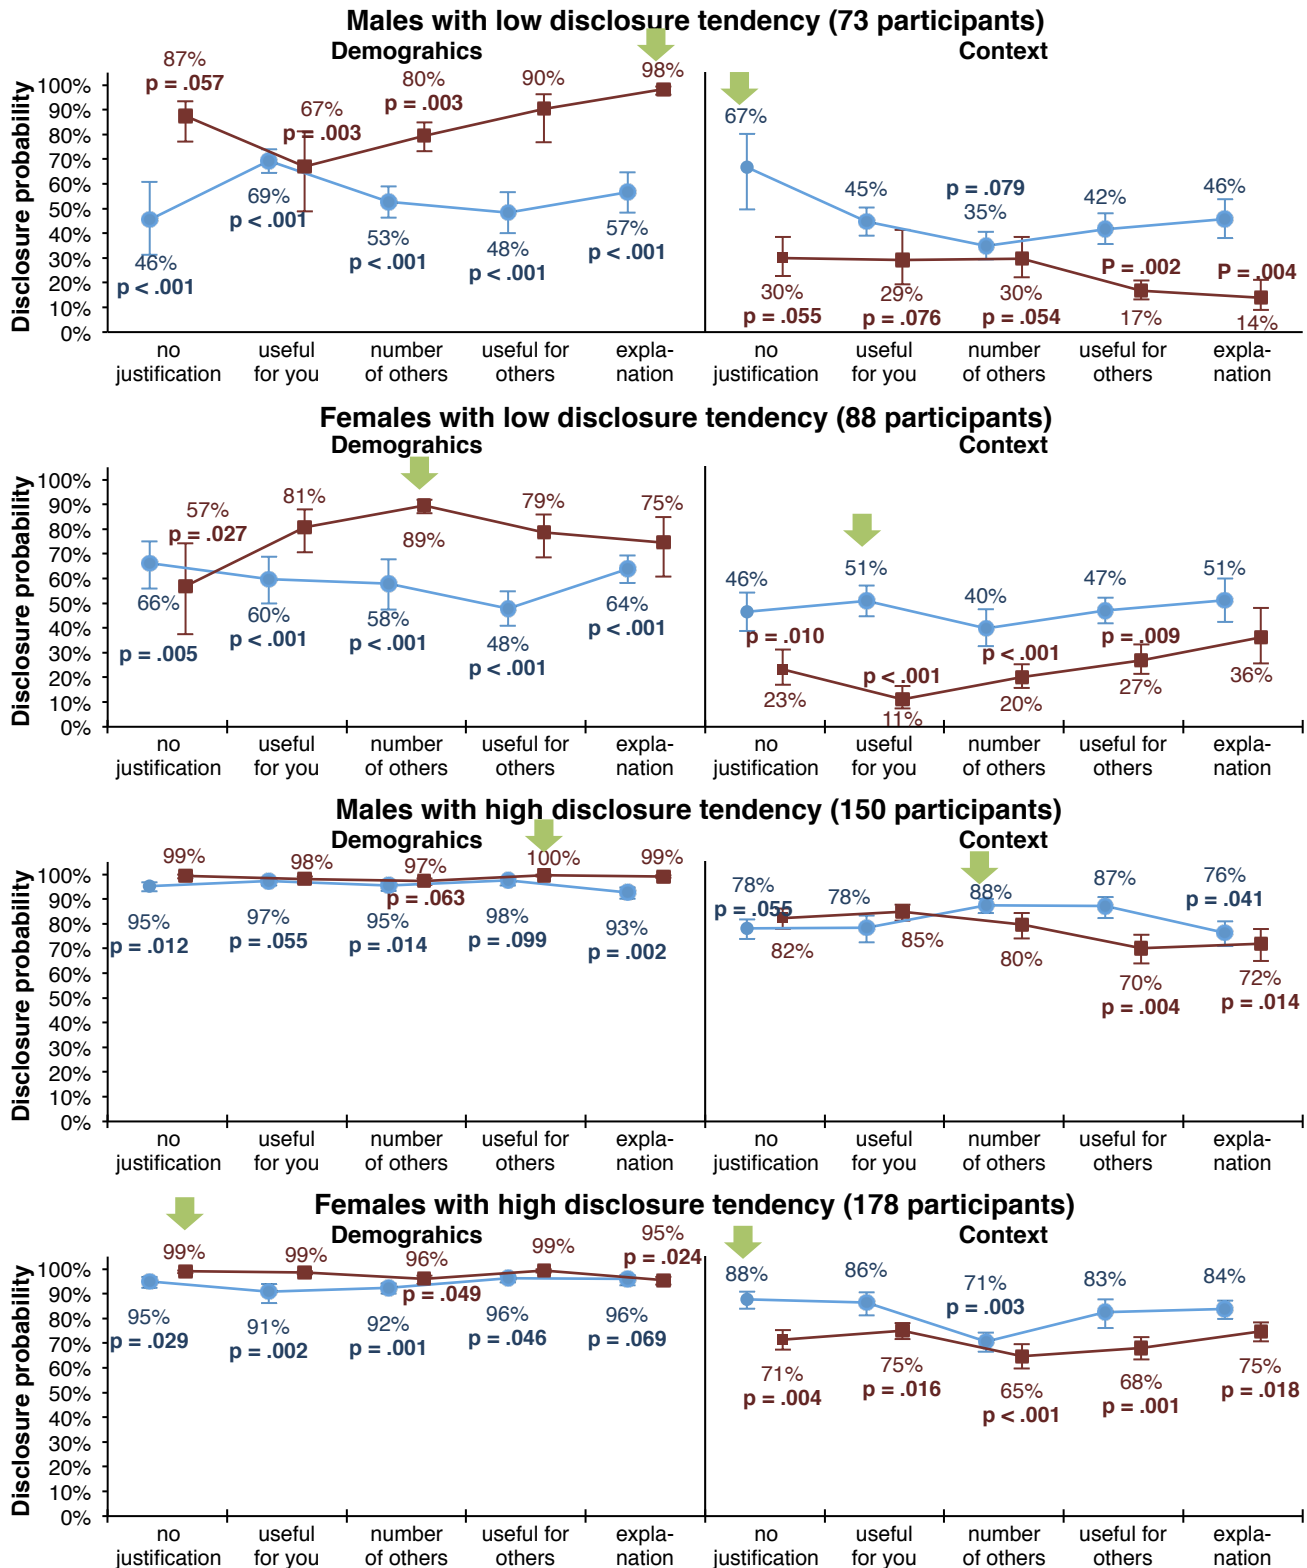




Figure 19: The effects of justification type and request order (blue: context first; brown: demographics first) on disclosure (for each type of data, gender, and disclosure tendency). Error bars are  $\pm 1$  standard error. The best strategy is labeled with an arrow; strategies with a p-value perform significantly worse.

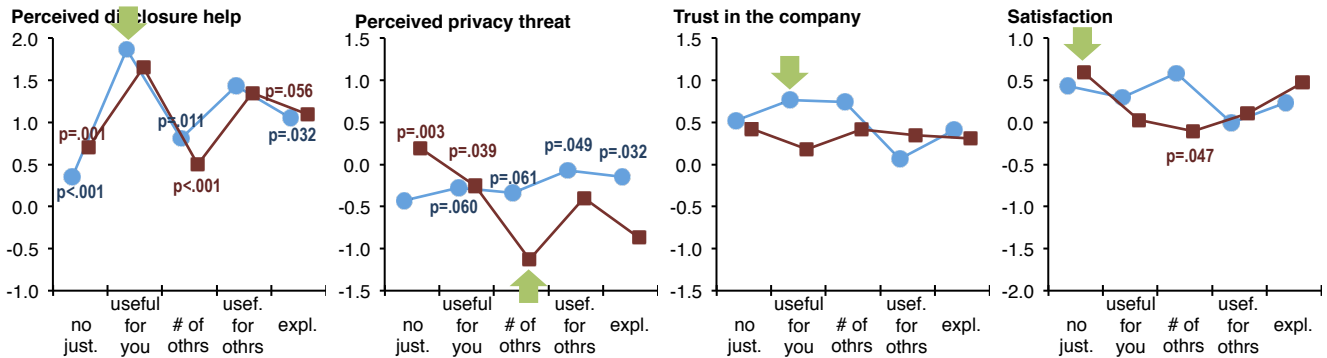
### Males with low disclosure tendency (73 participants)



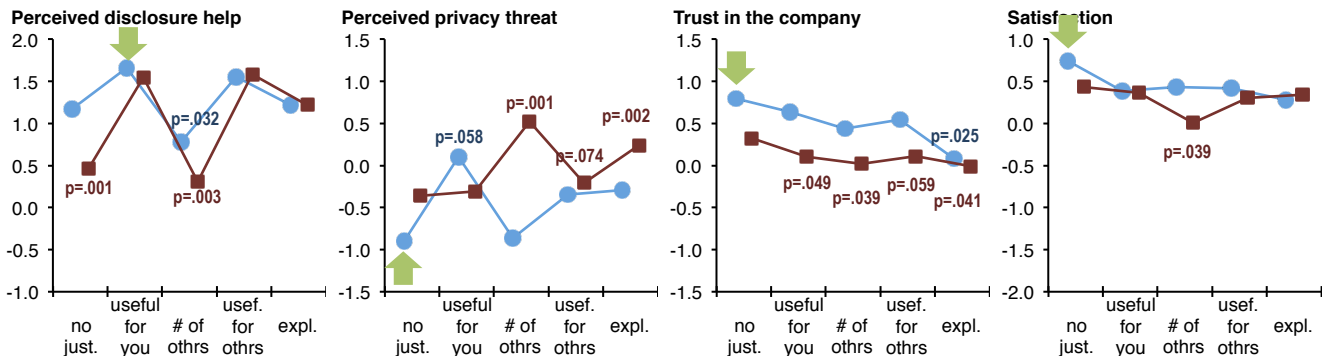
### Females with low disclosure tendency (88 participants)



### Males with high disclosure tendency (150 participants)



### Females with high disclosure tendency (178 participants)



**Figure 20: The estimated effects of justification type and request order (blue: context first; brown: demographics first) on the subjective constructs for each gender and disclosure tendency. Since the outcomes are scale-free factor scores, the y-axis is scaled in sample standard deviations, and the value for [male, low disclosure tendency, context first, no justification] is set to zero.**

**Table 9: Heuristics to find the best strategy, based on user characteristics (gender, disclosure tendency) and optimization goals.**

<b>Best strategies for MALES with LOW disclosure tendency</b>	
<i>Goal</i>	<i>Best strategy</i>
High demographics disclosure	Demographics first, 'explanation' justification.
High context data disclosure	Context first, no justification.
High overall disclosure	Context first, 'useful for you' justification.
High satisfaction	Context first, 'useful for others' justification or demographics first, 'useful for you' justification.
All of the above	Demographics first, no justification.
<b>Best strategies for FEMALES with LOW disclosure tendency</b>	
<i>Goal</i>	<i>Best strategy</i>
High demographics disclosure	Demographics first, 'number of others' justification.
High context data disclosure	Context first, 'useful for you' justification.
High overall disclosure	Demographics first, 'explanation' justification.
High satisfaction	Context first, 'useful for you' justification.
All of the above	Demographics first, 'explanation' justification.
<b>Best strategies for MALES with HIGH disclosure tendency</b>	
<i>Goal</i>	<i>Best strategy</i>
High demographics disclosure	Demographics first with any justification except 'number of others'.
High context data disclosure	Context first, 'number of others' or 'useful for others' justification.
High overall disclosure	Demographics first with no justification or the 'useful for you' justification, or context first with 'useful for others' justification.
High satisfaction	Demographics first, 'useful for others' or 'explanation' justification.
All of the above	Demographics first, 'useful for you' justification.
<b>Best strategies for FEMALES with HIGH disclosure tendency</b>	
<i>Goal</i>	<i>Best strategy</i>
High demographics disclosure	Demographics first with no justification, the 'useful for you' justification, or the 'useful for others' justification.
High context data disclosure	Context first with no justification.
High overall disclosure	Context first with no justification.
High satisfaction	Context first with no justification.
All of the above	Context first with no justification.

The results show that the best strategy depends on users' disclosure tendency and gender. It also depends on the goal of the system: some strategies increase disclosure of one type of data but not the other, and some increase disclosure but at the same time reduce users' satisfaction. We therefore suggest that the strategy should be adapted to the optimization goal of the system and the characteristics of the user. Table 9 outlines heuristics for selecting the best strategy

for each type of user, given a certain system goal, as well as a suggestion for a compromise between optimizing disclosure rates and satisfaction simultaneously.

To follow the heuristics, a system would have to discover the users' characteristics before or during the interaction. Gender can just be the first item to request. In fact, gender disclosure was the highest of all items in our study (94.9%), and hence we expect that asking for it first will not raise any concerns. To correctly determine the users' disclosure tendency, the system would have to first ask a number of potentially invasive questions, which is not desirable. Alternatively, one could ask about (or otherwise determine) the users' stated privacy concerns, mobile Internet usage and/or tech-savvyness, since these characteristics are related to users' disclosure tendency.

The most important take-away from this study is that while none of the presented justifications work well for everyone (which led to our earlier conclusion that justifications generally fail as a nudge), Table 9 shows there are subsets of users for which certain justifications work better than providing no justifications at all. In other words, an *adaptive justification*, based on a careful consideration of user characteristics, can significantly improve strategies for helping users with information disclosure decisions.

*Refer to Knijnenburg and Kobsa (2013b) for expanded analysis and discussion.*

### **5.3 Smart defaults for social network sites (original work)**

In the social network privacy settings study described in Section 4.5.3 we also tested smart default settings. Particularly, the defaults of the settings interface


(Figure 18) were manipulated at three levels: no checkboxes would be checked in the *private-by-default* condition, all checkboxes would be checked in the *checked-by-default* condition, and a subset of the checkboxes would be checked in the *smart default* condition. This subset was determined based on the results of our categorization study (Section 4.5.2). Specifically, we checked a box if participants in that study shared that item with members of that category at a rate of at least 70%.

We hypothesized that users' disclosure in a smart default setting should fall between the private-by-default setting (which may cause under-disclosure) and the disclosed-by-default setting (which may cause over-disclosure). Figure 21 shows that this was generally indeed the case. Specifically, compared to the private-by-default condition, the odds of disclosure are estimated to be 2.1 times as high in the smart default condition, and 3.9 times as high in the disclose by-default-condition. However, this effect is smaller for participants with low interpersonal privacy concerns and when categories are ordered weaker ties first<sup>4</sup>. As a result, for participants with low privacy concerns in the weaker ties first condition, the private-by-default condition shows a level of disclosure that is significantly *higher* than the smart default condition.

---

<sup>4</sup> We used two different presentation orders for the privacy settings interface: The stronger ties first condition started with family and friends (typically stronger ties), then classmates and colleagues (typically weaker ties), and then acquaintances (typically the weakest ties). The weaker ties first condition used the opposite order.





**Figure 21: Observed disclosure levels in each default condition, split by privacy concerns and category order.**

We furthermore hypothesized that users of the disclosed-by-default condition would perceive a higher level of over-disclosure threat than users of the private-by-default and smart default conditions. Our results confirm this hypothesis, but the effect disappears when participants can make exceptions; in that case disclosed-by-default results in no higher over-disclosure threat than the other conditions.

Concluding, the smart default setting did not stand out as the superior solution: it did not result in higher disclosure rates than the private-by-default setting (except for people with high privacy concerns in the stronger ties first condition), and it did not reduce over-disclosure threat or increase ease of use. Although we based the smart default on previous data using a somewhat arbitrary threshold, this threshold seemed to be quite accurate: analyzing users' deviations from the default, we found that the number of disclosures turned into non-disclosures (median: 11, mean: 21.9) was about equal to the number of non-

disclosures turned into disclosures (median: 10, mean: 15.9). Also, the total number of changes in the smart default condition (median: 25, mean: 37.7) was much lower than the number of changes in both the private-by-default condition (median: 257, mean: 251.9) and the disclosed-by-default condition (median: 91, mean: 103.3). This means that users in the smart default condition indeed used far fewer clicks to optimally set their privacy settings, but that this presumptive reduction in physical effort was not accompanied by a reduction in cognitive effort (i.e., perceived ease of use).

*Refer to Knijnenburg and Kobsa (2014) for expanded analysis and discussion.*

#### **5.4 Privacy recommendations for location-sharing (original work)**

The field of location-sharing services (LSS) has seen a number of papers proposing a prediction-based approach to support users in setting their privacy settings (Pallapa et al. 2014; Ravichandran et al. 2009; Sadeh et al. 2009). These works demonstrate that location sharing preferences can accurately be predicted, but they have to date not actually implemented their prediction algorithms in a location-sharing service and then tested it with real users. We filled this gap in the literature by implementing a simple (manual-input) based recommender in a mockup of a location sharing system and then evaluating its impact on users' sharing behavior and subjective experience (perceived help, decision freedom, over-disclosure threat, trust in the company, and satisfaction with the system).

The study considers a fictitious location-sharing service called "HotSpots", which recommends locations to visit based on previously visited locations and also

allows users to share their location on Facebook. The system allows users to choose among 8 different disclosure actions (based on Duckham and Kulik 2005; Li and Chen 2010; Tang et al. 2012):

- A1. Fully use the system
- A2. Restrict Facebook posts to friends that are nearby
- A3. Restrict Facebook posts to certain friends only
- A4. Restrict Facebook posts to only share city
- A5. Restrict Facebook posts to only share city block
- A6. Use the system for recommendations only
- A7. Turn the system to “private mode” (anonymous)
- A8. Turn the system off

Giving users these finer-grained options should reduce their privacy concerns (Consolvo et al. 2005), but also turns location-sharing into a rather complex decision that puts extra burden on the user (Compañó and Lusoli 2010). We therefore decided to help users in this decision by framing the decision in a way that matches users’ *evaluation* of the activity. The question “What do you think about this activity?” is arguably easier to answer than the question “How do you want to share this location?” Moreover, if this evaluation is strongly related to users’ sharing behavior, we can use it to *recommend* a (restricted set of) sharing action(s).

#### 5.4.1 *Determining the potential for adaptation*

Our first study (N=100) was an online user experiment to test the hypothesis that users’ evaluation of the activity is a good predictor of users’ sharing behavior.

We showed participants 10 scenarios, asked them to choose a sharing action, and then asked them to choose among the following 10 evaluations of the activity (based on Kairam et al. 2012; Sleeper et al. 2013):

- E1. is exciting
- E2. is interesting for others
- E3. makes me proud
- E4. makes me look interesting
- E5. needed a good recommendation
- E6. is private
- E7. embarrasses me
- E8. isn't useful for everyone
- E9. doesn't really represent me
- E10. may have unintended consequences when shared

Table 10 shows that there is a strong relation between the disclosure action and the evaluation of the activity. Given the evaluation, it is thus possible to *recommend* an action to the user. For instance, if we recommend only the most-selected action for each evaluation, we are recommending the “correct” sharing action to the user 43.2% of the time, which is considerably higher than the 12.5% we would get by recommending a random action. For practical use this is not very accurate, but if we recommend not one but a small set of actions, this set would contain the “correct” option more often than not. For example, if we recommend the

dark gray cells in Table 10, we can get 81.5% recall<sup>5</sup> with 2.3 actions on average per evaluation. Increasing the number of recommended actions to just under 4 actions on average per evaluation (dark and light gray cells in Table 10), we can get 95.1% recall.

**Table 10: The co-occurrence of actions and evaluations.**  
Gray cells show possible action recommendations for each reason.

	E1	E2	E3	E4	E5	E6	E7	E8	E9	E10
A1	34	88	14	25	24	0	0	1	1	1
A2	6	25	1	6	6	3	0	32	0	4
A3	5	16	6	9	6	17	3	41	1	8
A4	1	8	1	11	6	4	2	10	0	2
A5	0	3	4	1	1	2	1	1	1	5
A6	2	5	0	1	23	112	17	58	16	36
A7	0	0	1	0	0	80	18	20	19	40
A8	0	0	0	0	0	34	14	27	4	26

#### 5.4.2 Testing the recommenders

Based on the results of the first study we created a system that first asks the user to evaluate the activity and then recommends a subset of the sharing actions that users are likely to choose. Two questions need to be answered when designing such a “privacy recommender”:

**How many actions should it recommend?** Recommender systems researchers have found that list length is an important determinant of user satisfaction (Bollen et al. 2010). In our case, a longer list of recommendations would be less restrictive and would have a higher accuracy, but may not help the user enough in terms of simplifying her decision.

---

<sup>5</sup> “Recall” is a machine learning term for the likelihood that a set of recommendations contains the item actually selected by the user.

**How should it present recommendations?** The system could hide actions that are not recommended, thereby reducing visual clutter but also increasing the risk that the user cannot find her desired action. Alternatively, the system could *highlight* the recommended actions, keeping all options on the screen, but also increasing the complexity of the interface.

We explored these questions in an online experiment with the HotSpots mockup (N=368) by testing 5 versions of the privacy recommender against 2 baseline conditions (resulting in a total of 7 between-subjects conditions):

- C1. **No recommendation:** Regardless of the users' evaluation, all sharing actions are displayed (this is the "comparable baseline" condition).
- C2. **Long list, rest hidden:** The dark gray and light gray actions from Table 1 are listed as "recommended options"; the rest is hidden under a "more options" link.
- C3. **Short list, rest hidden:** The dark gray actions from Table 1 are recommended; the rest is hidden.
- C4. **One item, rest hidden:** Only the most popular action for that evaluation is displayed, the rest is hidden.
- C5. **Short list, highlighted:** All actions are displayed, but the dark gray actions from Table 1 are highlighted.
- C6. **One item, highlighted:** All actions are displayed, but the most popular for that evaluation is highlighted.

C7. **No evaluation:** Same as C1, but the user does not evaluate the activity (this is the “optimized baseline” condition, because no evaluation is needed if the system is not using it for recommendations).

In every condition (except for C7), the system first asks the user to evaluate the activity using one of 7 options.<sup>6</sup> Each recommender then tailors the display of the 8 sharing actions to the selected evaluation. In terms of evaluating these recommenders, we focused on the following aspects:

**How accurate is the recommender?** Using offline evaluations, previous work has shown relative success in predicting users’ binary (yes/no) sharing decisions (cf. Cranshaw et al. 2011; Toch et al. 2010). Our recommender has to predict among 8 actions though, which is considerably harder. Moreover, offline accuracy evaluations do not always agree with online evaluations (McNee et al. 2002). We thus purposefully evaluate the accuracy of our recommender in an online evaluation. The first line in Table 11 shows the *recall* of each recommender: the proportion of decisions that were in line with the recommended action. As expected, longer lists have a higher recall, but the short lists perform particularly well given the lower number of recommendations. Moreover, the recommenders that hide items have a higher recall than the recommenders that highlight items. The “rest hidden” recommenders are thus more persuasive than the “highlighted” recommenders (more on persuasion below). This is likely due to the additional effort it takes in these systems to select an option that is not initially listed.

---

<sup>6</sup> We combined E2/E4, E6/E10, and E7/E9 because they were similar evaluations and also showed very similar behavior (see Table 10).

**Table 11: Recall in the 5 recommender conditions (C2–C6).**

	<b>C2: Long list, rest hidden</b>	<b>C3: Short list, rest hidden</b>	<b>C4: One item, rest hidden</b>	<b>C5: Short list, highlighted</b>	<b>C6: One item, highlighted</b>
<b>Recall in study 2</b>	98.7%	92.2%	75.0%	86.6%	62.5%
<b>Recall in study 1 (ex-post)</b>	95.1%	81.5%	42.8%	81.5%	42.8%

**Is the recommender persuasive?** Merely calling an item a recommendation may increase the chances that users choose it (Cremonesi et al. 2012; Pathak et al. 2010). This would result in accuracy levels that are even higher than predicted based on study 1, especially when the recommender hides the other actions. This was indeed the case in our study: the “actual” recall in the recommender conditions (Table 11, line 1) is *higher* than the ex-post recall<sup>7</sup> (line 2): the mere fact that certain options were presented as “recommendations” increased their likelihood to be chosen. In other words, the system persuaded participants to choose one of the recommended actions. Companies can use this persuasive power to influence users’ behavior through recommendations. Note, however, that recommending items that the user clearly does not like is likely to result in *reactance* (behavior that explicitly counters the recommended action) and to lower satisfaction (Brehm 1966; Fitzsimons and Lehmann 2004). This argument is in line with Wilson et al. (2013), who also warn that the subset of available sharing options has to be “carefully considered” because it “can influence users to share significantly more without a substantial difference in comfort”.

---

<sup>7</sup> The recommendations in study 2 are selected in such a way that if users were to behave exactly the same as in study 1, the recall would be as high as possible (i.e., by recommending the options that were chosen most often in study 1). Therefore, the recall in study 1 is optimized “ex post”. Since users may not behave exactly the same in study 2, the “actual” recall is very likely to be lower, *except* when participants are persuaded to select a recommendation rather than their own preferred option.

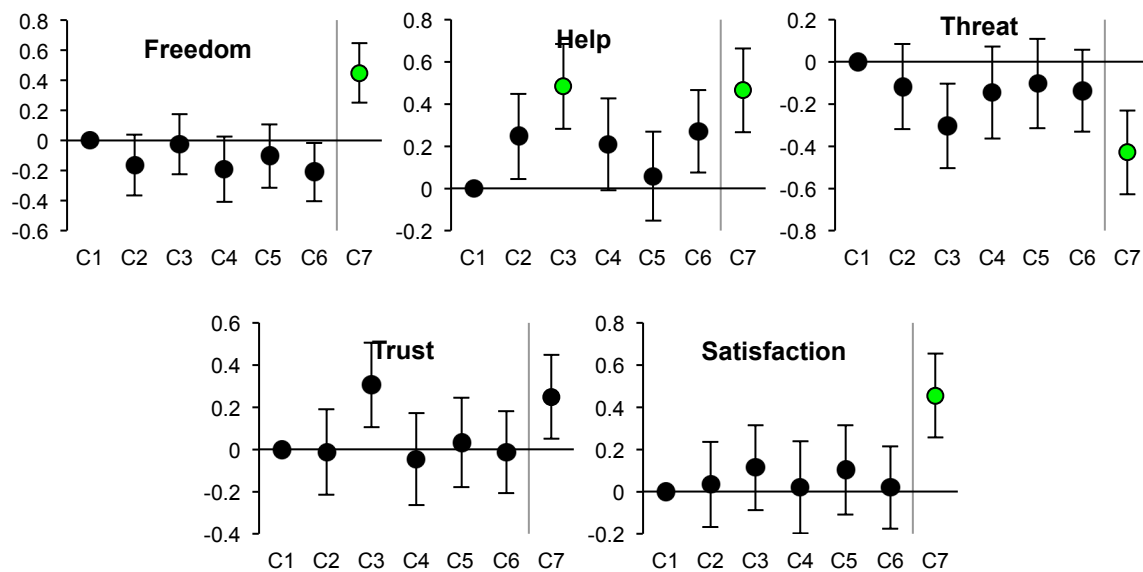


**Does the recommender increase satisfaction?** Accurate recommenders are not always more satisfying to the user, and researchers have thus called for a more comprehensive, subjective evaluation of recommender systems (Knijnenburg et al. 2012c). Recommenders may give users a sense that they are helped (Häubl and Trifts 2000), but they must leave users enough freedom to make their own decisions (Pariser 2012). Moreover, inaccurate recommendations may be perceived as nefarious (Fitzsimons and Lehmann 2004), which in our case may manifest itself as privacy threat and reduced trust. We thus evaluate the recommender with a comprehensive post-study questionnaire that measures users' subjective evaluations. Figure 22 compares the recommenders (C2–C6) against the two baseline conditions (C1 and C7) in terms of perceived decision freedom, perceived decision help, perceived threat, trust in the company, and system satisfaction.

Temporarily ignoring the optimized baseline (C7), we observe that although the recommenders result in somewhat lower (yet not significantly lower) perceived decision freedom, they do result in somewhat higher perceived decision help, especially the “short list, rest hidden” recommender (C3), which is perceived as significantly more helpful than the baseline system without recommendations. The recommenders also result in slightly lower perceived threat, and C3 seems to instill some trust in the company (albeit not significantly). In terms of system satisfaction, the recommenders are on par with C1.

Returning to the optimized baseline, Figure 22 shows that this system has a significantly higher decision freedom, higher decision help, lower threat, and higher satisfaction than baseline C1. The difference between C7 and the other conditions is

that participants in C7 are not asked to evaluate the described activity before choosing a sharing action. This poses an interesting dilemma: Although a recommendation (i.e., C3) can increase the perceived decision help, asking for the evaluation that is necessary to give such a recommendation actually ruins the positive effect of the recommendation itself. Asking for an evaluation thus thwarts the positive effect of the recommender system.



**Figure 22: The effect of the recommenders (C2–C6) on subjective measures.**  
The error bars are  $\pm 1$  SE of the comparison with C1

Concluding, we showed that users would feel assisted by privacy recommendations, but that the input required for these recommendations would counter the positive effects on their satisfaction. This of course reduces the practical applicability of our results, but from an academic perspective we see this as a valuable lesson for those who want to create adaptive privacy systems: building an accurate system is not enough, it needs to be accepted by users as well (cf. McNee et

al. 2006). In this specific case, these results may mean that our initial premise (evaluating the activity is easier than choosing a sharing action) is false. Alternatively, it may mean that users are made more aware of the dangers of a location-sharing service when asked to evaluate the activities described in our rather “risqué” scenarios. This, in turn, increases users’ perceived threat, despite the comparatively wide range of options our LSS offers to protect their privacy (this is in line with the ironic effect of transparency, see Section 2.3). Luckily, day-to-day location sharing rarely involves extreme, privacy-sensitive scenarios such as those presented in our study. The “inadvertent awareness effect” would thus arguably be smaller in a real-life implementation of our recommender, where users are generally much more likely to consider sharing “exciting” and “interesting” activities than “embarrassing” and “private” ones.

*Refer to Knijnenburg and Jin (2013a; 2013) for expanded analysis and discussion.*

## **CHAPTER 6: Adaptive request order for demographics-based recommender systems (original work<sup>8</sup>)**

### **6.1 Motivation**

In previous Chapters I have demonstrated a potential for recommendations as a means to alleviate some of the burden of privacy decision making while at the same time respecting users' individual privacy preferences. At the same time, I have shown that even rudimentary implementations of privacy recommendations can lead to unexpected effects on users' satisfaction. Are these effects intrinsic to the practical implementation of privacy recommendations, or are they due to the fact that our initial implementations were somewhat simplistic in nature? To complete this dissertation I present a series of studies that implements a more sophisticated Privacy Adaptation Procedure in a live recommender system giving energy-saving or healthy living related advice. The system recommends items (in our case: advice) based on the user's demographics, and at the same time adapts the sequence of demographic requests to the user's privacy preferences. This doubly adaptive approach allows us to explicitly model the privacy/benefits trade-off that is inherent in most privacy decisions.

---

<sup>8</sup> The studies presented in this chapter comprise the capstone work of this dissertation. These studies have not yet been published elsewhere, and are therefore described in full detail.

## 6.2 Background

Recommender systems filter a large number of alternatives to a shorter list of options (Bollen et al. 2010; Knijnenburg et al. 2012c; Resnick and Varian 1997). Recommender systems tailor this list to the users' preferences, thereby optimizing the chance that it contains the option that the user wants to choose. By removing the other options, recommender systems have the potential to reduce choice overload (Bollen et al. 2010; Willemsen et al. 2011).

### 6.2.1 Preference elicitation methods

The way in which recommender systems allow users to express their preferences, the *preference elicitation method* (PE-method), has been the topic of numerous studies (e.g., Chen and Pu 2009; Dooms et al. 2011; Gena et al. 2011; Lee and Benbasat 2011; Sparling and Sen 2011). Particularly, the PE-method seems to have an impact on users' satisfaction with the system (Chen et al. 2009; Knijnenburg et al. 2012c; Sparling and Sen 2011). In the context of energy-saving, we have investigated a series of preference elicitation methods for a multi-attribute utility theory (MAUT) based recommender system<sup>9</sup> (Knijnenburg et al. 2011, 2014b; Knijnenburg and Willemsen 2009, 2010):

---

<sup>9</sup> A MAUT-based recommender system is ideally suited for domains where items are meaningful metadata (i.e., attributes) and where interactions are "one-time" rather than continuous. For other domains (e.g., movies, music), existing services (e.g., Netflix, Amazon) use a collaborative filtering (CF) based recommender system, which use either implicit feedback (e.g., consumption/purchasing behavior) or item ratings (e.g., 5-star rating scales). Note that CF-based recommender systems can (and do) also use demographic information as input (cf. Linden et al. 2003; Vozalis and Margaritis 2007), e.g., to overcome the "cold start" problem (i.e., to provide recommendations when little other input is available); the current study may thus apply to CF-based recommender systems as well.

**Attribute-based PE:** the most-used PE method for MAUT-based recommenders (Häubl and Trifts 2000; Olson and Widing 2002; Roy et al. 2008). In this method, users directly indicate the importance of each of the attributes with which choice options are described.

**Case-based PE:** takes an indirect approach to discover users' attribute weights, namely by analyzing the users' evaluation of exemplary choice options (Chen and Pu 2009, 2011; McGinty and Smyth 2006; Smyth 2007). In case-based PE, users' positive (or negative) evaluation of an example is indicative of their preferences regarding its most prominent attribute, and this evaluation is therefore translated into a higher (or lower) weight for that attribute.

**Needs-based PE:** takes the indirect approach to PE a step further: In this method, users express their preferences not in terms of product attributes, but in terms of consumer needs (Randall et al. 2007).

**Implicit PE:** does not require users to actively express their preferences. Instead, it infers the attribute weights as a by-product of the user's browsing behavior. When a user inspects, selects, or discards a recommended item, the system uses the attribute values of this item to update the user's attribute weights accordingly (Knijnenburg et al. 2011).

**Hybrid PE:** combines implicit PE with attribute-based PE to give users the convenience of automatic preference elicitation while still allowing them to monitor and control the attribute weights.

Our previous work tested these PE-methods against two baselines that are non-personalized: The top-N baseline simply ranks the energy-saving methods in

decreasing order of popularity, while the Sort baseline allows users to sort the recommendations on one of the attributes.

### *6.2.2 The moderating effect of domain knowledge*

The results of our evaluations show that the best PE-method depends on users' level of domain knowledge. Domain knowledge is a personal characteristic that may significantly influence one's decision strategy. For example, compared to novices, energy-saving experts have more knowledge about the underlying attributes of energy-saving measures, and are therefore better capable of translating their needs into attribute weights (Hutton and Klein 1999), and making complex tradeoffs between them (Shanteau 1988; Xiao and Benbasat 2007). Novices, in contrast, lack the knowledge required to understand the impact of the attributes (Hutton and Klein 1999), and may thus not readily know how to express their preferences in terms of product attributes (Xiao and Benbasat 2007).

Because experts and novices differ in the way they make decisions, they arguably also prefer different PE-methods (Butler et al. 2008). Spiekerman and Paraschiv (2002) indeed note that many existing recommender systems fail to motivate user interaction because they limit their interaction to attribute-based PE and fail to adjust to the user's level of domain knowledge. They propose a strategy to integrate different knowledge levels in the system by offering a different interface for experts and novices. Similarly, Guttman et al. (1998) suggest that "matching the system's user interface with the consumer's manner of shopping will likely result in greater customer satisfaction." (p. 153). Following this, Randall et al. (2007)

demonstrate that experts are more satisfied with a system that employs a parameter-based PE-method (a variant of attribute-based PE), while novices are more satisfied with a system that uses a needs-based PE-method.

Based on our own evaluations (Knijnenburg et al. 2011, 2014b; Knijnenburg and Willemsen 2009, 2010), it seems that energy-saving experts prefer complex systems that allow direct control over the attributes weights (attribute-based and hybrid PE), while novices prefer systems that are tailored to their needs (needs-based PE), provide limited control (sort) or rather no control at all (top-N). Most importantly, we find that tailoring the interface of the recommender to users' level of domain knowledge not only increases their satisfaction with the system; it is also the key to make users save more energy, because they end up choosing more and more impactful energy-saving measures.

### 6.2.3 *Demographics-based preference elicitation*

Our previous work shows that while experts are capable of expressing their preferences in terms of attributes, novices prefer either non-personalized systems, or systems that give them simpler means of expressing their preferences. Needs-based PE is one of these simpler means; in the current series of studies I propose to investigate *demographics-based PE* (Lee and Park 2007; Lee and Lee 2007; Oh and Moon 2012; Zheng et al. 2012) as an alternative PE-method that is arguably simple enough for novice users. I test this PE-method in our energy-saving recommender, as well as a recent adaptation of the system that gives advice about healthy living



(Elsten 2012). This domain is arguably more privacy-sensitive than energy-saving (Pattaraintakorn et al. 2007; Sezgin and Ozkan 2013).

In both the energy and healthy living domains, demographics are arguably an important determinant of preferences. For example, household size and housing situation potentially influence the kind of energy-saving measures one can implement, medical history and age are important variables to consider when advising people about healthy living, and income could be an important variable to take into account when recommending (potentially costly) measures to take in either domain. In line with our previous work, I argue that demographics-based PE will be most beneficial for domain novices, since demographics are usually known and easy to report; unlike attribute-based PE, demographics questions do not require one to make a trade-off between abstract values related to the recommendation domain.

#### *6.2.4 Personalization-privacy paradox*

Demographics-based PE may induce privacy problems, though, because unlike personal preferences, many demographic variables are considered privacy-sensitive (Ackerman et al. 1999). Moreover, researchers have shown that privacy can play a limiting role in users' adoption of personalized services. For example, in a study on ubiquitous commerce (u-commerce), Sheng et al. (2008) showed that personalization induced privacy concerns, and that users consequently would feel less inclined to use personalized (rather than non-personalized) u-commerce services, unless the benefits were overwhelming (i.e., providing help in an

emergency). Similarly, Awad and Krishnan (2006) showed that users' privacy concerns inhibited their use of personalized services and advertising, and Sutanto et al. (2013) demonstrated that privacy concerns can prevent people from using a potentially beneficial personalized application. These findings have led the FTC to highlight what Awad and Krishnan call the "personalization-privacy paradox": despite the benefits of personalization, users may not agree with the data-collection required to make personalization work (FTC 2010).

On other hand, it may be that these concerns mainly exist when such services fail to provide useful benefits (Hagel and Rayport 1999). Indeed, people are willing to give up privacy for personalization (Hann et al. 2007; Olivero and Lunt 2004), as long as this gives them benefits (Phelps et al. 2000), such as content relevance, time savings, enjoyment and novelty (Ho and Tam 2006; Hui et al. 2006). Some even go as far as to say that privacy concerns do not matter as long as the benefits are clear (Knight 2010). In most cases, though research has shown that privacy and benefits are both important in determining users' willingness to adopt and provide personal information to personalized services, and researchers therefore claim that they should both meet a certain threshold (Treiblmaier and Pollach 2007), or that they at least should be in balance (Chellappa and Sin 2005; Xu et al. 2009, 2011).

The aforementioned work on the privacy-personalization paradox fails to truly investigate the tradeoff between privacy and benefits as a concrete behavioral decision, because their outcome measure is a more generic form of an intention (i.e., it is measured with generic questionnaire items such as "How likely would you provide your personal information (including your location) to use the M-Coupon

service?”). Such stated intentions arguably do not directly relate to observable privacy behaviors (cf. Spiekermann et al. (2001b) and Norberg et al. (2007), who show that privacy preferences and actual behavior tend to be weakly related at best). In our previous work, discussed in Section 3.2, we instead considered users’ detailed privacy decisions (a yes/no decision for multiple disclosures), which is more compatible with existing information disclosure research (cf. Acquisti et al. 2012). A similar approach was used in Kobsa et al. (2014). In both works, we demonstrated that disclosure behavior in a demographics- and context-based recommender system was determined by system satisfaction, trust in the company, perceived privacy threat, and general privacy concerns.

#### 6.2.5 *Adaptive request order*

Personalization may thus cause privacy concerns, especially when demographic data are used. Since not all demographic data are equally sensitive, we hypothesize that it would be better for a recommender to ask less sensitive questions first. This may seem counter-intuitive in light of Acquisti et al.’s (2012) finding that disclosure rates are *lower* when asking less sensitive questions first. Similarly, in Section 3.2 we find that people are more likely to disclose the information that is requested first, so if the goal is to get users to disclose as much sensitive information as possible, asking the most sensitive questions first would arguably be the most successful strategy. However, our goal here is not simply to increase disclosure, but rather to provide good recommendations without inciting privacy concerns. If we ask sensitive questions first, these requests may trigger

users' privacy concerns, possibly even to the extent that the user will stop answering questions altogether—this was not possible in Acquisti's study, but will very likely occur in our study. In fact, it is a *desirable feature* of our interaction design: by relegating the most sensitive questions to the end of the interaction, users will likely not see these questions at all if they stop answering questions at some point during the interaction. In sum, due to the flexible nature of our interaction paradigm, users will likely answer (or at least consider answering) more questions and have lower concerns when the least (instead of the most) sensitive questions are asked first.

The quality of the recommendations is however also determined by the usefulness of the item: some items may have a more significant impact on the recommendation quality than other items (see Section 6.5.4 for a more thorough explanation of this point). In this sense, a system should make a *tradeoff* in deciding what question to ask next: it should employ an “automatic privacy calculus” (cf. Section 2.1) by weighing the predicted privacy sensitivity of each item with its potential usefulness for the recommender system.

This is not a trivial task, for several reasons. First of all, both privacy and usefulness are dynamic concepts. As I have demonstrated in Section 4.3, people with different privacy profiles differ in what types of information they find sensitive. Similarly, the benefit of each question depends on the current user model and the potential changes to it after asking the question (McGinty and Smyth 2006; Mirzadeh et al. 2005; Rashid et al. 2002). Secondly, since privacy and benefit may contradict each other, some trade-off needs to be made between these two

variables, and the optimal value of this trade-off may depend on the user's privacy concerns. Concluding, this problem requires a sophisticated Privacy Adaptation Procedure that dynamically optimizes the trade-off between privacy and benefit to decide which question to ask next. The proposed studies develop, implement and test this Privacy Adaptation Procedure in an energy- and health-related advice recommender.

### 6.3 Pre-study: linking demographics to attributes

**The first step towards a privacy-aware demographics recommender is to gather the essential parameters for such a system. We therefore conducted a pre-study that collects data about user demographics (multiple choice), recommender attribute weights (scale: 0-very unimportant to 10-very important), and perceived privacy risk (1-very safe to 7-very risky). We invited Amazon Mechanical Turk participants (N=200) to provide their demographic information (57 multiple-choice items, see Table 12) and their personal preference in terms of the 7 or 8 attribute weights for the energy or health advice recommender (**

Table 13; the type of recommender is manipulated between subjects).

Finally, we asked them how privacy-sensitive (perceived risk) each demographic item is in the context of an energy-saving or healthy-living recommender.

**Table 12: Demographics questions asked, and sensitivity levels as determined for study 2 (see Section 6.5.5).**

<b>Demographics question</b>	<b>Sensitivity</b>
What is your age?	-2.170
What is your gender?	-2.154
What is your height?	-2.132
Do you have children?	-2.101
Do you have a gym membership?	-1.734
How often do you eat fast food?	-1.707
Do you eat organic food?	-1.405

Are you vegan/vegetarian?	-1.294
Are you on a diet?	-1.261
Do you ever regret your eating behavior?	-1.064
What is your education?	-0.891
In what type of area do you live?	-0.626
How frequently do you use the computer?	-0.618
How much do you weigh?	-0.533
Are you active in a sports competition?	-0.507
What is your current type of employment?	-0.467
What is your race?	-0.298
Do you have a criminal record?	-0.277
What is your relationship status?	-0.227
How often do you work out?	-0.194
In the last 3 months, how many times have you been cited for traffic violations?	-0.163
Are you a member of an environmental organization?	-0.143
Have you ever been evicted?	-0.138
How frequently do you watch TV?	0.024
What is your favorite genre of music?	0.026
How often do you use public transportation?	0.028
What is your housing situation?	0.032
How much do you read?	0.044
What is your favorite hobby?	0.079
Do you carpool?	0.097
How long do you usually shower?	0.097
How often do you ride your bike?	0.130
Do you separate your household trash?	0.139
Who do you vote for?	0.305
What, if any, is your most prominent medical condition?	0.327
What kind of car do you own?	0.411
What is your sexual orientation?	0.457
Have you ever cheated in a relationship?	0.529
What is your monthly energy bill?	0.626
Do you ever download movies illegally?	0.630
What is your field of work?	0.638
How often do you watch pornography?	0.759
What is your religion?	0.801
In which news category are you most interested?	0.841
What is your current mobile data plan?	0.861
From which of these social services, if any, do you receive most benefit?	0.863
What kind of toilet paper do you use?	0.883
What is your household income?	0.967
What is your phone's voice and text plan?	1.007
How often do you have sex?	1.059
How much are you in debt?	1.104
In what size do you typically buy your beverages?	1.113
Which charity, if any, do give the most financial support to?	1.214
Do you or your partner use any type of birth control?	1.392
How many sexual partners have you had so far?	1.403
What is your most practiced sport?	1.458
What is your amount of savings?	1.736

**Table 13: Attributes of the energy-saving and healthy-living measures used in the two recommender systems.**

<b>Attributes of energy-saving measures</b>	<b>Attributes of health measures</b>
Initial effort	Emphasis (nutrition or exercise)
Continuous effort	Frequency of activity
Initial costs	Calories burned/avoided
Total savings (monetary)	Exercise intensiveness
Energy savings	Costs
Return on investment (time)	Duration
Additional environmental effects	Social impact
Comfort	

### *6.3.1 Preference rules*

If we want to use demographics as input for a MAUT-based recommender, then we need to link specific values of the demographics to the attribute weights. Specifically, we can use differences in attribute weights between different values of each demographic item to create “preference rules” for that demographic item. To generate the preference rules for a certain answer option to a certain demographic item, we first calculate the mean attribute weights of participants who chose that answer option. If certain answer options have almost the same attribute weights, they are grouped together. We then calculate the deviance of each attribute weight from the mean of the entire sample. Any deviance larger than a certain threshold is multiplied by a certain scaling constant, and entered into the “rules” table of the recommender system as a “preference update rule”. Based on simulated interactions, we found that the behavior of the recommender is most similar to previous versions of the system if the threshold is set to 0.125, and the multiplier is set to 6.0.

As an example, Figure 23 shows how the demographic item “age” is related to the attributes of the energy recommender. The age question had 7 answer options:

<20, 20-25, 26-30, 31-40, 41-50, 51-60, and >60. Some of these options have been grouped together. Figure 24 subsequently shows how the deviances that reach the threshold are turned into preference update rules that link the values of the age question to increases or decreases in attribute weights<sup>10</sup>. These rules make sure that when a user of the energy recommender (experiment = 1) answers the demographics question (action = demo) for age (obj = 26) with a certain answer (details = 1 through 7), their preference (model = utility) for a certain attribute (modelaspect) is increased or decreased by a certain amount (modelvalue). Note that grouped answer options (e.g., details = 1 and 2; details = 5, 6 and 7) have exactly the same rules. Moreover, not all answers update the preference value for every single attribute (e.g., answer option 1 does not update the attribute “savingsKWH” (energy savings)).

This procedure of using Mechanical Turk to uncover the link between demographics and preferences could easily be automated for use in commercial systems. Alternatively, such a system could bootstrap this link by tracking users’ choices, and updating the links based on the attribute weights of the items they select. However, since the presented items are based on the current link, one has to be careful not to create an inescapable positive feedback loop when employing this method. As a solution, extra weight could be given to counterfactual behaviors.

---

<sup>10</sup> These preference rules inherit their attributes (columns) from the generic “rules” construct in the system (which also manages item selection and other interactive behaviors). A rule triggers when an action is performed on a certain object (optionally with certain details) by changing the value of a certain aspect of a certain model. This explains the somewhat overly generic terminology used to describe the table of Figure 24.






Figure 23: Example of the values of energy-saving attributes for different age groups.

id	action	obj	details	model	modelaspect	modelvalue	experiment
3647	demo	26	1	utility	comfort	-2.93939394	1
2497	demo	26	1	utility	costonce	-2.80000000	1
3433	demo	26	1	utility	ecoweight	-2.79393939	1
2255	demo	26	1	utility	effortcont	2.37575758	1
2023	demo	26	1	utility	effortonce	-0.84242424	1
2746	demo	26	1	utility	realsavingsE	-1.14545455	1
3648	demo	26	2	utility	comfort	-2.93939394	1
2498	demo	26	2	utility	costonce	-2.80000000	1
3434	demo	26	2	utility	ecoweight	-2.79393939	1
2256	demo	26	2	utility	effortcont	2.37575758	1
2024	demo	26	2	utility	effortonce	-0.84242424	1
2747	demo	26	2	utility	realsavingsE	-1.14545455	1
3649	demo	26	3	utility	comfort	2.47619048	1
2499	demo	26	3	utility	costonce	-3.22857143	1
2025	demo	26	3	utility	effortonce	1.28095238	1
3196	demo	26	3	utility	returninv	4.23809524	1
2978	demo	26	3	utility	savingsKWH	-2.81904762	1
3650	demo	26	4	utility	comfort	-2.12121212	1
2500	demo	26	4	utility	costonce	5.65454545	1
3435	demo	26	4	utility	ecoweight	-0.88484848	1
2026	demo	26	4	utility	effortonce	-2.20606061	1
2748	demo	26	4	utility	realsavingsE	2.12727273	1
3197	demo	26	4	utility	returninv	-2.69696970	1
2979	demo	26	4	utility	savingsKWH	2.01212121	1
3651	demo	26	5	utility	comfort	3.50980392	1
2501	demo	26	5	utility	costonce	1.25882353	1
3436	demo	26	5	utility	ecoweight	4.52156863	1
2257	demo	26	5	utility	effortcont	-1.65098039	1
2027	demo	26	5	utility	effortonce	1.77254902	1
2749	demo	26	5	utility	realsavingsE	-2.01176471	1
3198	demo	26	5	utility	returninv	-2.92156863	1
2980	demo	26	5	utility	savingsKWH	2.52549020	1
3652	demo	26	6	utility	comfort	3.50980392	1
2502	demo	26	6	utility	costonce	1.25882353	1
3437	demo	26	6	utility	ecoweight	4.52156863	1
2258	demo	26	6	utility	effortcont	-1.65098039	1
2028	demo	26	6	utility	effortonce	1.77254902	1
2750	demo	26	6	utility	realsavingsE	-2.01176471	1
3199	demo	26	6	utility	returninv	-2.92156863	1
2981	demo	26	6	utility	savingsKWH	2.52549020	1
3653	demo	26	7	utility	comfort	3.50980392	1
2503	demo	26	7	utility	costonce	1.25882353	1
3438	demo	26	7	utility	ecoweight	4.52156863	1
2259	demo	26	7	utility	effortcont	-1.65098039	1
2029	demo	26	7	utility	effortonce	1.77254902	1
2751	demo	26	7	utility	realsavingsE	-2.01176471	1
3200	demo	26	7	utility	returninv	-2.92156863	1
2982	demo	26	7	utility	savingsKWH	2.52549020	1

Figure 24: Example rules of the energy recommender when the user answers a question (database entries).

### 6.3.2 No privacy profiles; privacy tendency

We used the data regarding users' perceived risk (7-point scale) to determine the dimensional structure of users' privacy preferences. In contrast to our findings in Section 4.3, we find in this study that a unidimensional representation of users' privacy preferences is not an oversimplification. Using Exploratory Factor Analysis, we find that the first extracted factor has an eigenvalue of 28.0 (for the energy-saving recommender) or 25.2 (for the healthy-living recommender), with the eigenvalue of the second factor being 3-4 times lower. Moreover, any attempt to construct a multidimensional model resulted in poorly discriminant factors (with factor correlations of around 0.90). Consequently, for Study 2 we opted to model users' *privacy tendency* on a unidimensional scale, rather than their multidimensional *privacy profile*.

An advantage of the unidimensional nature of users' privacy preferences in this system is that we can use a Rasch model to track users' privacy tendency. Rasch models have been used extensively in adaptive systems (e.g., Computer-Adaptive Tests (Gershon 2005), such as the TOEFL and GRE tests), and therefore several efficient implementations of dynamically updating Rasch models exist.

## 6.4 Study 1: testing demographics vs. attribute-based recommenders

The goal of this study is to find out whether a demographics-based recommender can achieve accurate recommendations and improve user satisfaction compared to an attribute-based recommender. The study also compares users' trust

when using demographics-based versus attribute-based recommenders. Finally, it investigates whether these effects are moderated by domain knowledge.

#### *6.4.1 Study setup*

Participants recruited on Amazon Mechanical Turk (N=403; 197 females, 203 males, 3 not disclosed; median age: 32, ranging from 19 to 67) were asked to participate in a usability test of a new energy or health recommender system. To make privacy concerns a realistic scenario, the system was ostensibly provided by “software-coaches.com”, a fictitious software company (i.e., not the university researchers). Participants were randomly assigned to one of four conditions (see Manipulations). The study started with a short assessment of participants’ domain knowledge. They then watched a short demonstration video explaining how the recommender system works. At the end of the video, they were told that they needed to make a selection of good recommendations (in a “basket”) by iteratively inspecting the recommendations and answering additional questions or changing the attribute weights (depending on whether they use the attribute-based or demographics-based system). After interacting with the system for at least 5 minutes, they concluded the study by answering a post-experimental questionnaire (see Measurements). Throughout the study, participants were subjected to a number of attention checks. Based on these checks, 46 participants were removed from the sample.

### 6.4.2 Manipulations

The experiment employed a 2-by-2 between-subjects design. The *recommendation domain* was manipulated by showing users either the energy-saving or healthy-living recommender. The *preference elicitation method* was manipulated with an attribute-based and demographics-based PE method:

- In **attribute-based PE** (Figure 25) participants set the weight of 7-8 attributes according to their personal preferences. The importance of every attribute can be set by clicking either on the “--”, “-”, “+” or “++” signs. Changing the weights will update their recommendations. Participants are instructed to change the weights and inspect/choose recommendations iteratively.
- In **demographics-based PE** (Figure 26) participants answer the 57 multiple-choice demographics questions in random order. They have the option to skip each question should they so desire. Answering the questions changes internal attribute weights and thus updates the recommendations. If users exhaust the list of demographics questions, the skipped questions are presented again, in order of appearance, with a notice that all other questions have been exhausted.

Software-Coaches.com
Healthy Living Coach<sup>beta</sup>

### Indicate preference ?

Indicate your preferences on the right. You can adjust your preferences by clicking repeatedly.

less important	14%	Exercise rather than nutrition	more important	+	+	+
less important	14%	Burn/avoid more calories	more important	+	+	+
less important	14%	Low exercise intensity	more important	+	+	+
less important	14%	Less frequent	more important	+	+	+
less important	14%	Shorter duration	more important	+	+	+
less important	14%	Low costs	more important	+	+	+
less important	14%	High social benefits	more important	+	+	+

### Choose measures ?

Here are your **recommendations**; select the measures you want to do, or you are already doing now.

Name	Focus	Calories	Exercise intensity	Frequency	Duration	Costs	Social benefits
Walk a National Trail together	exercise	700 cal				none	
Register at fitlink to find an exercise buddy	exercise	none				none	
Attend a nordic walking class together	exercise	500 cal				\$ 10.00	
Take a 1 hour walk together	exercise	350 cal				none	
Go to a spinning class with a friend	exercise	750 cal				\$ 10.00	
Prepare healthy meals three times this week	nutrition	none				none	
Find an exercise buddy	exercise	none				none	
Take turns with colleagues to bring fruit	nutrition	none				\$ 2.00	

### Your choices ?

Here are the measures you have chosen!

You have now spent 0 minutes using the system. After you click stop you will be asked a few more questions. At the end you can print your choices.

stop

**I want to do this:**  
You haven't chosen any measures yet.

**I can burn/avoid (weekly):** none

**I already do this:**  
You haven't chosen any measures yet.

**I am already burning/avoiding (weekly):** none

**I don't want to do this:**  
You haven't chosen any measures yet.

Figure 25: The healthy-living recommender with the attribute-based PE-method.

Software-Coaches.com
Energy Saving Coach<sup>beta</sup>

### Indicate preference ?

The recommendations will automatically update based on your answers to the questions on the right.

### What is your age?

< 20	20-25	26-30	31-40
41-50	51-60	> 60	skip this question

### Choose measures ?

Here are your **recommendations**; select the measures you want to do, or you are already doing now.

Name	Initial effort	Continuous effort	Initial costs	Savings US\$/year	Savings kWh/year	Return on investment	Env. effects	Comfort
Install LED lamps			\$ 105.00	\$ 121.20	537 kWh	11 months		
Configure your PC's power management			none	\$ 67.20	320 kWh	immediate		
Install a heat exchanger on ventilation ducts			\$ 1139.50	\$ 261.06	3055 kWh	5 years		
Cook on a gas stove instead of an electric stove			\$ 190.00	\$ 75.00	210 kWh	31 months		
Choose a utility with a green power program			none	none	0 kWh	immediate		
Turn off the PC when not in use			none	\$ 96.28	458 kWh	immediate		
Save up a full load of laundry			none	\$ 57.75	275 kWh	immediate		
Use a laptop instead of a PC			\$ 95.00	\$ 31.50	150 kWh	3 years		

### Your choices ?

Here are the measures you have chosen!

Show totals in ☒ US\$ ☐ kWh

You have now spent 0 minutes using the system. After you click stop you will be asked a few more questions. At the end you can print your choices.

stop

**I want to do this:**  
You haven't chosen any measures yet.

**I can save (yearly):** none

**I already do this:**  
You haven't chosen any measures yet.

**I am already saving (yearly):** none

**I don't want to do this:**  
You haven't chosen any measures yet.

Figure 26: The energy-saving recommender with the demographics-based PE-method.

### 6.4.3 *Measurements*

Any click made in the system was saved to the database of the experiment for subsequent analysis. Of specific interest are the number of items that people add to their list of recommendations and the “size” of these items (in terms of energy saved or calories burned/avoided). Moreover, for users of the demographics-based PE method, we are interested in their question answering/skipping behavior. The post-experimental questionnaire contained the questions listed in Table 14. All scales have been validated in at least two previous studies. In line with our framework for the user experience of (Knijnenburg et al. 2012c), we separate our constructs into Objective System Aspects (OSA), Subjective System Aspects (SSA), Experience (EXP), Interaction (INT) and Personal Characteristics (PC).

In line with our framework for the user experience of recommender systems (Knijnenburg et al. 2012c) and the uses and gratifications theory (McGuire 1974; Rubin 2002; Stafford et al. 2004) we consider both system satisfaction (a form of process gratification) and choice satisfaction (a form of content gratification).

**System satisfaction (EXP)** is a positive self-relevant evaluation of a system (Hassenzahl 2005). Satisfaction is not only determined by tangible aspects, such as service quality, but also by intangible ones, such as feelings of joy, fear, and frustration associated with the service experience (Johnson and Grayson 2005). The concept of system satisfaction relates to the perceived usefulness construct in TAM (Davis 1989) and the performance expectancy construct in UTAUT (Venkatesh et al. 2003). The questions used here were developed and extensively confirmed in a series of experiments (Knijnenburg et al. 2011, 2012c, 2014b; Knijnenburg and

Willemsen 2009, 2010) and applied directly in several privacy studies (Knijnenburg et al. 2013b; Knijnenburg and Kobsa 2013a, 2014; Kobsa et al. 2014).

**Choice satisfaction (EXP)** is a positive self-relevant evaluation of the *outcome* of using a system (Bechwati and Xia 2003; Bollen et al. 2010; Pedersen 2000). It is related to users' decision confidence (Hostler et al. 2005; Krishnan et al. 2008; Vijayasarathy and Jones 2001). Whereas system satisfaction provides users with an outcome expectation, choice satisfaction is an actual evaluation of those outcomes. The questions used here were developed and extensively confirmed in a series of experiments (Knijnenburg et al. 2011, 2012c, 2014b; Knijnenburg and Willemsen 2009, 2010). In the field of privacy and personalization similar constructs have been used, such as "value of personalization" (Chellappa and Sin 2005) and "personalization benefits" (Sutanto et al. 2013).

The concept of **perceived recommendation quality (SSA)** is also an integral part of our framework for the user experience of recommender systems (Knijnenburg et al. 2012c). This concept was developed and extensively confirmed in a series of experiments on recommender systems (Bollen et al. 2010; Knijnenburg et al. 2010, 2012a, 2012c; Willemsen et al. 2011).

Several researchers have argued that **understandability (SSA)** and **perceived control (SSA)** are important qualities of a recommender system (Czarkowski and Kay 2000, 2003; Herlocker et al. 2000; Kay and Lum 2005; Knijnenburg et al. 2012a; McNee et al. 2003; Tintarev and Masthoff 2011). We use questions from questionnaires developed in one of our studies on energy-saving

recommenders (Knijnenburg et al. 2011) and our study on social recommenders (Knijnenburg et al. 2012a).

**Trust in the provider (EXP)** measures users' privacy concerns regarding the company that provides the system. It is developed in Knijnenburg and Kobsa (2013a) and Sutanto et al. (2013) as a provider-specific analogy to the "trusting beliefs" factor of Malhotra et al. (2004), which is based on Jarvenpaa et al. (1999).

**General online privacy concern (PC)** is taken from directly from Malhotra et al. (2004) and Smith et al. (1996).

**Domain knowledge (PC)** was developed and extensively validated in our previous work on energy advice recommenders (Knijnenburg et al. 2011, 2014b; Knijnenburg and Willemsen 2009, 2010), and is here adapted to the health advice domain as well.

**Familiarity with recommenders (PC)** was developed in Knijnenburg et al. (2012a) and can be seen as a "trust-building factor" (cf. Chellappa and Sin 2005).

The questionnaire items were subjected to a Confirmatory Factor Analysis (CFA) with ordered categorical indicators and a weighted least squares estimator, estimating 9 factors. A separate CFA was conducted for the energy recommender and the health recommender. Items with low factor loadings, high cross-loadings, or high residual correlations were removed from the analysis. Factor loadings of the included items as well as the Average Variance Extracted (AVE) for each factor are shown in Table 14. AVEs are all adequate ( $AVE > 0.5$ ), indicating good convergent validity. Several factor correlations are however higher than the square root of the AVE, indicating that not all factors show sufficient discriminant validity (meaning



that some factors could be collapsed into a single factor). This is likely due to the small sample size in these studies (N~180 for each recommender). Previous work has demonstrated the discriminant validity of these factors, so we will maintain them as separate factors here as well.

**Table 14: AVEs and loadings for the questionnaire items in the two conditions of Study 1. (<domain> = “energy-saving” or “healthy-living”, <coach> “Energy Saving Coach” or “Healthy Living Coach”). Items without loadings were removed from the analysis.**

<b>Domain knowledge</b> Taken from Knijnenburg et al. (2011, 2014b) and Knijnenburg and Willemsen (2009, 2010). <i>AVE (E): 0.666; AVE (H): 0.570</i>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I know the energy consumption of all devices in my household / I know the effect of my diet and exercise routine on my health	0.709	
I understand difference between different types of <domain> measures	0.889	0.676
I know <domain> measures that most others haven't even heard of	0.738	0.604
I know which <domain> measures are useful to implement	0.914	0.740
I am able to choose the right <domain> measures	0.917	0.953
I sometimes doubt whether I chose good <domain> measures		
I don't understand most <domain> measures	-0.695	
<b>System satisfaction</b> Taken from Knijnenburg et al. (2011, 2014b) and Knijnenburg and Willemsen (2009, 2010); applied to privacy in Knijnenburg et al. (2013b), Knijnenburg and Kobsa (2013a, 2014), and Kobsa et al. (2014). <i>AVE (E): 0.679; AVE (H): 0.688</i>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
Overall, I am satisfied with the system	0.895	0.903
Using the <coach> made me happy	0.807	0.850
Using the <coach> is annoying	-0.744	-0.865
Using the <coach> was a pleasant experience	0.895	0.901
The <coach> was useless	-0.873	-0.875
I can make better <domain> choices with the <coach>	0.784	0.723
The <coach> made me more aware of my options	0.743	0.757
I can find better measures using the <coach>	0.777	0.675
The <coach> made me more energy-conscious	0.848	0.774
I would quickly abandon using the <coach>	-0.736	-0.826
I would use the <coach> more often if possible	0.838	0.903
I would recommend the <coach> to others	0.920	0.867
<b>Perceived recommendation quality</b> Taken from Bollen et al. (2010) and Knijnenburg et al. (2010, 2012a, 2012c). <i>AVE (E): 0.694; AVE (H): 0.717</i>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I liked the measures recommended by the <coach>	0.915	0.886
The recommended measures fitted my preference	0.874	0.897
The recommended measures were relevant	0.837	0.832
The system recommended too many bad measures	-0.718	-0.830
I didn't like any of the recommended measures	-0.775	-0.782

<b>Choice satisfaction</b> Taken from Knijnenburg et al. (2011, 2014b) and Knijnenburg and Willemsen (2009, 2010); related to “value of personalization” (Chellappa and Sin 2005) and “personalization benefits” (Sutanto et al. 2013). <div style="text-align: right;"><i>AVE (E): 0.674; AVE (H): 0.670</i></div>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I like the measures I’ve chosen	0.876	0.902
The chosen measures exactly fit my preference	0.798	0.785
I would recommend some of the measures I chose to others	0.836	0.715
I think I chose the best measures from the list	0.731	
I am excited about the measures I’ve chosen	0.855	0.860
How many measures will you implement?		
<b>Perceived control</b> Taken from Knijnenburg et al. (2011, 2012a). <div style="text-align: right;"><i>AVE (E): 0.652; AVE (H): 0.664</i></div>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I had limited control over the way the <coach> made recommendations		0.845
The <coach> does what I want	0.901	-0.895
The <coach> restricted me in my choice of measures	-0.714	0.757
I would like to have more control over the recommendations		
I had full control over the <coach>	0.797	-0.752
<b>Understandability</b> Taken from Knijnenburg et al. (2011, 2012a). <div style="text-align: right;"><i>AVE (E): 0.665; AVE (H): 0.616</i></div>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I understand how the <coach> came up with the recommendations	0.898	0.781
The <coach> is difficult to understand		-0.758
I am unsure how the recommendations were generated	-0.764	-0.857
The recommendation process is clear to me	0.854	0.904
I understood how to indicate my preference	0.755	0.719
I looked primarily at the names of the measures, not at the attributes		
How difficult/easy was comparing measures?	0.840	0.734
How difficult/easy was stating your preference?	0.814	0.809
How difficult/easy was comparing attributes?	0.772	0.693
<b>Familiarity with recommenders</b> Taken from Knijnenburg et al. (2012a), related to “trust-building factors” (Chellappa and Sin 2005). <div style="text-align: right;"><i>AVE (E): 0.647; AVE (H): 0.709</i></div>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I am familiar with online recommender systems	0.950	0.932
My experience with recommender systems is limited	-0.747	-0.840
I have occasionally followed the advice of a recommender system	0.759	0.834
To me, recommender systems are a new phenomenon		
I use recommender systems on a regular basis	0.743	0.752
<b>Trust in provider</b> Taken from Knijnenburg and Kobsa (2013a) and Sutanto et al. (2013); provider-specific analogy to “trusting beliefs” (Jarvenpaa et al. 1999; Malhotra et al. 2004). <div style="text-align: right;"><i>AVE (E): 0.879; AVE (H): 0.897</i></div>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
I believe software-coaches.com (the company that runs this website) is trustworthy in handling my information	0.937	0.965
I believe software-coaches.com tells the truth and fulfills promises related to the information I provide	0.911	0.948

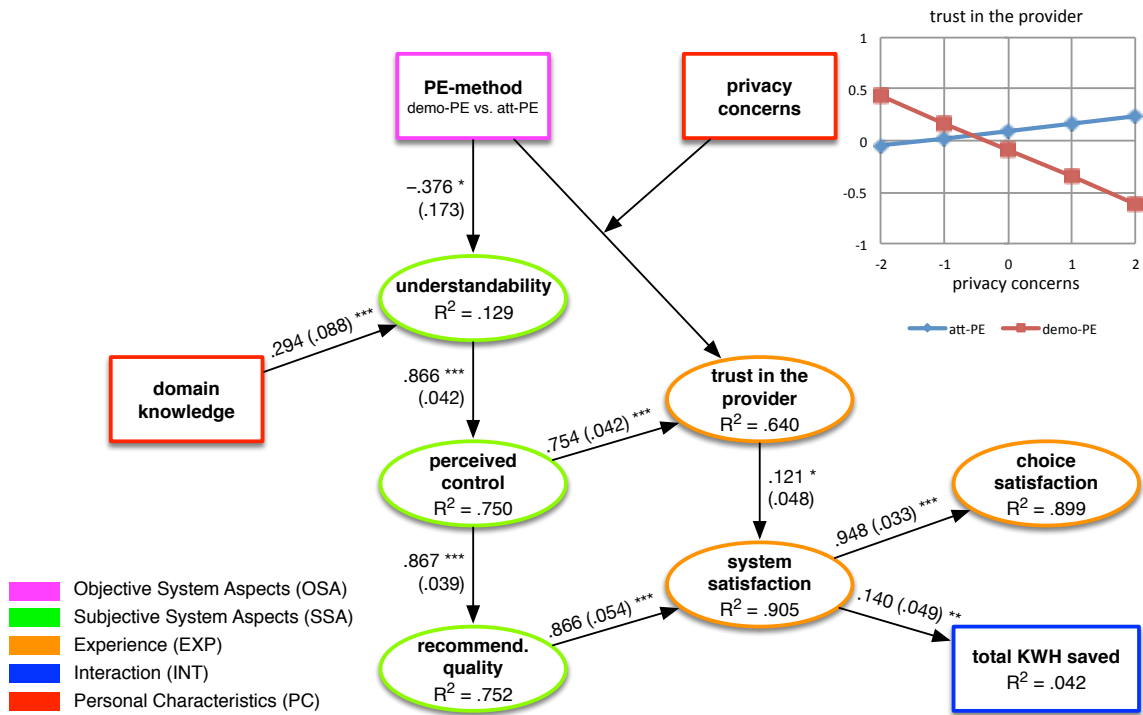
I believe software-coaches.com is predictable and consistent regarding the usage of my information	0.922	0.937
I believe software-coaches.com is honest when it comes to using the information I provide	0.954	0.953
I believe software-coaches.com keeps my best interests in mind when dealing with my information	0.963	0.932
<b>General online privacy concern</b> Taken from Malhotra et al. (2004) and Smith et al. (1996). <div style="text-align: right;"><i>AVE (E): 0.619; AVE (H): 0.759</i></div>		
<i>Items</i>	<i>Load. (E)</i>	<i>Load. (H)</i>
All things considered, the Internet would cause serious privacy problems		
Compared to others, I am more sensitive about the way online companies handle my personal information	0.750	0.883
To me, it is the most important thing to keep my privacy intact from online companies	0.799	0.923
I believe other people are too concerned with online privacy issues		
Compared with other subjects on my mind, personal privacy is very important	0.783	0.828
I am concerned about threats to my personal privacy today	0.812	0.846

#### 6.4.4 Results

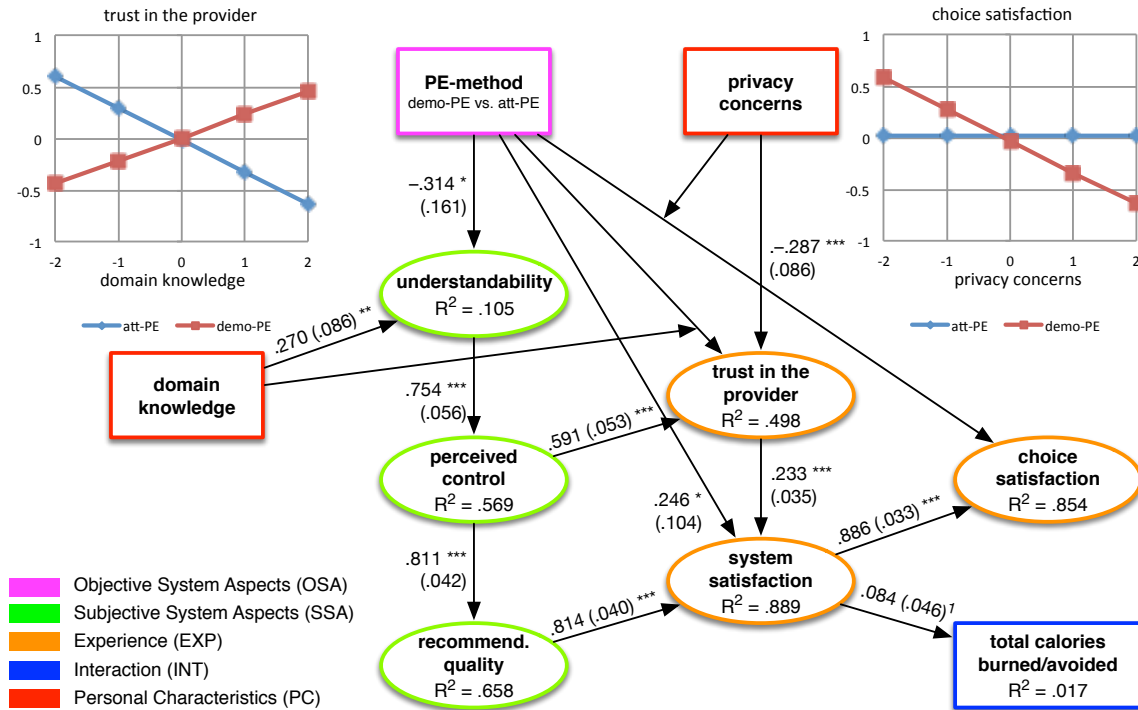
The subjective constructs, recommendation selection behaviors, and experimental manipulations were subsequently subjected to Structural Equation Modeling (SEM). The subjective constructs “Domain knowledge” and “General online privacy concern” are expected to interact with the PE-method manipulation. These factors were therefore turned into standardized weighted index scales. “Familiarity” was left out of the models, as it was not involved in any interesting effects.

Separate analyses were conducted for the energy-saving and healthy-living recommender. Figure 27 shows the resulting model for the energy-saving recommender, while Figure 28 shows the resulting model for the healthy-living recommender. Models were pruned to maintain significant effects that are

consistent across recommenders. The final models have a reasonably good<sup>11</sup> fit  
(Energy-saving recommender:  $\chi^2(805) = 1307$ ,  $p < .001$ ;  $RMSEA = 0.059$ , 90% CI:  
[0.053, 0.065],  $CFI = 0.974$ ,  $TLI = 0.972$ . Healthy-living recommender:  
 $\chi^2(881) = 1450$ ,  $p < .001$ ;  $RMSEA = 0.060$ , 90% CI: [0.055, 0.066],  $CFI = 0.965$ ,  
 $TLI = 0.963$ .



<sup>11</sup> For parsimonious results, we left out a few just-significant effects that were inconsistent across the two recommender systems. This resulted in a relatively high  $RMSEA$ .



**Figure 28: Structural Equation Model for the healthy-living recommender system. For more explanation, see Figure 27.**

The main structure of the two models is the same: the demographics-based system is overall less understandable ( $\beta = -.376$ ,  $p < .05$  in the energy-saving recommender and  $\beta = -.314$ ,  $p < .05$  in the healthy-living recommender, see Figure 27 and Figure 28 respectively) than the attribute-based system. Domain knowledge also has a positive effect on understandability, with domain experts understanding either system better than domain novices. Understandability is positively related to perceived control, which is in turn positively related to perceived recommendation quality. This is in line with our previous work on understandability and control in recommender systems (Knijnenburg et al. 2012a). Perceived control is also positively related to trust in the provider: participants who feel in control over the recommendation process are more trusting of the provider of the recommender


system. This is interesting, given that trust in this study was measured as a *privacy*-related construct, while control is a *recommendation*-related construct. This suggests a new variant of the “misplaced confidence” effect of control (Brandimarte et al. 2013): even perceptions of control that have nothing to do with privacy may increase users’ privacy-related trust in the system.

Both trust in the provider and recommendation quality positively influence participants’ system satisfaction. The former confirms the same consistent finding of our work on recommender systems (Knijnenburg et al. 2012c; Knijnenburg and Willemsen 2009, 2010), while the latter confirms our work on privacy (Knijnenburg and Kobsa 2013a; Kobsa et al. 2014). System satisfaction, finally, increases choice satisfaction and the total amount of energy saved (or calories burned/avoided) by the selected measures. This is again in line with our earlier work on recommender systems (Knijnenburg et al. 2012c, 2014b; Knijnenburg and Willemsen 2009, 2010).

Beyond the aforementioned effects, which are all consistent between the two recommenders, a few effects of the PE-method—in some cases interacting with privacy concerns or domain knowledge—are different between the two recommenders. In the energy-saving recommender, there is a significant interaction effect of PE-method and privacy concern on trust in the provider ( $\beta = -.330$ ,  $p = .044$ ). Specifically, people with high privacy concerns who use demographics-based PE seem to trust the provider less, while there is no such effect for the attribute-based PE-method.

In the healthy-living recommender, a similar interaction of PE-method and privacy concern affects participants’ choice satisfaction ( $\beta = -.324$ ,  $p = .042$ ). Again,

people with high privacy concerns who use demographics-based PE seem to be less satisfied with their choices, while there is no such effect for the attribute-based PE-method. Additionally, in the healthy-living recommender there is an interaction effect of PE-method and domain knowledge on trust in the provider ( $\beta = .506$ ,  $p = .009$ ). Specifically, domain novices seem to trust the provider more when using attribute-based PE, while domain experts seem to trust the provider more when using demographics-based PE. Note that when the indirect effects via understandability and control are also taken into account, the total effect is most pronounced for the demographics-based PE method, which is more trusted by experts and less trusted by novices (see Figure 29). Also note that a direct effect of the PE-method on system satisfaction cancels out some of the negative effect of demographics-based PE on understandability.



**Figure 29: Total effects of domain knowledge and PE-method on trust in the provider.**

#### 6.4.5 Discussion

What do these results mean? The lower understandability of the demographics-based PE-method may be due to the random order of the demographic information requests in this study. Since these questions cover a very wide range of demographics that are, importantly, not all equally useful in improving the preference model, users may be confused as to why the system would ask certain questions. Prioritizing questions that are typically perceived as more useful could potentially reduce this effect.

This may also be why novices are less trusting of the demographics-based PE-method in the healthy-living recommender: arguably, novices are less likely to come up with reasons for why certain questions are asked, and this may consequently reduce their trust in a system that asks demographic questions in a random order.

As for users with high privacy concerns, the (occasionally rather sensitive) demographic information requests reduce their trust in the provider (for the energy-saving recommender), or their choice satisfaction (for the healthy-living recommender). The latter may occur because they evaluate the outcome of the recommendation process against the information they had to disclose to generate those outcomes. Again, the root cause of this problem could very well be the fact that the system asks questions in a random order instead of prioritizing the less private and more useful items: Despite the fact that users have the explicit option to skip questions, privacy concerned users may still be irked by the occasional sensitive question. Using a request order that avoids these sensitive questions



would arguably reduce this effect. Note that in line with these findings the odds of disclosure for moderately concerned participants (+1 SD) were about 2.51 times<sup>12</sup> lower than for moderately unconcerned participants (−1 SD), a strongly significant difference ( $p = .006$ ).

Concluding, the demographics-based PE-method did not live up to its expectations, especially for domain novices and users with high privacy concerns. Given the nature of the observed effects, we argue that this is likely due to the random order of the demographic information requests, which disregards both the sensitivity of the requested demographic information and its contribution to the recommendation process.

## 6.5 Towards a better item request order

One of the main problems identified in study 1 is that asking demographics questions in a random order disregards the fact that:

- not all items are equally useful to the recommendation process;
- not all items are equally sensitive;
- not all users are equally private.

In this section we develop a number of alternative request orders that take one or more of these facts into account to create a more efficient and less privacy-sensitive demographics-based PE-method.

---


<sup>12</sup> This figure is based on the likelihood of answering any item (even those the user did not end up seeing). When we consider the *seen* items only, the odds ratio is even larger (3.26 times,  $p = < .001$ ).

### 6.5.1 *Choosing the domain: Health*

As we are to develop several alternative request orders (eventually leading to 8 experimental conditions), we chose to reduce the complexity of the second study by restricting our domain to healthy-living. Intuitively, we argue that the more useful demographic items for the healthy-living recommender are likely also the more sensitive ones. The trade-off between usefulness and privacy is thus arguably more prominent in the healthy-living domain.

Also, inspecting the results of study 1, we see a more interesting opportunity for improvement in the healthy-living recommender: For this recommender, both novices and highly concerned individuals have issues with the demographics-based recommender. This means that there is more room for improvement in this recommendation domain.

Finally, running a mixed logistic regression model with random intercepts for participants and demographic items, we observe that the variance in item disclosure levels in study 1 is higher in the healthy-living recommender (variance of random intercept: 2.041) than in the energy-saving recommender (variance of random intercept: 1.632). This is visually portrayed in Figure 30, which shows less of a “cusp” for the health recommender than for the energy recommender. This more even spread of item disclosure levels creates a better opportunity for a dynamic sensitivity/usefulness trade-off model to result in different types of experiences for different types of users.



**Figure 30: Ordered demographic item disclosure rates for the energy- and health-recommender.**

### 6.5.2 Possible request orders

We can make several possible improvements to a completely random request order. Given that not all items are equally useful to the recommendation process, one could prioritize requests by usefulness. First and foremost, this would make the recommendation process more efficient, thereby requiring fewer answers to demographic questions to get good recommendations. Based on the results of study 1 we can hypothesize a psychological benefit as well: If this *most-useful-first* request order prioritizes items that are ostensibly related to the recommendation domain, then this could increase the understandability of the system, and also increase novices' trust in the provider of the system.

Given that not all items are equally sensitive, one could alternatively prioritize requests by sensitivity. Assuming that most users do not go through *all* questions during their interaction with the system (an assumption that is correct for 78.5% of the study 1 participants), this *least-sensitive-first* request order would avoid asking users the most sensitive questions at all. This, in turn, could increase trust, disclosure, and choice satisfaction for users with high privacy concerns.

Optimizing both of these objectives simultaneously is difficult, though, because some of the sensitive items may also be very useful—this is likely to be the case for a healthy-lifestyle recommender. A system that makes a reasonable *trade-off* between these two objectives could potentially be efficient, understandable *and* trustworthy.

This leaves one crucial question to be answered: if the request order were to be a trade-off between usefulness and sensitivity, then how should each of these objective be weighted? Acknowledging that not all users are equally private, the best answer to this question would be “make it dependent on the user’s level of privacy”. An *adaptive request order* would entail tracking users’ disclosure tendency, and then allowing the system to ask more useful-but-sensitive questions to users with low privacy concerns, while strictly prioritizing the least sensitive questions for users with high privacy concerns.

In sum, we consider four generic strategies for prioritizing the order of demographic information requests:

- Most-useful-first
- Least-sensitive-first

- A static trade-off between usefulness and sensitivity (with a certain trade-off weight that is the same for all users, the value of which is still to be determined)
- An adaptive request order that sets the trade-off weight dynamically based on users' disclosure tendency (the starting value of this trade-off weight is still to be determined)

### 6.5.3 Defining a trade-off

Two of the four generic strategies defined in the previous subsection require a weighted trade-off between the two possibly conflicting attributes of usefulness and sensitivity. Such a trade-off can be conceptualized in several different ways. Bettman, Luce and Payne (1998) suggest *weighted adding* as one of the normatively most accurate trade-off strategies. This strategy is also in line with the overwhelming assumption in privacy research (Hui et al. 2007; Mothersbaugh et al. 2012; Nehf 2005) that privacy decisions are compensatory (cf. the word “privacy calculus” suggests that people trade-off privacy and usefulness in a calculative manner). In weighted adding, demographics items would be ordered by a linear function of sensitivity and usefulness:

$$r_i = u_i - \alpha \delta_i$$

where  $r_i$  is the request-priority of the item,  $u_i$  is its usefulness,  $\delta_i$  is its sensitivity, and  $\alpha$  is the relative weight of sensitivity, which can either be static, or dynamically estimated for each user ( $\alpha_n$ ).

An alternative conceptualization of the trade-off strategy would be a *non-compensatory* trade-off. In this trade-off strategy, demographic items would have to meet a certain threshold level of one attribute, before optimizing the other attribute.<sup>13</sup> In our system, this could be implemented by selecting the most useful demographics item that has a sensitivity level below a certain threshold (i.e., it applies the most-useful-first ordering to these items). If no items below the threshold are left, the system could simply select the least sensitive item among the items that are left (i.e., it applies the least-sensitive-first ordering to the remaining items). Formulaically:

$$r_i = \begin{cases} u_i & \text{if } \delta_i < \alpha, \\ -\delta_i & \text{if } \delta_i > \alpha. \end{cases}$$

where  $\alpha$  is a threshold, which can either be static, or dynamically estimated for each user ( $\alpha_n$ ).

Although the non-compensatory trade-off strategy may seem like a less elegant solution at first, it has a few advantages over the weighted adding strategy. First, the non-compensatory strategy is in most cases computationally less intensive, because it requires the usefulness (which, as we will explain in the next subsection, depends on the user's current preference model, and thus has to be continuously updated) only to be calculated for demographics items that fall below the sensitivity threshold. Second, it may lead to a more robust request behavior, because it provides a guarantee for the maximum level of sensitivity that will be

---

<sup>13</sup> The first part of this strategy is essentially the *elimination by aspects* strategy. Bettman, Luce and Payne (Bettman et al. 1998) note that elimination by aspects is often used as a first step in the decision process to reduce the list of potential options.

considered.<sup>14</sup> In our Rasch modeling framework (see Section 6.5.5), this guarantee can be translated into a “minimum disclosure probability” for requested items. Third, since the items above the threshold are requested in the order of increasing sensitivity (if they are requested at all), the non-compensatory trade-off strategy makes sure that users will always see the most sensitive item last (unless the threshold is higher than this most sensitive item). In the weighted adding strategy, the most sensitive item can conceivably be requested very early in the process—even when the sensitivity threshold is very low—as long as the item is deemed very useful. Finally, if we adapt the threshold to the user, the non-compensatory strategy defaults to the most-useful-first strategy for users with a very high threshold ( $\alpha_n > \max(\delta_i)$ ), and to the least-sensitive-first strategy for users with a very low threshold ( $\alpha_n < \min(\delta_i)$ ). Due to these advantages, we chose to implement the non-compensatory trade-off strategy in our recommender system.

#### 6.5.4 *Determining item usefulness $u_i$*

The usefulness of a demographic question to the recommendation process depends on how much it influences the attribute weights. The questions in our system vary in the extent to which they alter the weight of each attribute. Moreover, by answering consecutive questions, users’ preferences for certain attributes will become more pronounced (more extremely high or low weights) than for others (more moderate weights). Consequently, the usefulness of a demographic question

---

<sup>14</sup> That is, unless the user has run out of items below the threshold, at which point items are requested in order of sensitivity. An even stronger guarantee could be given if one would simply end the requests at this point.

for the recommendation process is a dynamic construct: if the user’s preference for a certain attribute is more pronounced, then a question that primarily influences this attribute’s weight is less important than a question that more strongly influences an attribute for which the user currently has a more moderate weight.

In our system, we define the usefulness of each demographics question ( $u_i$ ) as the weighted sum of the usefulness ( $u_o$ ) of each answer option ( $o_i$ ). The weight of an answer option is the probability ( $p_o$ ) at which that answer was given in study 1.<sup>15</sup> The usefulness of the answer option, in turn, is the weighted sum of the absolute amount of the change in model value ( $v_r$ ) of the “preference update rules” triggered by this answer ( $r_{oa}$ ). The weight in this case is the inverse of how much the weight of the attribute that would be updated ( $a$ ) deviates from the grand mean attribute weight ( $d_{an}$ ), plus a small regularization constant. Formulaically:

$$u_i = \sum_{o_i} p_o u_o$$

where

$$u_o = \sum_{r_{oa}} \frac{v_r}{d_{an}}$$

and

$$d_{an} = \text{abs}(w_{an} - \bar{w}_n) + .0001$$

The most useful item is the one with the highest  $u_i$ .

---

<sup>15</sup> We use this “marginal answering probability” as an estimate of how likely the user is to provide each answer. A more accurate estimate would be the probability of each answer to the current question *conditional* upon the user’s answers to all previous questions. Our current study 1 dataset is not large enough to robustly support this more accurate estimate, but it would be a conceivable improvement in commercial settings.



### 6.5.5 Determining item sensitivity $\delta_i$

Given that disclosure tendency is unidimensional,<sup>16</sup> we can create a Rasch model to determine the item sensitivities. A Rasch model defines the probability that user  $n$  discloses item  $i$  as follows:

$$p_{ni} = \frac{e^{\beta_n - \delta_i}}{1 + e^{\beta_n - \delta_i}}$$

where  $\beta_n$  is the user's disclosure tendency, and  $\delta_i$  is the item sensitivity. We calculate the item sensitivities pre-hoc, based on disclosure behaviors observed in the demographics-based PE condition of the healthy-living recommender of study 1.<sup>17</sup> Note that some items had a 100% disclosure rate; for these items we use the overall item sensitivity estimate (based on a Rasch model across both recommenders) and then subtracting 1 (this makes the estimated disclosure about half as likely).

To identify the Rasch model, either  $\beta_n$  or  $\delta_i$  needs to be anchored. We anchor  $\delta_i$  to have a mean of zero. The resulting item sensitivities are listed in Table 12.

### 6.5.6 Determining the threshold $\alpha$ (static, or based on user disclosure tendency)

The final component of the request order that needs to be defined is the threshold ( $\alpha$ ) of the trade-off-based strategies (see Section 6.5.3). This threshold can either be static, or dynamically estimated for each user ( $\alpha_n$ ).

---

<sup>16</sup> Analysis of the disclosure tendency in study 1 reconfirmed the result from our pre-study (see Section 6.3.2) that a single factor model is not an oversimplification.

<sup>17</sup> In a commercial setting, these values could be determined and/or updated on the fly.

The adaptive threshold can be based on the user's disclosure tendency,  $\beta_n$  in the Rasch model. To estimate  $\beta_n$  "on the fly", during the user's interaction with the system, we use the Normal Approximation Estimation Algorithm (PROX), a very lightweight method to estimate  $\beta_n$  with missing data<sup>18</sup> when item sensitivities are known (Cohen 1979). The PROX estimate for  $\beta_n$  in this particular case is given by:

$$\beta_n = \text{mean}_n(\delta) + \sqrt{1 + \text{var}_n(\delta)/2.9} * \ln\left(\frac{|D_n|}{|L_n| - |D_n|}\right)$$

where  $L_n$  is the set of items presented to user  $n$ ,  $D_n$  is the subset of items disclosed by user  $n$ , and  $\text{mean}_n(\delta)$  and  $\text{var}_n(\delta)$  are the mean and variance of the sensitivity of the items  $L$  presented to user  $n$ :

$$\text{mean}_n(\delta) = \left(\sum_{i \in L_n} \delta_i\right) / |L_n|$$

$$\text{var}_n(\delta) = \left(\sum_{i \in L_n} (\delta_i - \text{mean}_n(\delta))^2\right) / (|L_n| - 1)$$

For memory efficiency reasons, we do not want to keep track of the actual sets  $L_n$  and  $D_n$ . Instead, we keep track of  $|L_n|$ ,  $|D_n|$ ,  $\text{mean}_n(\delta)$ , and  $\text{var}_n(\delta)$  by creating recursively-updating versions of these metrics:

$$|L_n|_{t+1} = |L_n|_t + 1$$

$$|D_n|_{t+1} = \begin{cases} |D_n|_t & \text{if } d_{n,t} = 0, \\ |L_n|_t + 1 & \text{if } d_{n,t} = 1. \end{cases}$$

$$\text{mean}_n(\delta)_{t+1} = \frac{\text{mean}_n(\delta)_t * |L_n|_t + \delta_{n,t}}{|L_n|_t + 1}$$

---

<sup>18</sup> Allowing for missing data is crucial, because we estimate  $\beta_n$  on the fly, meaning that there is always a subset of demographic items for which the user's disclosure behavior is not known, because they have not been asked yet.

$$\text{var}_n(\delta)_{t+1} = \frac{|L_n|_t - 1}{|L_n|_t} \text{var}_n(\delta)_t + \frac{(\delta_{n,t} - \text{mean}_n(\delta)_t)^2}{|L_n|_t + 1}$$

with  $\text{var}_n(\delta)_1 = 0$  by definition.

One final problem to be solved is the regularization of  $\beta_n$ . If user  $n$  initially discloses either none or all of the items (one of which has to happen by definition at  $t = 1$ , but possibly even after that), then  $\beta_n$  resolves to  $\pm\infty$ . We therefore slightly adjust the model with two regularization parameters,  $D_{reg}$  and  $L_{reg}$ :

$$\beta_n = \text{mean}_n(\delta) + \sqrt{1 + \text{var}_n(\delta)/2.9} * \ln\left(\frac{|D_n| + D_{reg}}{|L_n| + L_{reg} - |D_n| - D_{reg}}\right)$$

Based on extensive simulations, we choose  $L_{reg} = 3$  and  $D_{reg} = L_{reg} * \bar{p}$ .

Where  $\bar{p}$  is the average disclosure probability over all users and all items in the health recommender in study 1, which is 0.9031. The regularization parameter thus adds 3 items at the average level of disclosure to the model.

We can now base the threshold  $\alpha_n$  on  $\beta_n$ . We don't want to set  $\alpha_n = \beta_n$ , because in that case the most sensitive item below the threshold has a disclosure probability of only 50% (since  $e^0/(1 + e^0) = 0.5$ ). As we try to prevent non-disclosures from happening at all, we prefer a (much) more conservative threshold. After extensive simulations, we choose the following two thresholds:

$$\alpha_n^H = \beta_n - 1.5$$

and

$$\alpha_n^L = \beta_n - 2.5$$

Assuming that  $\beta_n$  is accurately estimated, the user discloses items below the  $\alpha_n^H$  threshold with a probability of at least 81.8% ( $e^{1.5}/(1 + e^{1.5}) = 0.818$ ), and items below the  $\alpha_n^L$  threshold with a probability of at least 92.4%.

Table 15 shows the results of simulations with the two dynamic thresholds on simulated users disclosing 35%, 51%, 75% and 91% of the items. Of note are the rejections in the  $\delta_i < \alpha$  phase (when the most-useful-first strategy is employed on items that are supposed to all fall below the threshold) and the disclosures in the  $\delta_i > \alpha$  phase (when the least-sensitive-first strategy is employed on items that are supposed to all fall above the threshold). As expected, the high threshold results in more “early rejections” but fewer “late disclosures” than the low threshold. Overall, though, the models fairly accurately reflect users’ disclosure tendencies, minimizing disclosures during the most-useful-first phase, and switching to the least-sensitive-first strategy just around the time when the subset of items that the user is willing to disclose is about to be depleted.

**Table 15: Simulations of model user behaviors under the two adaptive thresholds**

Disclosures	Threshold	Rejections in the $\delta_i < \alpha$ phase	Disclosures in the $\delta_i > \alpha$ phase
20 (35%)	Low	1 (5.0%)	7 (35.0%)
20 (35%)	High	2 (10.0%)	6 (30.0%)
29 (51%)	Low	1 (3.4%)	13 (44.8%)
29 (51%)	High	4 (13.8%)	2 (6.9%)
43 (75%)	Low	2 (4.7%)	5 (11.6%)
43 (75%)	High	5 (11.6%)	0 (0.0%)
52 (91%)	Low	2 (3.8%)	0 (0.0%)
52 (91%)	High	5 (9.7%)	0 (0.0%)

Finally, for the static threshold, we simply set  $\beta_n$  to the average disclosure tendency in the healthy-living recommender in study 1, which is  $\beta_n = \ln(\bar{p}/(1 - \bar{p}))$ . This leads to the two static thresholds  $\alpha^H = 0.732$  and  $\alpha^L = -0.268$ . Referring to Table 12, this means that 41 out of the 57 demographics items (72%) fall below the high threshold, while 18 items (32%) fall below the low

threshold. To wit, in the health recommender of study 1 almost all participants (97%) disclosed more than 18 items, while only about 51% of the participants disclosed more than 41 items. Participants in the low threshold condition are thus likely to end up in the least-sensitive-first fallback scenario after disclosing the first 18 items. This fallback scenario is less likely to occur in the high threshold condition.

With all aspects of the request order in place, we can now turn to study 2, where we investigate the effect of the different request orders on users' experience and behavior.

## **6.6 Study 2: testing different request orders**

The goal of this study is to find out if there is an effect of request order in the demographics-based PE-method on rate of disclosure, recommendation accuracy, privacy threat, and user satisfaction. The idea is that the system will be most effective if it prioritizes demographic items that are most likely to change the recommendations (most-useful-first request order), but that it will reduce users' privacy concerns if it prioritizes request that are least sensitive (least-sensitive-first request order). The system that makes a careful tradeoff between these two strategies is expected to be the most satisfying overall.

### **6.6.1 Study setup**

The study was conducted on the health recommender, as was argued in Section 6.5.1. Participants recruited on Amazon Mechanical Turk (N=672; 338 females, 328 males, 6 not disclosed; median age: 30, ranging from 18 to 70) were

randomly assigned to one of the 8 experimental conditions (see Manipulations).

Procedures were the same as in study 1 (see Section 6.4.1). In this study, 58 participants were removed from the sample based on attention checks.

### 6.6.2 Manipulations

The study employed a between-subjects design with 8 conditions. The baseline condition was the **attribute-based PE** method (see Figure 25). The remaining conditions used the **demographics-based PE** method (see Figure 26) with different request orders:

- **Most-sensitive-first:** the items are ordered by decreasing sensitivity (see Section 6.5.5)<sup>19</sup>
- **Least-sensitive-first:** the items are ordered by increasing sensitivity (see Section 6.5.5)
- **Most-useful-first:** the items are ordered by decreasing usefulness (see Section 6.5.4)
- **Static trade-off, low threshold:** the items are ordered most-useful-first for items with a sensitivity below  $\alpha^L$  (the 18 least sensitive items), and least-

---

<sup>19</sup> The most-sensitive-first condition is of course far from optimal, but we include it to establish an upper bound on users' perceived privacy threats. Moreover, Acquisti et al. (Acquisti et al. 2012) demonstrated that asking privacy-sensitive questions in a decreasing order of intrusiveness could increase overall levels of disclosure, because subsequent requests compare favorably to the previous more intrusive requests, and users will therefore be more likely to answer them. Disclosure may thus in fact be *higher* in the most-sensitive first condition. However, Acquisti et al. did *not* measure users' satisfaction with the disclosure procedure, which is likely to be lower in this scenario, or users' perceived privacy threat, which is likely to be higher. Moreover, unlike Acquisti et al., users in our study are free to stop answering questions at any point in time. The increased privacy threat may cause users to abandon the demographics disclosure part of the interface, which would result in a lower overall level of disclosure in our case.

sensitive-first for items above this threshold (the 39 remaining items; see Section 6.5.6)

- **Static trade-off, high threshold:** the items are ordered most-useful-first for items with a sensitivity below  $\alpha^H$  (the 41 least sensitive items), and least-sensitive-first for items above this threshold (the 16 remaining items; see Section 6.5.6)
- **Adaptive request order, low threshold:** the items are ordered most-useful-first for items with a sensitivity below  $\alpha_n^L$ , and least-sensitive-first for items above this threshold. The threshold is dynamically adapted to the user's disclosure tendency  $\beta_n$ :  $\alpha_n^L = \beta_n - 2.5$  (see Section 6.5.6)
- **Adaptive request order, high threshold:** the items are ordered most-useful-first for items with a sensitivity below  $\alpha_n^H$ , and least-sensitive-first for items above this threshold. The threshold is dynamically adapted to the user's disclosure tendency  $\beta_n$ :  $\alpha_n^H = \beta_n - 1.5$  (see Section 6.5.6)

### 6.6.3 Measurements

As in study 1, any click made in the system was saved to the database of the experiment for subsequent analysis. Of specific interest are the number of items that people add to their list of recommendations and the “size” of these items (in terms of calories burned/avoided). Moreover, for all users except those in the attribute-based PE condition, we are interested in their question answering/skipping (i.e., “information disclosure”) behavior.

The post-experimental questionnaire contained the questions listed in

Table 16. These scales measure the same constructs as in study 1, plus an additional scale for **perceived privacy threat**, which was developed in Knijnenburg and Kobsa (2013a) and Kobsa et al. (2014) as a system-specific analogy to the “collection” factor of the Internet Users Information Privacy Concerns scale (Malhotra et al. 2004), which is in turn adapted from the “collection” factor of the Concern For Information Privacy scale (Smith et al. 1996). Xu et al. (2008) similarly developed their “privacy intrusion” scale as a system-specific combination of the “collection”, “unauthorized secondary use” and “improper access” factors of the Internet Users Information Privacy Concerns scale. We limit ourselves to the “collection” factor.

The items were subjected to a Confirmatory Factor Analysis with ordered categorical indicators and a weighted least squares estimator, estimating 10 factors. Items with low factor loadings, high cross-loadings, or high residual correlations were removed from the analysis. Factor loadings of the included items as well as the Average Variance Extracted (AVE) for each factor are shown in



Table 16.

**Table 16: Questionnaire used in study 2. Items without a loading were removed from the analysis.**

<b>Domain knowledge</b>		<b>AVE: 0.605</b>
<i>Items</i>		<i>Loading</i>
I understand difference between different types of healthy-living measures		0.719
I know <domain> measures that most others haven't even heard of		
I know which <domain> measures are useful to implement		0.813
I am able to choose the right healthy-living measures		0.798
<b>System satisfaction</b>		<b>AVE: 0.718</b>
<i>Items</i>		<i>Loading</i>
Overall, I am satisfied with the system		0.872
Using the Healthy Living Coach made me happy		0.861
Using the Healthy Living Coach is annoying		-0.806
Using the Healthy Living Coach was a pleasant experience		0.882
The Healthy Living Coach was useless		-0.858
I can make better healthy living choices with the Healthy Living Coach		0.778
The Healthy Living Coach made me more aware of my options		0.834
I can find better measures using the Healthy Living Coach		0.779
The Healthy Living Coach made me more energy-conscious		0.807
I would quickly abandon using the Healthy Living Coach		-0.861
I would use the Healthy Living Coach more often if possible		0.924
I would recommend the Healthy Living Coach to others		0.890
<b>Perceived recommendation quality</b>		<b>AVE: 0.744</b>
<i>Items</i>		<i>Loading</i>
I liked the measures recommended by the Healthy Living Coach		0.926
The recommended measures fitted my preference		0.897
The recommended measures were relevant		0.889
The system recommended too many bad measures		-0.796
I didn't like any of the recommended measures		-0.796
<b>Choice satisfaction</b>		<b>AVE: 0.691</b>
<i>Items</i>		<i>Loading</i>
I like the measures I've chosen		0.851
The chosen measures exactly fit my preference		0.785
I would recommend some of the measures I chose to others		0.804
I think I chose the best measures from the list		
I am excited about the measures I've chosen		0.881
<b>Perceived control</b>		<b>AVE: 0.635</b>
<i>Items</i>		<i>Loading</i>
I had limited control over the way the Healthy Living Coach made recommendations		0.715
The Healthy Living Coach does what I want		-0.939
The Healthy Living Coach restricted me in my choice of measures		0.733
I had full control over the Healthy Living Coach		-0.781
<b>Understandability</b>		<b>AVE: 0.587</b>
<i>Items</i>		<i>Loading</i>
I understand how the Healthy Living Coach came up with the recommendations		0.872
The Healthy Living Coach is difficult to understand		-0.760
I am unsure how the recommendations were generated		-0.734
The recommendation process is clear to me		0.778
I understood how to indicate my preference		0.790
How difficult/easy was comparing measures?		0.703
How difficult/easy was stating your preference?		0.796
How difficult/easy was comparing attributes?		0.677

<b>Familiarity with recommenders</b>		<b>AVE: 0.678</b>
<i>Items</i>		<i>Loading</i>
I am familiar with online recommender systems		0.912
My experience with recommender systems is limited		-0.751
I have occasionally followed the advice of a recommender system		0.819
I use recommender systems on a regular basis		0.802
<b>Perceived privacy threat</b>		<b>AVE: 0.717</b>
Taken from (Knijnenburg and Kobsa 2013a; Kobsa et al. 2014), system-specific analogy to “collection” and “control” (Malhotra et al. 2004)		
<i>Items</i>		<i>Loading</i>
The Healthy Living Coach has too much information about me.		0.819
The Healthy Living Coach does not know anything I would be uncomfortable sharing with it.		-0.678
I felt tricked into disclosing more information about myself than I wanted.		0.874
The Healthy Living Coach has information about me that I consider private.		0.820
The Healthy Living Coach knows more about me than I am comfortable with.		0.906
The Healthy Living Coach has information about me that they could use to invade my privacy.		0.871
I find the questions the Healthy Living Coach asked me intrusive.		0.933
<b>Trust in provider</b>		<b>AVE: 0.922</b>
<i>Items</i>		<i>Loading</i>
I believe software-coaches.com (the company that runs this website) is trustworthy in handling my information		0.957
I believe software-coaches.com tells the truth and fulfills promises related to the information I provide		0.971
I believe software-coaches.com is predictable and consistent regarding the usage of my information		0.931
I believe software-coaches.com is honest when it comes to using the information I provide		0.977
I believe software-coaches.com keeps my best interests in mind when dealing with my information		0.965
<b>General online privacy concern</b>		<b>AVE: 0.634</b>
<i>Items</i>		<i>Loading</i>
All things considered, the Internet would cause serious privacy problems		
Compared to others, I am more sensitive about the way online companies handle my personal information		0.783
To me, it is the most important thing to keep my privacy intact from online companies		0.804
Compared with other subjects on my mind, personal privacy is very important		0.755
I am concerned about threats to my personal privacy today		0.841

#### 6.6.4 Results

The subjective constructs, recommendation selection behavior, and experimental manipulations were subsequently subjected to Structural Equation Modeling (SEM). The subjective constructs “Domain knowledge” and “General online privacy concern” were again turned into standardized weighted index scales and allowed to interact with the experimental manipulation. “Familiarity” was again

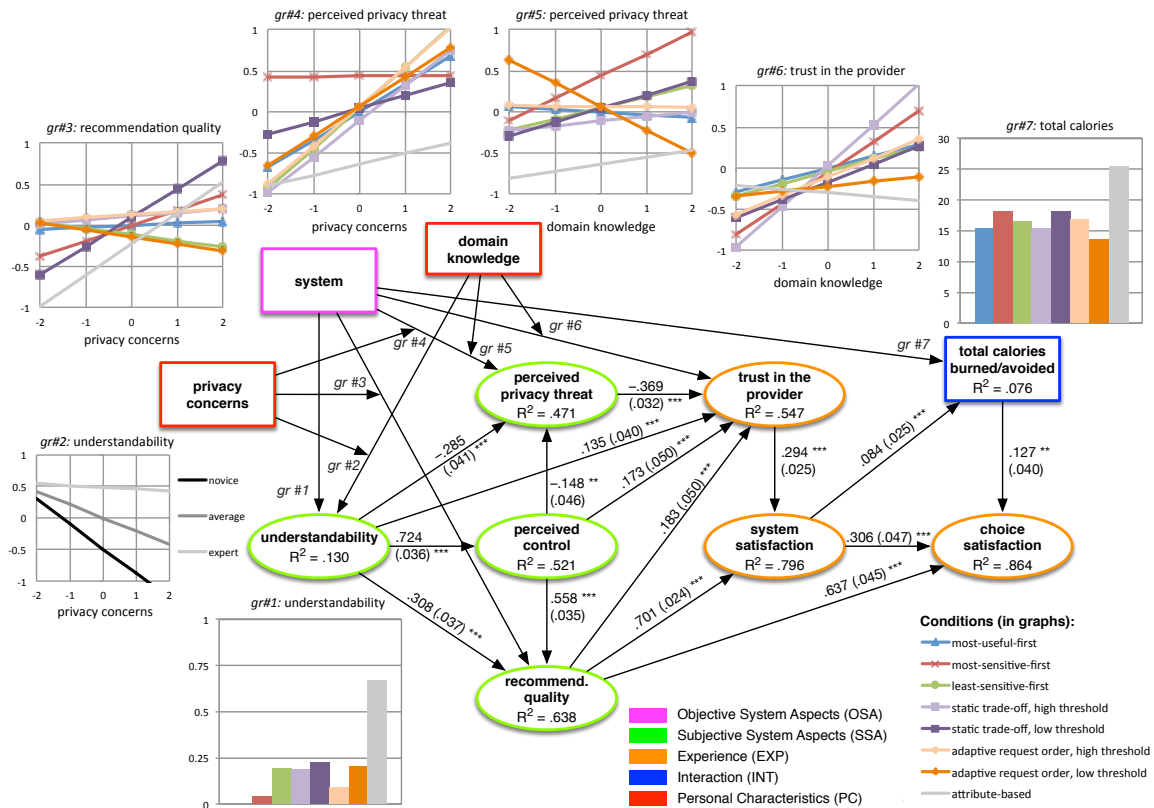
left out of the models, as it was not involved in any important effects.<sup>20</sup> Figure 31 shows the final model, which has an excellent fit:  $\chi^2(2009) = 3239, p < .001$ ;  $RMSEA = 0.032$ , 90% CI: [0.030, 0.034],  $CFI = 0.984$ ,  $TLI = 0.983$ . Furthermore, Figure 32–Figure 43 show the total effects of the experimental manipulations (crossed with domain knowledge and/or privacy concerns, where appropriate) on the subjective constructs and recommendation selection behavior. These total effects closely match the marginal effects, indicating that the model accurately represents the true effects.

Structurally, the model is similar to the models of study 1. Understandability is positively related to perceived control, which is in turn positively related to recommendation quality (cf. Knijnenburg et al. 2012a), although we now also find a direct effect of understandability on recommendation quality. Perceived control is positively related to recommendation quality and trust in the provider, and negatively to perceived privacy threat (there is also a direct negative effect of understandability on perceived privacy threat), which partially mediates the effect on trust in the provider (cf. Knijnenburg and Kobsa 2013a). Both trust in the provider and perceived recommendation quality positively influence system satisfaction (cf. Knijnenburg et al. 2012c; Knijnenburg and Kobsa 2013a; Knijnenburg and Willemsen 2009, 2010; Kobsa et al. 2014). System satisfaction in turn increases choice satisfaction and the total amount of calories burned/avoided;

---

<sup>20</sup> One notable finding for familiarity is that participants who are more familiar with recommender systems perceive simultaneously higher trust in the provider ( $\beta = .126, p < .001$ ), but also higher privacy threat ( $\beta = .141, p < .001$ ). Arguably, their previous experiences with recommender systems has made them more trusting, but also more aware of the potential privacy problems that recommender systems can cause.

the latter now also has a positive effect on choice satisfaction (cf. Knijnenburg et al. 2014b), as does recommendation quality.



**Figure 31: Structural Equation Model for study 2. Numbers on the arrows indicate standardized effect sizes (and standard errors), and p-values: <sup>1</sup>  $p < .10$ . \*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ . The main and interaction effects of the experimental manipulation are displayed in the graphs around the model.**

In terms of the experimental manipulation and its interaction with domain knowledge and privacy concerns, we find a rather large number of interesting effects, depicted in the graphs gr#1–gr#7 in Figure 31:

There is a significant main effect of PE-method on understandability ( $\chi^2(7) = 19.49, p = .007$ ). As gr#1 in Figure 31 shows, participants find the attribute-based PE-method more understandable than any of the demographics-based PE


conditions (all  $ps < .05$ ; no significant differences between other conditions). As Figure 32 shows, these effects on understandability also reflect on perceived control.

There is a significant interaction effect of domain knowledge and privacy concerns on understandability ( $\beta = .090, p = .040$ ). As gr#2 in Figure 31 shows, participants with high privacy concerns find the system less understandable unless they are domain experts, and domain novices find the system less understandable unless they have low privacy concerns.


There is a interaction effect of privacy concerns and PE-method on perceived recommendation quality ( $\chi^2(7) = 14.16, p = .048$ ), as depicted in gr#3 in Figure 31. First, let us consider the main effects: perceived recommendation quality seems generally higher in the demographics-based PE conditions than in the attribute-based PE condition (i.e., the grey line is shifted somewhat down compared to the others), especially when questions are asked most-useful-first (difference:  $p = .060$ ), most-sensitive-first ( $p = .032$ ), either of the static trade-off versions (high threshold:  $p = .008$ ; low threshold:  $p = .019$ ), and the adaptive request order with a high threshold ( $p = .004$ ). Note though, that this effect does not account for the positive effect on perceived recommendation quality mediated by understandability. When this effect is taken into account (as in the total effects graph in Figure 33), the attribute-based PE condition is on par with the others.

Furthermore, gr#3 shows that users with different levels of privacy concern perceive significantly different levels of recommendation quality in the different experimental conditions. Specifically, perceived recommendation quality increases

with privacy concerns for the attribute-based PE-method (slope:  $p = .012$ ) and the static trade-off condition with a low threshold (slope:  $p = .050$ ).



**Figure 32: Total effects of the manipulation on perceived control.**




**Figure 33: Total effects of the interaction of the manipulation and privacy concerns on perceived recommendation quality.**


There are significant interactions of privacy concerns and PE-method, and domain knowledge and PE-method, on perceived privacy threat (interaction with privacy concerns:  $\chi^2(7) = 21.96, p = .003$ ; interaction with domain knowledge:  $\chi^2(7) = 25.25, p < .001$ ). Starting with the main effects, gr#4 and gr#5 in Figure 31 both clearly show that participants perceive significantly less threat in the attribute-based PE condition (all  $ps < .001$ ), and significantly more threat in the most-sensitive-first condition (all  $ps < .01$ ), compared to all other conditions.

Gr#4 shows that users with different levels of privacy concern perceive significantly different levels of privacy threat in the different experimental conditions. Specifically, people with higher privacy concerns generally perceive more threat, except in the attribute-based PE condition (for which threat is low regardless of concerns), the most-sensitive-first condition (for which threat is high

regardless of concerns), and the static trade-off condition with a low threshold (for which threat is moderate regardless of concerns; all other slopes:  $p_s < .001$ ). The total effects reflect these findings (see Figure 34), and they carry over to trust (Figure 36), system satisfaction (Figure 38) and choice satisfaction (Figure 40) as well. Gr#5 further shows that novices and experts also perceive significantly different levels of privacy threat in the different experimental conditions. Specifically, novices perceive more threat from the adaptive request order with a low threshold (slope:  $p = .005$ ), while experts perceive more threat from the most-sensitive-first request order (slope:  $p = .014$ ). The total effects reflect these findings (see Figure 35).



**Figure 34: Total effects of the interaction of the manipulation and privacy concerns on perceived privacy threat.**




**Figure 35: Total effects of the interaction of the manipulation and domain knowledge on perceived privacy threat.**

There is a significant interaction effect of domain knowledge and PE-method on trust in the provider ( $\chi^2(7) = 17.43, p = .015$ ). There is not much of a main effect,




other than that participants in the attribute-based PE condition are somewhat less trusting; an effect that is canceled out once indirect effects (predominantly through perceived privacy threat) are taken into account (see Figure 36 and Figure 37).


Gr#6 in Figure 31 shows that the interaction effect mainly causes differences for experts. Specifically, experts are more trusting of the most-sensitive-first request order (slope:  $p = .005$ ) and the static trade-off with a high threshold (slope:  $p = .002$ ). As the total effects graph in Figure 37 shows, the most-sensitive-first request order ends up on the low end of the trust scale once indirect effects through perceived privacy threat are taken into account. The static trade-off with a high threshold seems favorable for experts, though, and this effect carries over to system satisfaction (Figure 39) as well.




**Figure 36: Total effects of the interaction of the manipulation and privacy concerns on trust in the provider.**




**Figure 37: Total effects of the interaction of the manipulation and domain knowledge on trust in the provider.**




**Figure 38: Total effects of the interaction of the manipulation and privacy concerns on system satisfaction.**



**Figure 39: Total effects of the interaction of the manipulation and domain knowledge on system satisfaction.**




**Figure 40: Total effects of the interaction of the manipulation and privacy concerns on choice satisfaction.**



**Figure 41: Total effects of the interaction of the manipulation and domain knowledge on choice satisfaction.**

Finally, there is a significant main effect of PE-method on recommendation selection behavior ( $\chi^2(7) = 29.26, p < .001$ ). Specifically, as gr#7 in Figure 31 clearly shows, participants in the attribute-based PE condition would burn/avoid substantially more calories with the measures they selected than in the other conditions (all  $ps < .01$ ). This main effect also dominates the total effects (see Figure 42 and Figure 43).



**Figure 42: Total effects of the interaction of the experimental manipulation and privacy concerns on total calories burned/avoided.**


**Figure 43: Total effects of the interaction of the experimental manipulation and domain knowledge on total calories burned/avoided.**

#### 6.6.5 Results for disclosure behavior

One variable of interest, disclosure behavior, was not included in the structural model presented above. This is because disclosure behaviors are only applicable to the demographics-based PE conditions. First, let us look at differences in disclosure behavior between the different request orders. Disclosure behavior is a repeated measure (57 measurements per participant; one for each demographics item) that can take on three values:

- The system showed the item to the participant, and participant disclosed the item.
- The system showed the item to the participant, and the participant skipped the item.
- The system did not show the item to the participant (or rather, the participant ended the interaction with the system before making a decision about the item).

Figure 44 shows the average percentage of items seen and disclosed per request order: both colored parts have been seen: the darker part has been disclosed, while the lighter part was skipped. Participants see the fewest items in the most-useful-first condition, and the most in the least-sensitive-first condition (comparison with most-useful-first:  $p = .010$ ) and the static trade-off with a low threshold condition (comparison with most-useful-first:  $p = .008$ ). Of the items that they see, participants disclose the most in the static trade-off with a high threshold condition (comparison with most-useful-first:  $p = .009$ ), and the least in the most-sensitive-first condition (comparison with most-useful-first:  $p = .019$ ). However, the overall level of disclosure is highest in the least-sensitive-first condition (comparison with most-useful-first:  $p = .004$ ) and the static trade-off with a low threshold condition (comparison with most-useful-first:  $p = .007$ ).




**Figure 44: Average percentage of items seen (colored part) and disclosed (darker part) per request order.**

To investigate the effect of disclosure behavior on the rest of the model, we created a separate Structural Equation Model without the attribute-based PE condition. The resulting model has an excellent fit ( $\chi^2(2001) = 2929, p < .001$ ;  $RMSEA = 0.004$ , 90% CI: [0.004, 0.004],  $CFI = 0.986$ ,  $TLI = 0.985$ ). The model excludes the effects from the original model that were primarily driven by the attribute-based PE condition, but otherwise has the same structure as the original model. Figure 45 shows the model, highlighting the effects involving disclosure behavior. In the model, disclosure is measured as a percentage of all items.

Most interesting is the negative feedback loop involving trust, threat, and disclosure. Specifically, trust in the provider increases participants' disclosure tendency, but this in turn also increases their perceived threat. Aside from this, we find that disclosure increases the perceived quality of the recommendations, but only for participants with high privacy concerns (see gr#1 in Figure 45). For participants with low privacy concerns, an increase in disclosure actually leads to *lower* perceived recommendation quality. This is likely due to the fact that disclosure yields diminishing returns; after a few disclosures the user's preferences are adequately represented, and further disclosures do not further increase the perceived recommendation quality. In other words: the more careful disclosure behavior of concerned users may be more efficient than the more lavish behavior of unconcerned users.

Finally, there is a significant interaction effect of domain knowledge and PE-method on disclosure behavior ( $\chi^2(6) = 15.77, p = .015$ ). As mentioned above, participants disclose most in the least-sensitive-first condition (comparison with

most-useful-first:  $p = .006$ ) and the static trade-off with a low threshold condition (comparison with most-useful-first:  $p = .01$ ). Also, experts disclose more in the static trade-off with a low threshold condition (slope:  $p = .019$ ).



**Figure 45: Structural Equation Model for study 2, including disclosure behavior. The attribute-based PE condition is left out of this model, and only the effects involving disclosure behavior are highlighted.**

#### 6.6.6 Discussion

In study 1 we found that the demographics-based PE-method had several shortcomings compared to the attribute-based PE-method, especially for novices and users with high privacy concerns. We hoped to remedy these shortcomings by introducing different request orders, arguing that a request order that adapts itself to the user would result in the highest overall satisfaction. Unfortunately, the adaptive request order did not clearly stand out in study 2, and generally speaking the attribute-based PE-method still excels, especially for novices and users with

high privacy concerns. That said, the results of study 2 show several interesting effects that suggest promising venues for future improvements.

First of all, we hoped that putting useful, non-sensitive demographic questions first would make the system intuitively more understandable. Unfortunately, changing the request order did not significantly increase understandability: users still find the attribute-based PE-method easier to understand. This is especially problematic for concerned and novice users, as they find the system less understandable overall. One alternative way to improve the understandability of the demographics-based PE-method would be to *explain* why certain questions are asked. Wang and Benbasat (2007) demonstrate that this may also increase users' trust in the system. Note though, that the explanations used in the study presented in Section 3.2 actually *decreased* trust, satisfaction, and disclosure. A possible solution to this problem is to employ multiple types of justifications, and to adapt the justification type to the user (as demonstrated in Section 5.2).

Manipulating the request order *did* have an effect on recommendation quality. In line with our expectations, request orders that trade off usefulness and sensitivity with a high threshold (thereby giving more weight to usefulness) result in a higher recommendation quality than the low-threshold versions. Interestingly, though, the most-useful-first request order did not result in the highest level of recommendation quality. This can be explained by looking at disclosure behavior: in the most-useful-first condition, users stop answering questions sooner than in the other conditions. Arguably, this is because the recommendations attain an

acceptable level of quality rather quickly in this condition. This may mean, however, that users stop *too soon*, possibly due to the combined effect of 1) one or more sensitive questions popping up and the 2) recommendations being already “good enough” (yet not as good as they could be; a behavior called “satisficing”, cf. Willemsen et al. (2011)). Consequently, they may have inadvertently ended up with worse recommendations than users in the other conditions. The trade-off conditions with a high threshold, on the other hand, arguably create a request order that is efficient enough to get good recommendations quickly, yet non-sensitive enough to encourage users to continue answering questions even when the recommendations are already “good enough”.

An interesting solution to users stopping the preference elicitation too soon is to encourage them to answer (or at least review) more questions. It is important not to overdo this, though, because ideally one would like users to end the preference elicitation when all remaining items are deemed too sensitive anyway. In future work we could implement an *adaptive nudge* that encourages users to review more questions until the remaining questions are all above the sensitivity threshold.

Considering the total effects on recommendation quality, Figure 33 shows that the attribute-based PE-method and the static trade-off with a low threshold condition lead to the highest recommendation quality for people with high privacy concerns, while the least-sensitive-first condition, the static trade-off with a high threshold condition, and both adaptive request order conditions lead to the highest recommendation quality for people with low privacy concerns. The static trade-off with a low threshold condition may be better for users with high concerns because



the overall level of disclosure is highest in this condition (see Figure 44), and for concerned users the amount of disclosure has a significant positive impact on recommendation quality (see gr#1 in Figure 45). The fact that users with low privacy concerns are the only ones who can get better recommendation quality in the high threshold versions also stands to reason; they trust the system enough to deal with the occasional sensitive question and disclose an amount of information that is sufficient to increase the perceived recommendation quality.

Looking at disclosure in more detail, we already noted that it makes sense that users see fewer questions in the most-useful-first condition: this is arguably due to satisficing. It also stands to reason that users inspect more items in the least-sensitive-first and static trade-off with a low threshold conditions: both of these conditions show a subset of very non-sensitive items upfront. The static trade-off with a high threshold condition has the highest disclosure among the items that the user sees, which makes this the most efficient condition (in this case, “efficient” means “least requests wasted on items that users will not disclose anyway”). In terms of overall disclosure, though, the least-sensitive-first condition and the static trade-off with a low threshold condition are the best. In the latter case, this increased disclosure behavior in fact leads to a higher recommendation quality for users with high privacy concerns.

Turning to perceived privacy threat, it makes sense that the most-sensitive-first condition leads to the highest perceived threat: participants skipped the most demographics items in this condition, meaning that they encountered the most items that were so sensitive that they did not wish to disclose them. It is interesting

that even people with low privacy concerns perceive this high level of threat; in all other demographics-based PE conditions (except for the static trade-off with a low threshold condition) threat increases with privacy concerns. In the attribute-based PE condition threat is low regardless of privacy concerns, making this by far the least threatening condition for people with high privacy concerns. For people with low privacy concerns, any version will do, except for the most-sensitive-first version and the static trade-off with a low threshold condition.

We noted in study 1 that novices were less trusting of the demographics-based PE method, and we hoped that changing the request order would reduce their privacy threat and consequently increase their trust in the provider. Unfortunately, we find that novices hardly distinguish among the different request orders in terms of threat. This is especially disconcerting regarding the most-sensitive-first request order, which surprisingly results in lower threat levels for novices. We had hoped that novices and experts alike would be able to detect when the system is helping them versus just trying to get sensitive information out of them. Novices may become better adept at assessing levels of threat and trust if we give them (adaptive) justifications for each disclosure (Knijnenburg and Kobsa 2013a, 2013b; Wang and Benbasat 2007).

In terms of overall levels of trust, it seems that the attribute-based version is best for novices and people with high concerns, while the static trade-off with a high threshold condition is better for experts and people with low concerns (see Figure 36 and Figure 37). These effects carry over to satisfaction (see Figure 38 and Figure 39), and—to a very limited extent—to choice satisfaction (see Figure 40 and Figure

41). As mentioned before, this means that we did not attain our goal of creating a demographics-based PE-method that outperforms the standard attribute-based PE-method for novices or people with high privacy concerns. Note, though, that for people with high privacy concerns, the demographic-based PE method using a static trade-off with a low threshold performs almost as well as the attribute-based PE-method on most user experience measures (i.e., trust, system satisfaction, and choice satisfaction).

Finally, in terms of calories saved, the attribute-based PE-method is by far the most effective version. Arguably, users in the demographics-based PE conditions may be distracted by demographics questions, and thus end up spending less time selecting measures. A possible solution to this problem would be to create a *hybrid method* that employs an attribute-based PE-method, but seeds the system with a few demographics questions. Again, this hybrid method could be *adaptive*, in that it could ask more questions to novices (who may need more help setting the initial attribute weights) and privacy-unconcerned users (who do not mind answering more questions) before switching to the attribute-based PE-method.

## **6.7 Conclusion regarding the adaptive request order studies**

Concluding, we can select a “best” condition for each type of user:

- For novices, this would be the attribute-based PE-method.
- For experts, the demographics-based PE-method is best, as long as it employs the static trade-off request order with a high threshold.
- For people with low privacy concerns any method will do, except for the

most-sensitive-first request order, the static trade-off with a low threshold, and in some cases, the attribute-based version.

- For people with high privacy concerns, the attribute-based PE-method and the static trade-off request order with a low threshold are both good.

Unfortunately, the adaptive request orders did not end up among the “best” versions. We designed the adaptive versions to seamlessly combine the benefits of the two static trade-off versions (i.e., an inherent focus on usefulness for novices, a low threshold for concerned users, and a high threshold for unconcerned users), but somehow these benefits did not materialize. One possible reason why the static trade-off versions are still better may be that they provide a guaranteed upper bound on the sensitivity of items that fall below the threshold, while in the adaptive versions this threshold fluctuates as the system tries to estimate the user’s disclosure tendency. This fluctuation may still cause the occasional sensitive question to be asked relatively early on in the interaction. A possible remedy would be to put an upper and lower bound on the estimated threshold, commensurate with the low and high thresholds used in the static trade-off versions.

Other things that could improve the demographics-based PE-method are: *(adaptive) justifications* that help users understand why a certain question is being asked, *adaptive nudges* to encourage users to explore more demographics questions (as long as they fall below the estimated threshold), and creating an *(adaptive) hybrid recommender* that starts with a few demographics questions and then switches to the attribute-based PE-method. These and other improvements can be tested in future work.

## **CHAPTER 7: General conclusion**

### **7.1 Summary**

Early research on users' information disclosure behavior revolved around the privacy calculus: the idea that users would make a careful and objective tradeoff between perceived risks and benefits of disclosure (Culnan and Armstrong 1999; Culnan and Bies 2003; Laufer et al. 1973; Laufer and Wolfe 1977; Milne and Gordon 1993; Petronio 2002). It soon became apparent, though, that objective privacy decision making is a chimera: people's privacy decisions fall prey to all sorts of decision biases (Acquisti et al. 2009, 2012; Acquisti and Grossklags 2005, 2008; John et al. 2011; Johnson et al. 2002; Lai and Hui 2006; Tsai et al. 2010) and most privacy decisions are too complex (Antón et al. 2004; Cate 2006; Consumer Reports 2012; Kelley et al. 2010; Liu et al. 2011; Madejski et al. 2012; McDonald et al. 2009; Strater and Lipford 2008; Turow et al. 2005) for people to fathom. In effect, many people refrain from exploiting the provided transparency and control altogether (Adkinson et al. 2002; Berendt et al. 2005; Bergmann 2009; Besmer et al. 2010; Compañó and Lusoli 2010; Gross and Acquisti 2005; Harris 2001; Jensen et al. 2005; Kelley et al. 2010; Larose and Rifon 2007; Singleton and Harper 2002; Turner and Varghese 2002). Contemporary privacy scholars have therefore moved beyond the idea of transparency and control (Barocas and Nissenbaum 2009; Nissenbaum 2011; Solove 2013).

Nudging then appeared as an alternative to transparency and control, which relieves some of the burden of privacy decision making (Acquisti 2009; Balebako et

al. 2011; Wang et al. 2013, 2014). However, nudges in the form of justifications (Acquisti et al. 2012; Besmer et al. 2010; Egelman et al. 2009; Hui et al. 2007; Kobsa and Teltzrow 2005; Metzger 2006; Patil et al. 2011; Rifon et al. 2005; Xu et al. 2009) and request orders (Section 3.2) had disappointing effects.

In this dissertation I therefore argued that privacy scholars need to move beyond the “one-size-fits-all” approach to privacy embodied in both nudges and transparency and control. I argued that because of the high variability and context-dependency of people’s privacy decisions, nudges need to be *tailored* to the user and her context (Kobsa 2001; Wang and Kobsa 2007).

In several studies, I contextualized users’ privacy decisions by showing how disclosure depended on the person’s privacy profile, the type of information, and the recipient of the information (Chapter 4). Then, I presented the idea of a “Privacy Adaptation Procedure” and demonstrated its merit in Chapter 5. Finally, I tested a complete implementation of the Privacy Adaptation Procedure in Chapter 6.

Although the adaptive request order conditions in the final study of this dissertation did not result in the hypothesized benefits, other versions that automatically traded off usefulness and sensitivity of the items to be disclosed did indeed improve users’ experience. We can therefore still conclude that automatic means to relieve some of the burden of controlling one’s privacy settings is a promising endeavor. Future work may further improve the truly adaptive versions, so that this automatic method works optimally for all kinds of users.

The manipulation of demographic information request order as presented in the final study is a good example of “realistic empowerment”: by prioritizing useful

and less-sensitive items, it helps users reduce privacy threat without reducing their ultimate control. An improved adaptive version would have the added benefit of not making any moral judgments about what the “right” level of privacy should be for each individual user (arguably, the “staticness” of the trade-off is what currently caused different types of users to prefer different versions of the static trade-off system). I believe that personalized advice is the best way forward to support people’s privacy decisions in an increasingly complex landscape of online social and commercial applications that gather and use all sorts of private information.

From a manager’s or a designer’s perspective, some additional practical factors need to be taken into account when implementing the Privacy Adaptation Procedure. Moreover, researchers are encouraged to test additional applications and variations of the Privacy Adaptation Procedure. Below I therefore outline some practical considerations and several venues for future work.

## **7.2 Practical considerations**

### *7.2.1 Implementation scenarios for privacy adaptation*

An important practical consideration for personalized privacy is who should provide the adaptation procedure. Users may rightfully be skeptical of a Privacy Adaptation Procedure that is provided by the very company that they perceive as trying to invade their privacy. Who is to say that the procedure acts honestly and ethically in applying the supposedly personalized nudges, and does not instead nudge users a little more towards sharing more information?

Because of this inherent conflict of interest, the Privacy Adaptation Procedure mainly works in systems where the service controlling the procedure is separate from the recipient of the personal information. Examples are a social network (service) that recommends how to share with your friends (recipient), an app store (service) that recommends what permissions to give to specific apps (recipient) or a browser (service) that recommends what information to disclose on a web form of a certain website (recipient).

This division of interests also prevents an essential “catch-22” inherent in privacy adaptation: the goal of a Privacy Adaptation Procedure is to reduce users’ concerns with the collection of personal information, but the procedure needs to learn the users’ privacy profile (via behavioral tracking or inference on demographics) to adapt its practices. This additional tracking may defeat the very purpose of the procedure. Note that this problem does not occur in systems that are meant to collect information on the fly anyway, such as a location-sharing service or a recommender system. In such systems no additional information needs to be collected: the incoming information is simply used to “throttle” further disclosure decisions.

More generally speaking, the Privacy Adaptation Procedure reverses the control sequence and thus the power balance of traditional privacy scenarios: rather than the control being in the hands of the service—who may then decide to deliver it to its users—the control is now in the hands of the user—who may delegate it to a service (possibly mediated by a third party that acts as a “privacy server”).



### *7.2.2 Privacy adaptation and the cold start problem*

But what if there is no data available to predict users' disclosure behavior? This is the so-called "cold start problem" (Schein et al. 2002). Some of our work shows that this may not be a very prominent problem in predicting users' disclosure behavior; when we ran a number of recommender algorithms on some of our disclosure data to predict users' disclosures, we found that the recommender is able to predict disclosure fairly accurately even when it only uses the preceding five (instead of all) items as a basis for learning (Wu et al. 2014).

If there is no data available initially, though, one could always use average levels of disclosure to make a recommendation (i.e., smart defaults) rather than personalized levels (i.e., adaptive defaults). A system can also always fall back on smart defaults if it lacks confidence about the user's disclosure profile.

### *7.2.3 How to gradually implement privacy adaptation*

I noted that users may sometimes show reactance towards recommendations (Fitzsimons and Lehmann 2004), which is especially true in domains where recommendations are not expected (Olson and Widing 2002; Pazzani and Billsus 2002). Users may thus not unanimously welcome privacy adaptation. A good solution to this problem is to start off with smart (not adaptive) defaults, and then gradually introduce personalized suggestions with good justifications (the latter is in line with Facebook's recently-introduced "Privacy Dinosaur" (Oremus 2014)). These suggestions help users, but give them explicit control over whether to accept them or not. Once users have gained sufficient trust in the Privacy Adaptation

Procedure, they may choose to accept subsequent suggestions automatically, thereby reducing their burden even further.

#### *7.2.4 Implications for the privacy calculus*

My work confirms recent findings that the privacy calculus (and the privacy-personalization paradox—the privacy calculus embodied in personalization scenarios) is a very poor descriptive framework for privacy decision-making (Kehr et al. 2013, 2015). People rarely take a truly calculative approach to privacy decision making, and are often prone to (or necessitated to) take mental shortcuts instead (Acquisti and Grossklags 2005; Angst and Agarwal 2009; Lowry et al. 2012; Wilson and Valacich 2012).

It is surprising that the privacy calculus has been accepted as the de facto standard descriptive theory of privacy decision making for such a long time. The reason may be that most research on privacy decision making either tests a very high-level conceptualization of the privacy calculus, or uses a statistical framework that simply assumes the privacy calculus to be correct (cf. Chellappa and Sin 2005; Hann et al. 2007; Ho and Tam 2006; Hui et al. 2006; Knight 2010; Olivero and Lunt 2004; Phelps et al. 2000; Xu et al. 2009, 2011). Neither of these approaches is able to invalidate the privacy calculus. Work that specifically tests privacy decision externalities (cf. Acquisti et al. 2012; Acquisti and Grossklags 2005, 2008; Adjerid et al. 2013; John et al. 2011; Johnson et al. 2002; Lai and Hui 2006; Tsai et al. 2010) usually ends up invalidating its core assumptions.

A separate yet related question is whether the privacy calculus is a valid *prescriptive* theory for good privacy decision making. In this dissertation I have implicitly answered this question positively in the implementation of the static and adaptive “trade-off” request orders, which automatically trade off the benefits and perceived risks of disclosure.

Arguably, the validity of the privacy calculus as a prescriptive theory depends on the specifics of its operationalization. In my implementation of the Privacy Adaptation Procedure, I have made several assumptions about the privacy calculus that need to be tested in future work. Specifically:

- How should disclosure risk be determined? My implementation relied on previous disclosure tendency (behavior) as a yardstick. Behavior may however be related to factor other than risk (e.g., it may be confounded with benefit). Individual risk perceptions (which may differ per user) or expert opinions (which come closer to a measure of “objective risk”) are valid alternative solutions.
- How should benefits be determined? My implementation relied on an objective benefits calculation, driven by the preference model of the recommender system. As recommenders can capitalize on unanticipated correlations between demographic information and preferences, this objective benefit may sometimes be quite different from the *perceived* benefit of disclosure. However, as users’ decisions are in most cases arguably governed by perceived benefit, this may create a conflicting situation. Giving users adequate explanations or justifications can reduce this conflict.

However, just as we can use either perceived or (more or less) objective versions of risk as input for the privacy calculus, we can similarly use *perceived* benefit as an alternative input for the privacy calculus.

- How should the trade-off between benefits and risks be modeled? My implementation used a non-compensatory threshold model, rather than the more commonly used weighted additive model. This turned out to be a good decision due to its predictably bounded behavior; in fact, one of our recommendations for future work was to curtail this model even more by adding a static upper bound to the dynamically defined threshold. Other operationalizations of the trade-off should however also be tested in future work.
- Finally, what other contextual variables may be included in the privacy calculus? Researchers increasingly agree that privacy decisions are highly context-dependent (Nissenbaum 2009), and this dissertation investigates several contextual variables that are important determinants of users' privacy decision behavior: the user ("who"), the information ("what"), and the recipient ("to whom"). My implementation of the Privacy Adaptation Procedure used a simple unidimensional preference tracking model, but a more complex model could result in more accurate results. Moreover, in other scenarios the recipient as well as other contextual variables may be included in the model as well.

Concluding, while the privacy calculus is not a good descriptive model of users' privacy decisions, I argue that it could be a useful prescriptive model. Its use

as a prescriptive model is however critically dependent on a correct operationalization of its parameters; something which has not adequately been investigated to date. This dissertation offers a first foray into this endeavor; future work should further improve upon these results to develop the most appropriate prescriptive model of users' privacy calculus.

### **7.3 Future directions**

As the Privacy Adaptation Procedure crucially depends on an adequate operationalization of (a prescriptive model of) the privacy calculus, the parameters of this operationalization (as discussed above) are a very important direction of future work. Beyond this, I envision several other extensions of the current work, which are listed below.

Although I extensively tested adaptive *request orders* in Chapter 6, the evidence regarding the benefits of adaptive *justifications* and adaptive *defaults* remains limited to the studies presented in Chapter 5. These studies have several limitations. Future work should test adaptive justifications and adaptive defaults in complete implementations as well.

Moreover, even the implementation of adaptive request orders described in Chapter 6 is still limited; a true test of the benefits of privacy adaptation would implement it in a scenario where users have to make difficult privacy decisions on a daily basis, e.g., in an app store or a social network. Future work should implement a Privacy Adaptation Procedure in such a setting. Preferably, the implementation

should have a large scale, so that the adaptation procedure can use sophisticated recommendation algorithms.

Finally, as I noted in Section 2.2, transparency and control are currently the norm in the policy landscape. Given that users' disclosure behaviors are highly context-dependent and their decision-making abilities limited, I believe that this paradigm is unrealistic. Can privacy adaptation be implemented in a regulatory framework? This would imply a "flexible standard" of privacy legislation that requires different measures to be taken in different contexts. A law in the spirit of privacy adaptation could for instance require online companies to improve the transparency and control of their sites/apps for users who want it, without mandating it in cases where transparency and control would merely cause confusion. Independent privacy studies could provide the regulatory "ground rules" for such a paradigm, e.g., in terms of reasonable default requirements (e.g., whether certain websites should provide opt-in or opt-out mechanisms) and on the extent to which sites should be able to overrule normative requirements (e.g., what kind of exceptions to the general privacy mandate on cookies should a website be able to request based on demonstrated negative user test results for their site). Moreover, in the spirit of food safety standards set by the FDA, online companies could test and validate the implications of their privacy practices on their user base as a way to "prove" that they meet minimal privacy standards. Future work should further study the idea of a regulatory framework in the spirit of privacy adaptation.

## REFERENCES

- Aagaard, M. 2013. "How Privacy Policy Affects Sign-Ups – Surprising Data From 4 A/B Tests," *ContentVerve.com*.
- Accenture. 2012. "Today's Shopper Preferences: Channels, Social Media, Privacy and the Personalized Experience," Accenture Interactive.
- Ackerman, M. S., Cranor, L. F., and Reagle, J. 1999. "Privacy in e-commerce: examining user scenarios and privacy preferences," in *Proceedings of the 1st ACM conference on electronic commerce*, EC '99, Denver, CO: ACM Press, pp. 1–8.
- Acquisti, A. 2004. "Privacy in electronic commerce and the economics of immediate gratification," in *Proceedings of the 5th ACM conference on Electronic commerce* New York, NY: ACM, pp. 21–29.
- Acquisti, A. 2009. "Nudging Privacy: The Behavioral Economics of Personal Information," *IEEE Security and Privacy* (7), pp. 82–85.
- Acquisti, A., and Grossklags, J. 2005. "Privacy and Rationality in Individual Decision Making," *IEEE Security & Privacy* (3:1), pp. 26–33.
- Acquisti, A., and Grossklags, J. 2008. "What Can Behavioral Economics Teach Us About Privacy?," in *Digital Privacy: Theory, Technologies, and Practices*, A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis (eds.), New York/London: Auerbach Publications, pp. 363–377.
- Acquisti, A., and Gross, R. 2006. "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook," in *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, G. Danezis and P. Golle (eds.), (Vol. 4258) Springer Berlin / Heidelberg, pp. 36–58.
- Acquisti, A., John, L. K., and Loewenstein, G. 2012. "The Impact of Relative Standards on the Propensity to Disclose," *Journal of Marketing Research* (49:2), pp. 160–174.
- Acquisti, A., John, L., and Loewenstein, G. 2009. "What is privacy worth?," in *Proceedings of the Twenty First Workshop on Information Systems and Economics* Phoenix, AZ, December.
- Adjerid, I., Acquisti, A., Brandimarte, L., and Loewenstein, G. 2013. "Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency," in *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, New York, NY, USA: ACM, pp. 9:1–9:11.
- Adkinson, W. F., Eisenach, J. A., and Lenard, T. M. 2002. *Privacy Online: A Report on the Information Practices and Policies of Commercial Web Sites*, Privacy & Freedom Foundation.
- Adomavicius, G., and Tuzhilin, A. 2011. "Context-Aware Recommender Systems," in *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor (eds.), Boston, MA: Springer US, pp. 217–253.
- Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*, Brooks/Cole Publishing Company, Monterey, CA.
- Andersen, A. 2000. *Internet Privacy Survey 2000: A Survey of the Privacy Practices of Australia's Most Popular Web Sites*, Andersen Legal.
- Angst, C. M., and Agarwal, R. 2009. "Adoption of electronic health records in the presence of privacy concerns: the elaboration likelihood model and individual persuasion," *MIS Q.* (33:2), pp. 339–370.
- Antón, A. I., Earp, J. B., He, Q., Stufflebeam, W., Bolchini, D., and Jensen, C. 2004. "Financial privacy policies and the need for standardization," *IEEE Security Privacy* (2:2), pp. 36–45.

- Ardagna, C. A., Capitani di Vimercati, S., and Samarati, P. 2011. "Privacy Models and Languages: Access Control and Data Handling Policies," in *Digital Privacy*, J. Camenisch, R. Leenes, and D. Sommer (eds.), (Vol. 6545) Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 309–329.
- Awad, N. F., and Krishnan, M. S. 2006. "The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled Online for Personalization," *MIS Quarterly* (30:1), pp. 13–28.
- Bagozzi, R. P., and Phillips, L. W. 1991. "Assessing construct validity in organizational research," *Administrative Science Quarterly* (36:3), pp. 421–458.
- Balebako, R., Leon, P. G., Mugan, J., Acquisti, A., Cranor, L. F., and Sadeh, N. 2011. "Nudging users towards privacy on mobile devices," in *CHI 2011 workshop on Persuasion, Influence, Nudge and Coercion Through Mobile Devices* Vancouver, Canada, pp. 23–26.
- Barocas, S., and Nissenbaum, H. 2009. "On notice: The trouble with Notice and Consent," in *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information*.
- Bechwati, N. N., and Xia, L. 2003. "Do Computers Sweat? The Impact of Perceived Effort of Online Decision Aids on Consumers' Satisfaction With the Decision Process," *Journal of Consumer Psychology* (13:1–2), pp. 139–148.
- Benisch, M., Kelley, P. G., Sadeh, N., and Cranor, L. F. 2011. "Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs," *Personal Ubiquitous Computing* (15:7), pp. 679–694.
- Berendt, B., Günther, O., and Spiekermann, S. 2005. "Privacy in E-Commerce: Stated Preferences vs. Actual Behavior," *Communications of the ACM* (48:4), pp. 101–106.
- Bergmann, M. 2009. "Testing Privacy Awareness," in *The Future of Identity in the Information Society*, V. Matyáš, S. Fischer-Hübner, D. Cvrček, and P. Švenda (eds.), (Vol. 298) Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 237–253.
- Besmer, A., Watson, J., and Lipford, H. R. 2010. "The impact of social navigation on privacy policy configuration," in *Proceedings of the Sixth Symposium on Usable Privacy and Security* Redmond, Washington, July, p. Article 7.
- Bettman, J. R., Luce, M. F., and Payne, J. W. 1998. "Constructive consumer choice processes," *Journal of consumer research* (25:3), pp. 187–217.
- Bhatnagar, A., Misra, S., and Rao, H. R. 2000. "On risk, convenience, and Internet shopping behavior," *Communications of the ACM* (43:11), pp. 98–105.
- Bicakci, K., Atalay, N. B., and Kiziloz, H. E. 2011. "Johnny in internet café: user study and exploration of password autocompleteness in web browsers," in *Proceedings of the 7th ACM workshop on Digital identity management* Chicago, IL, pp. 33–42.
- Binder, J., Howes, A., and Sutcliffe, A. 2009. "The problem of conflicting social spheres: effects of network structure on experienced tension in social network sites," in *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, New York, NY, USA: ACM, pp. 965–974.
- Böhmer, M., Bauer, G., and Krüger, A. 2010. "Exploring the Design Space of Context-aware Recommender Systems that Suggest Mobile Applications," in *2nd Workshop on Context-Aware Recommender Systems* Barcelona, Spain, September.
- Bokhove, W., Hulsebosch, B., Van Schoonhoven, B., Sappelli, M., and Wouters, K. 2012. "User Privacy in Applications for Well-being and Well-working," in *AMBIENT 2012, The Second International Conference on Ambient Computing, Applications, Services and Technologies*, September 23, pp. 53–59.



- Bollen, D., Knijnenburg, B. P., Willemsen, M. C., and Graus, M. 2010. "Understanding choice overload in recommender systems," in *Proceedings of the fourth ACM conference on Recommender systems* Barcelona, Spain, pp. 63–70.
- Bonneau, J., and Preibusch, S. 2010. "The Privacy Jungle: On the Market for Data Protection in Social Networks," in *Economics of Information Security and Privacy*, T. Moore, D. Pym, and C. Ioannidis (eds.), New York, NY: Springer US, pp. 121–167.
- Brandimarte, L., Acquisti, A., and Loewenstein, G. 2013. "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychological and Personality Science* (4:3), pp. 340–347.
- Brehm, J. W. 1966. *A Theory of Psychological Reactance*, Oxford, England: Academic Press.
- Brodie, C., Karat, C.-M., and Karat, J. 2004. "How Personalization of an E-Commerce Website Affects Consumer Trust," in *Designing Personalized User Experience for eCommerce*, C.-M. Karat, J. O. Blom, and J. Karat (eds.), Dordrecht, Netherlands: Kluwer Academic Publishers, pp. 185–206.
- Brown, C. L., and Krishna, A. 2004. "The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice," *Journal of Consumer Research* (31:3), pp. 529–539.
- Buchanan, T., Paine, C., Joinson, A. N., and Reips, U.-D. 2007. "Development of Measures of Online Privacy Concern and Protection for Use on the Internet," *Journal of the American Society for Information Sciences and Technology* (58:2), pp. 157–165.
- Bulgurcu, B. 2012. "Understanding the information privacy-related perceptions and behaviors of an online social network user," Ph.D. Thesis, Vancouver, BC: University of British Columbia.
- Bustos, L. 2012. "Best Practice Gone Bad: 4 Shocking A/B Tests," *GetElastic*.
- Butler, J. C., Dyer, J. S., Jia, J., and Tomak, K. 2008. "Enabling e-transactions with multi-attribute preference models," *European Journal of Operational Research* (186:2), pp. 748–765.
- Cate, F. H. 2006. "The Failure of Fair Information Practice Principles," in *Consumer Protection in the Age of the "Information Economy"*, J. K. Winn (ed.), Burlington, VT: Ashgate Publishing Company.
- Cavusoglu, H., Phan, T., and Cavusoglu, H. 2013. "Privacy Controls and Content Sharing Patterns of Online Social Network Users: A Natural Experiment," in *ICIS 2013 Proceedings* Milan, Italy.
- Chellappa, R. K., and Sin, R. 2005. "Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma," *Information Technology and Management* (6:2-3), pp. 181–202.
- Chen, J., Geyer, W., Dugan, C., Muller, M., and Guy, I. 2009. "Make new friends, but keep the old: recommending people on social networking sites," in *Proceedings of the 27th international conference on Human factors in computing systems* Boston, MA, USA: ACM, pp. 201–210.
- Chen, L., and Pu, P. 2009. "Interaction design guidelines on critiquing-based recommender systems," *User Modeling and User-Adapted Interaction* (19:3), pp. 167–206.
- Chen, L., and Pu, P. 2011. "Critiquing-based recommenders: survey and emerging trends," *User Modeling and User-Adapted Interaction* (22:1-2), pp. 125–150.
- Cho, H. 2010. "Determinants of Behavioral Responses to Online Privacy: The Effects of Concern, Risk Beliefs, Self-Efficacy, and Communication Sources on Self-Protection Strategies," *Journal of Information Privacy and Security* (6:1), pp. 3–27.
- Cho, H., Rivera-Sánchez, M., and Lim, S. S. 2009. "A multinational study on online privacy: global concerns and local responses," *New Media & Society* (11:3), pp. 395–416.
- Cialdini, R. B., Vincent, J. E., Lewis, S. K., Catalan, J., Wheeler, D., and Darby, B. L. 1975. "Reciprocal concessions procedure for inducing compliance: The door-in-the-face technique," *Journal of Personality and Social Psychology* (31:2), pp. 206–215.
- Cohen, L. 1979. "Approximate expressions for parameter estimates in the Rasch model," *British Journal of Mathematical and Statistical Psychology* (32:1), pp. 113–120.

- Compañó, R., and Lusoli, W. 2010. "The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas," in *Economics of Information Security and Privacy*, T. Moore, D. Pym, and C. Ioannidis (eds.), New York, NY: Springer US, pp. 169–185.
- Connolly, T., and Zeelenberg, M. 2002. "Regret in decision making," *Current directions in psychological science* (11:6), pp. 212–216.
- Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J., and Powledge, P. 2005. "Location disclosure to social relations: why, when, & what people want to share," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* Portland, OR, pp. 81–90.
- Consumer Reports. 2012. "Facebook & your privacy: Who sees the data you share on the biggest social network?," *Consumer Reports*.
- Coupey, E., Irwin, J. R., and Payne, J. W. 1998. "Product Category Familiarity and Preference Construction," *Journal of Consumer Research* (24:4), pp. 459–468.
- Cranor, L. F., Reagle, J., and Ackerman, M. S. 1999. "Beyond Concern: Understanding Net Users' Attitudes About Online Privacy," Technical Report No. TR 99.4.3, AT&T Labs - Research.
- Cranor, L., Langheinrich, M., and Marchiori, M. 2002. "A P3P Preference Exchange Language 1.0 (APPEL1.0)," W3C Working Draft, .
- Cranshaw, J., Mugan, J., and Sadeh, N. 2011. "User-Controllable Learning of Location Privacy Policies with Gaussian Mixture Models," in *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence* San Francisco, CA, pp. 1146–1152.
- Cremonesi, P., Garzotto, F., and Turrin, R. 2012. "Investigating the Persuasion Potential of Recommender Systems from a Quality Perspective: An Empirical Study," *ACM Transactions on Interactive Intelligent Systems* (2:2), pp. 11:1–11:41.
- Culnan, M. J. 1993. "'How Did They Get My Name?': An Exploratory Investigation of Consumer Attitudes toward Secondary Information Use," *MIS Quarterly* (17:3), pp. 341–363.
- Culnan, M. J., and Armstrong, P. K. 1999. "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation," *Organization Science* (10:1), pp. 104–115.
- Culnan, M. J., and Bies, R. J. 2003. "Consumer Privacy: Balancing Economic and Justice Considerations," *Journal of Social Issues* (59:2), pp. 323–342.
- Czarkowski, M., and Kay, J. 2000. "Bringing Scrutability to Adaptive Hypertext Teaching," in *Intelligent Tutoring Systems 2000*, Lecture Notes in Computer Science, G. Gauthier, C. Frasson, and K. VanLehn (eds.), (Vol. 1839) Berlin: Springer, pp. 423–433.
- Czarkowski, M., and Kay, J. 2003. "How to Give the User a Sense of Control Over the Personalization of AH?," in *AH2003: Workshop on Adaptive Hypermedia and Adaptive Web-Based Systems* Budapest, Hungary; Johnstown, PA; Nottingham, England.
- Davidsson, C., and Moritz, S. 2011. "Utilizing implicit feedback and context to recommend mobile applications from first use," in *Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation* Palo Alto, California: ACM Press, February, pp. 19–22.
- Davis, F. D. 1989. "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly* (13:3), p. 319.
- Deuker, A. 2012. "Friend-to-Friend Privacy Protection on Social Networking Sites: A Grounded Theory Study," in *AMCIS 2012 Proceedings* Seattle, WA.
- DeVellis, R. F. 2011. *Scale development: theory and applications*, Thousand Oaks, Calif.: SAGE.
- Dinev, T., Bellotto, M., Hart, P., Russo, V., Serra, I., and Colautti, C. 2006. "Privacy calculus model in e-commerce - a study of Italy and the United States," *European Journal of Information Systems* (15:4), pp. 389–402.

- Dinev, T., and Hart, P. 2006. "An Extended Privacy Calculus Model for E-Commerce Transactions," *Information Systems Research* (17:1), pp. 61–80.
- Dooms, S., De Pessemier, T., and Martens, L. 2011. "An online evaluation of explicit feedback mechanisms for recommender systems," in *7th International Conference on Web Information Systems and Technologies (WEBIST-2011)* Noordwijkerhout, The Netherlands, pp. 391–394.
- Duckham, M., and Kulik, L. 2005. "A Formal Model of Obfuscation and Negotiation for Location Privacy," in *Pervasive Computing, Lecture Notes in Computer Science*, H.-W. Gellersen, R. Want, and A. Schmidt (eds.), Springer Berlin Heidelberg, pp. 152–170.
- Egelman, S., Tsai, J., Cranor, L. F., and Acquisti, A. 2009. "Timing is everything?: the effects of timing and placement of online privacy indicators," in *Proceedings of the 27th international conference on Human factors in computing systems*, pp. 319–328.
- Elsten, T. 2012. "Supporting the selection of health improvement measures by means of a recommender system.," Eindhoven University of Technology.
- EU. 2012. "Proposal for a directive of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data by competent authorities for the purposes of prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and the free movement of such data," No. 2012/0010 (COD), .
- Fang, L., and LeFevre, K. 2010. "Privacy Wizards for Social Networking Sites," in *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, Raleigh, NC: ACM, pp. 351–360.
- Featherman, M. S., and Pavlou, P. A. 2003. "Predicting e-services adoption: a perceived risk facets perspective," *International Journal of Human-Computer Studies* (59:4), pp. 451–474.
- Fishbein, M., and Ajzen, I. 1975. *Belief, attitude, intention, and behavior: an introduction to theory and research*, Reading, MA: Addison-Wesley Pub. Co.
- Fitzsimons, G. J., and Lehmann, D. R. 2004. "Reactance to Recommendations: When Unsolicited Advice Yields Contrary Responses," *Marketing Science* (23:1), pp. 82–94.
- Fjermestad, J., and Nicholas Romano, J. 2009. "An Integrated Model for Personalization, Privacy and Security in eCommerce," *AMCIS 2009 Proceedings*.
- Fogel, J., and Nehmad, E. 2009. "Internet social network communities: Risk taking, trust, and privacy concerns," *Computers in Human Behavior* (25:1), pp. 153–160.
- Fox, S., and Lewis, O. 2001. "Fear of Online Crime: Americans support FBI interception of criminal suspects' email and new laws to protect online privacy," Pew Internet Tracking Report, Pew Internet & American Life Project.
- FTC. 2010. "Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers," Federal Trade Commission.
- Ganzach, Y. 1995. "Attribute Scatter and Decision Outcome: Judgment versus Choice," *Organizational Behavior and Human Decision Processes* (62:1), pp. 113–122.
- van de Garde-Perik, E., Markopoulos, P., de Ruyter, B., Eggen, B., and Ijsselstein, W. 2008. "Investigating Privacy Attitudes and Behavior in Relation to Personalization," *Social Science Computer Review* (26:1), pp. 20–43.
- Gardner, J. 2012. "12 Surprising A/B Test Results to Stop You Making Assumptions," *Unbounce*.
- Garrido, A., Rossi, G., and Distant, D. 2011. "Refactoring for Usability in Web Applications," *IEEE Software* (28:3), pp. 60–67.
- Gena, C., Brogi, R., Cena, F., and Vernerio, F. 2011. "The Impact of Rating Scales on User's Rating Behavior," in *User Modeling, Adaption and Personalization*, J. A. Konstan, R. Conejo, J. L. Marzo, and N. Oliver (eds.), (Vol. 6787) Berlin, Heidelberg: Springer, pp. 123–134.

- Gershon, R. C. 2005. "Computer Adaptive Testing," *Journal of Applied Measurement* (6:1), pp. 109–127.
- Girardello, A., and Michahelles, F. 2010. "AppAware: which mobile applications are hot?," in *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services* Lisbon, Portugal: ACM Press, September, pp. 431–434.
- Goncalves, J., Gomes, D., and Aguiar, R. L. 2012. "Low-latency privacy-enabled Context Distribution Architecture.," in *ICCIEEE*, pp. 1917–1922.
- Gross, R., and Acquisti, A. 2005. "Information revelation and privacy in online social networks," in *Proceedings of the 2005 ACM workshop on Privacy in the electronic society* Alexandria, VA, USA: ACM, pp. 71–80.
- Guttman, R. H., Moukas, A. G., and Maes, P. 1998. "Agent-mediated Electronic Commerce: A Survey," *Knowledge Engineering Review* (13:2), pp. 147–159.
- Hagel, J., and Rayport, J. F. 1999. "The Coming Battle for Customer Information," in *Creating value in the network economy* Boston, MA: Harvard Business School Press, pp. 159–171.
- Hann, I.-H., Hui, K.-L., Lee, S.-Y., and Png, I. 2007. "Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach," *Journal of Management Information Systems* (24:2), pp. 13–42.
- Hann, I. H., Hui, K. L., Lee, T. S., and Png, I. P. . 2002. "Online information privacy: Measuring the cost-benefit trade-off," in *23rd International Conference on Information Systems*.
- Harris. 2000. "A Survey of Consumer Privacy Attitudes and Behaviors," Privacy and American Business Newsletter, Harris Interactive, Inc.
- Harris. 2001. "Privacy Notices Research: Final Results," No. Study No. 15338, Harris Interactive, Inc.
- Harris, L., Westin, A. F., and associates. 2003a. "Consumer Privacy Attitudes: A Major Shift Since 2000 and Why," No. 10, , Privacy and American Business Newsletter, Harris Interactive, Inc.
- Harris, L., Westin, A. F., and associates. 2003b. "Most People Are 'Privacy Pragmatists' Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits," Equifax Inc.
- Hassenzahl, M. 2005. "The thing and I: understanding the relationship between user and product," in *Funology, From Usability to Enjoyment*, M. A. Blythe, K. Overbeeke, A. F. Monk, and P. C. Wright (eds.), Dordrecht, The Netherlands: Kluwer Academic Publishers, pp. 31–42.
- Häubl, G., and Trifts, V. 2000. "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science* (19:1), pp. 4–21.
- Herlocker, J. L., Konstan, J. A., and Riedl, J. 2000. "Explaining collaborative filtering recommendations," in *Proc. of the 2000 ACM conference on Computer supported cooperative work* Philadelphia, PA: ACM Press, pp. 241–250.
- Hoffman, D. L., Novak, T. P., and Peralta, M. 1999. "Building consumer trust online," *Communications of the ACM* (42:4), pp. 80–85.
- Hostler, R. E., Yoon, V. Y., and Guimaraes, T. 2005. "Assessing the impact of internet agent on end users' performance," *Decision Support Systems* (41:1), pp. 313–323.
- Ho, S. Y., and Tam, K. 2006. "Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes," *Management Information Systems Quarterly* (30:4).
- Hsu, C. 2006. "Privacy concerns, privacy practices and web site categories: Toward a situational paradigm," *Online Information Review* (30:5), pp. 569–586.
- Huber, G. P. 1974. "Multi-Attribute Utility Models: A Review of Field and Field-Like Studies," *Management Science* (20:10), pp. 1393–1402.

- Huber, J., and Puto, C. 1983. "Market Boundaries and Product Choice: Illustrating Attraction and Substitution Effects," *Journal of Consumer Research* (10:1), pp. 31–44.
- Hui, K.-L., Tan, B. C. Y., and Goh, C.-Y. 2006. "Online information disclosure: Motivators and measurements," *ACM Transactions on Internet Technology* (6:4), pp. 415–441.
- Hui, K.-L., Teo, H. H., and Lee, S.-Y. T. 2007. "The Value of Privacy Assurance: An Exploratory Field Experiment," *MIS Quarterly* (31:1), pp. 19–33.
- Hutton, R. J. B., and Klein, G. 1999. "Expert decision making," *Systems Engineering* (2:1), pp. 32–45.
- Inman, J. J., and Zeelenberg, M. 2002. "Regret in repeat purchase versus switching decisions: The attenuating role of decision justifiability," *Journal of Consumer Research* (29:1), pp. 116–128.
- Internet Society. 2012. "Global Internet User Survey 2012," Internet Society.
- Jacoby, J., and Kaplan, L. B. 1972. "The Components of Perceived Risk," in *Proceedings of the Third Annual Conference of the Association for Consumer Research*, M. Venkatesan (ed.), Chicago, IL: Association for Consumer Research, pp. 382–393.
- Jarupunphol, P., and Mitchell, C. J. 2002. "E-commerce and the Media - Influences on Security Risk Perceptions," in *Proceedings of the IFIP TC6/WG6.4 Workshop on Internet Technologies, Applications and Social Impact*, WITASI '02, Deventer, The Netherlands, The Netherlands: Kluwer, B.V., pp. 163–174.
- Jarvenpaa, S. L., Tractinsky, N., and Saarinen, L. 1999. "Consumer Trust in an Internet Store: A Cross-Cultural Validation," *Journal of Computer-Mediated Communication* (5:2).
- Jedrzejczyk, L., Price, B. A., Bandara, A. K., and Nuseibeh, B. 2010. "On the impact of real-time feedback on users' behaviour in mobile location-sharing applications," in *Proceedings of the Sixth Symposium on Usable Privacy and Security - SOUPS '10* Redmond, Washington, p. 1.
- Jensen, C., Potts, C., and Jensen, C. 2005. "Privacy Practices of Internet Users: Self-Reports versus Observed Behavior," *International Journal of Human-Computer Studies* (63:1-2), pp. 203–227.
- John, L. K., Acquisti, A., and Loewenstein, G. 2011. "Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information," *Journal of consumer research* (37:5), pp. 858–873.
- Johnson, D., and Grayson, K. 2005. "Cognitive and affective trust in service relationships," *Journal of Business research* (58:4), pp. 500–507.
- Johnson, E. J., Bellman, S., and Lohse, G. L. 2002. "Defaults, Framing and Privacy: Why Opting In ≠ Opting Out," *Marketing Letters* (13:1), pp. 5–15.
- Johnson, E. J., and Goldstein, D. 2003. "Do Defaults Save Lives?," *Science* (302:5649), pp. 1338–1339.
- Johnson, M., Egelman, S., and Bellovin, S. M. 2012. "Facebook and privacy: it's complicated," in *Proc. of the 8th Symposium on Usable Privacy and Security* Pittsburgh, PA: ACM.
- Joinson, A. N., Paine, C., Buchanan, T., and Reips, U.-D. 2008. "Measuring self-disclosure online: Blurring and non-response to sensitive items in web-based surveys," *Computers in Human Behavior* (24:5), pp. 2158–2171.
- Joinson, A. N., Reips, U.-D., Buchanan, T., and Schofield, C. B. P. 2010. "Privacy, Trust, and Self-Disclosure Online," *Human-Computer Interaction* (25:1), p. 1.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy* (98:6), pp. 1325–1348.
- Kairam, S., Brzozowski, M., Huffaker, D., and Chi, E. 2012. "Talking in circles: Selective Sharing in Google+," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* Austin, TX: ACM Press, pp. 1065–1074.
- Kay, J., and Lum, A. 2005. "Ontology-based user modelling for the Semantic Web," in *Workshop on Personalisation on the Semantic Web* Edinburgh, UK, pp. 15–23.

- Kehr, F., Kowatsch, T., Wentzel, D., and Fleisch, E. 2015. "Thinking Styles and Privacy Decisions: Need for Cognition, Faith into Intuition, and the Privacy Calculus," Osnabrück, Germany: Springer.
- Kehr, F., Wentzel, D., and Mayer, P. 2013. "Rethinking the Privacy Calculus: On the Role of Dispositional Factors and Affect," in *ICIS 2013 Proceedings* Milan, Italy.
- Keith, M. J., Babb, J. S., Lowry, P. B., Furner, C. P., and Abdullat, A. 2011. "The Roles of Privacy Assurance, Network Effects, and Information Cascades in the Adoption of and Willingness to Pay for Location-Based Services with Mobile Applications," in *2011 Dewald Roode Information Security Workshop* Blacksburg, VA, September.
- Kelley, P. G., Brewer, R., Mayer, Y., Cranor, L. F., and Sadeh, N. 2011. "An Investigation into Facebook Friend Grouping," in *INTERACT, Lecture Notes in Computer Science*, P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler (eds.), (Vol. 6948) Lisbon, Portugal: Springer Heidelberg, pp. 216–233.
- Kelley, P. G., Cesca, L., Bresee, J., and Cranor, L. F. 2010. "Standardizing privacy notices: an online study of the nutrition label approach," in *Proceedings of the 28th International Conference on Human Factors in Computing Systems, CHI 2010* Atlanta, Georgia: ACM Press, pp. 1573–1582.
- Kelley, P. G., Cranor, L. F., and Sadeh, N. 2013. "Privacy As Part of the App Decision-making Process," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, Paris, France: ACM, pp. 3393–3402.
- Kim, D. J., Ferrin, D. L., and Rao, H. R. 2008. "A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents," *Decision Support Systems* (44:2), pp. 544–564.
- Knapp, H., and Kirk, S. A. 2003. "Using pencil and paper, Internet and touch-tone phones for self-administered surveys: does methodology matter?," *Computers in Human Behavior* (19:1), pp. 117–134.
- Knight, L. 2010. "Social experiment: online privacy vs. personalization paradox," Upshot.
- Knijnenburg, B. P. 2013. "On The Dimensionality Of Information Disclosure Behavior in Social Networks," in *CSCW 2013 workshop on measuring networked privacy* San Antonio, TX.
- Knijnenburg, B. P. 2014. *Information Disclosure Profiles for Segmentation and Recommendation*, Submitted to the SOUPS2014 workshop on Privacy Personas and Segmentation.
- Knijnenburg, B. P., Bostandjiev, S., O'Donovan, J., and Kobsa, A. 2012a. "Inspectability and control in social recommenders," in *Proceedings of the sixth ACM conference on Recommender systems, RecSys '12*, New York, NY, USA: ACM, pp. 43–50.
- Knijnenburg, B. P., and Bulgurcu, B. 2015. "Form Auto-completion Tools Designed for Elaboration: Overcoming the Deleterious Effects of Decisional Heuristics on Users' Privacy," *submitted for journal publication*.
- Knijnenburg, B. P., and Jin, H. 2013. *The Persuasive Effect of Privacy Recommendations for Location Sharing Services*, unpublished manuscript.
- Knijnenburg, B. P., and Kobsa, A. 2013a. "Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems," *ACM Transactions on Interactive Intelligent Systems* (3:3), pp. 20:1–20:23.
- Knijnenburg, B. P., and Kobsa, A. 2013b. "Helping users with information disclosure decisions: potential for adaptation," in *Proceedings of the 2013 ACM international conference on Intelligent User Interfaces* Santa Monica, CA: ACM Press, March, pp. 407–416.
- Knijnenburg, B. P., and Kobsa, A. 2014. "Increasing Sharing Tendency Without Reducing Satisfaction: Finding the Best Privacy-Settings User Interface for Social Networks," in *ICIS 2014 Proceedings* Auckland, New Zealand.

- Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013a. "Preference-based location sharing: are more privacy options really better?," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* Paris, France: ACM, pp. 2667–2676.
- Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013b. "Counteracting the Negative Effect of Form Auto-completion on the Privacy Calculus," in *ICIS 2013 Proceedings* Milan, Italy.
- Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013c. "Dimensionality of information disclosure behavior," *International Journal of Human-Computer Studies* (71:12), pp. 1144–1162.
- Knijnenburg, B. P., Kobsa, A., and Jin, H. 2014a. *Segmenting the Recipients of Personal Information*, Submitted to the SOUPS2014 workshop on Privacy Personas and Segmentation.
- Knijnenburg, B. P., Kobsa, A., and Saldamli, G. 2012b. "Privacy in Mobile Personalized Systems: The Effect of Disclosure Justifications," in *Proceedings of the SOUPS 2012 Workshop on Usable Privacy & Security for Mobile Devices* Washington, DC, pp. 11:1–11:5.
- Knijnenburg, B. P., Reijmer, N. J. M., and Willemsen, M. C. 2011. "Each to his own: how different users call for different interaction methods in recommender systems," in *Proceedings of the fifth ACM conference on Recommender systems* Chicago, IL: ACM Press, pp. 141–148.
- Knijnenburg, B. P., and Willemsen, M. C. 2009. "Understanding the effect of adaptive preference elicitation methods on user satisfaction of a recommender system," in *Proceedings of the third ACM conference on Recommender systems* New York, NY, pp. 381–384.
- Knijnenburg, B. P., and Willemsen, M. C. 2010. "The effect of preference elicitation methods on the user experience of a recommender system," in *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems* Atlanta, GA, pp. 3457–3462.
- Knijnenburg, B. P., Willemsen, M. C., and Broeders, R. 2014b. "Smart Sustainability through System Satisfaction: Tailored Preference Elicitation for Energy-saving Recommenders," in *AMCIS 2014 proceedings* Savannah, GA.
- Knijnenburg, B. P., Willemsen, M. C., Gantner, Z., Soncu, H., and Newell, C. 2012c. "Explaining the user experience of recommender systems," *User Modeling and User-Adapted Interaction* (22:4-5), pp. 441–504.
- Knijnenburg, B. P., Willemsen, M. C., and Hirtbach, S. 2010. "Receiving Recommendations and Providing Feedback: The User-Experience of a Recommender System," in *E-Commerce and Web Technologies*, F. Buccafurri and G. Semeraro (eds.), (Vol. 61) Berlin, Heidelberg: Springer, pp. 207–216.
- Kobsa, A. 2001. "Tailoring Privacy to Users' Needs (Invited Keynote)," in *User Modeling 2001*, Lecture Notes in Computer Science, M. Bauer, P. J. Gmytrasiewicz, and J. Vassileva (eds.), Springer Verlag, pp. 303–313.
- Kobsa, A., Cho, H., and Knijnenburg, B. P. 2014. *An attitudinal and behavioral model of personalization at different providers*, manuscript, under review.
- Kobsa, A., and Teltzrow, M. 2005. "Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior," in *Privacy Enhancing Technologies: Revised Selected Papers of the 4th International Workshop, PET 2004, Toronto, Canada, May 26-28, 2004*, LNCS, D. Martin and A. Serjantov (eds.), (Vol. 3424) Springer Berlin Heidelberg, pp. 329–343.
- Kolter, J., and Pernul, G. 2009. "Generating User-Understandable Privacy Preferences," in *Conf. on Availability, Reliability and Security* Fukuoka, Japan: IEEE Computer Society, pp. 299–306.
- Koshimizu, T., Toriyama, T., and Babaguchi, N. 2006. "Factors on the sense of privacy in video surveillance," in *Proceedings of the 3rd ACM workshop on Continuous archival and retrieval of personal experiences*, CARPE '06, New York, NY, USA: ACM, pp. 35–44.

- Krasnova, H., Hildebrand, T., and Guenther, O. 2009. "Investigating the Value of Privacy in Online Social Networks: Conjoint Analysis," in *ICIS 2009 Proceedings* Phoenix, AZ.
- Krasnova, H., Spiekermann, S., Koroleva, K., and Hildebrand, T. 2010. "Online social networks: why we disclose," *Journal of Information Technology* (25:2), pp. 109–125.
- Krishnan, V., Narayanashetty, P. K., Nathan, M., Davies, R. T., and Konstan, J. A. 2008. "Who Predicts Better?: Results from an Online Study Comparing Humans and an Online Recommender System," in *Proceedings of the 2008 ACM Conference on Recommender Systems* Zurich, Switzerland: ACM, pp. 211–218.
- Lai, Y.-L., and 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* Claremont, CA, pp. 253–263.
- Larose, R., and Rifon, N. J. 2007. "Promoting i-Safety: Effects of Privacy Warnings and Privacy Seals on Risk Assessment and Online Privacy Behavior," *Journal of Consumer Affairs* (41:1), pp. 127–149.
- Laufer, R. S., Proshansky, H. M., and Wolfe, M. 1973. "Some Analytic Dimensions of Privacy," in *Proceedings of the Lund Conference on Architectural Psychology*, R. Küller (ed.), Lund, Sweden: Dowden, Hutchinson & Ross.
- Laufer, R. S., and Wolfe, M. 1977. "Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory," *Journal of Social Issues* (33:3), pp. 22–42.
- Lederer, S., Hong, J. I., Dey, A. K., and Landay, J. A. 2004. "Personal privacy through understanding and action: five pitfalls for designers," *Personal and Ubiquitous Computing* (8:6), pp. 440–454.
- Lederer, S., Mankoff, J., and Dey, A. K. 2003. "Who wants to know what when? privacy preference determinants in ubiquitous computing," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* Ft. Lauderdale, FL: ACM Press, pp. 724–725.
- Lee, D.-H., Im, S., and Taylor, C. R. 2008. "Voluntary self-disclosure of information on the Internet: A multimethod study of the motivations and consequences of disclosing information on blogs," *Psychology & Marketing* (25:7), pp. 692–710.
- Lee, H. J., and Park, S. J. 2007. "MONERS: A news recommender for the mobile web," *Expert Systems with Applications* (32:1), pp. 143–150.
- Lee, J., and Lee, J. 2007. "Context Awareness by Case-Based Reasoning in a Music Recommendation System," in *Proc. of the 4th International Symposium on Ubiquitous Computing Systems*, Lecture Notes in Computer Science, H. Ichikawa, W.-D. Cho, I. Satoh, and H. Youn (eds.), Tokyo, Japan: Springer Berlin Heidelberg, November 25, pp. 45–58.
- Lee, Y. E., and Benbasat, I. 2011. "The Influence of Trade-off Difficulty Caused by Preference Elicitation Methods on User Acceptance of Recommendation Agents Across Loss and Gain Conditions," *Information Systems Research* (22:4), pp. 867–884.
- Li, H., Sarathy, R., and Xu, H. 2010. "Understanding situational online information disclosure as a privacy calculus," *Journal of Computer Information Systems* (51:1), pp. 62–71.
- Li, H., Sarathy, R., and Xu, H. 2011. "The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors," *Decision Support Systems* (51:3), pp. 434–445.
- Li, N., and Chen, G. 2010. "Sharing location in online social networks," *IEEE Network* (24:5), pp. 20–25.
- Linden, G., Smith, B., and York, J. 2003. "Amazon.com recommendations: item-to-item collaborative filtering," *IEEE Internet Computing* (7:1), pp. 76–80.



- Lin, J., Benisch, M., Sadeh, N., Niu, J., Hong, J., Lu, B., and Guo, S. 2012. "A comparative study of location-sharing privacy preferences in the United States and China," *Personal and Ubiquitous Computing* (17:4), pp. 697–711.
- Lipford, H. R., Besmer, A., and Watson, J. 2008. "Understanding Privacy Settings in Facebook with an Audience View," in *Proc. of the 1st Conference on Usability, Psychology, and Security* Berkeley, CA, USA: USENIX Association.
- Liu, Y., Gummadi, K. P., Krishnamurthy, B., and Mislove, A. 2011. "Analyzing Facebook privacy settings: user expectations vs. reality," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference* Berlin, Germany: ACM, pp. 61–70.
- Li, X., and Santhanam, R. 2008. "Will it be Disclosure or Fabrication of Personal Information? An Examination of Persuasion Strategies on Prospective Employees," *International Journal of Information Security and Privacy* (2:4), pp. 91–109.
- Li, Y. 2012. "Theories in online information privacy research: A critical review and an integrated framework," *Decision Support Systems* (54:1), pp. 471–481.
- Lowry, P. B., Moody, G., Vance, A., Jensen, M., Jenkins, J., and Wells, T. 2012. "Using an elaboration likelihood approach to better understand the persuasiveness of website privacy assurance cues for online consumers," *Journal of the American Society for Information Science and Technology* (63:4), pp. 755–776.
- Lusoli, W., Bacigalupo, M., Lupiáñez-Villanueva, F., Andrade, N., Monteleone, S., and Maghiros, I. 2012. "Pan-European Survey of Practices, Attitudes and Policy Preferences as Regards Personal Identity Data Management," SSRN Scholarly Paper No. ID 2086579, Rochester, NY: Social Science Research Network.
- Mabley, K. 2000. *Privacy vs. Personalization: Part III*, Cyber Dialogue, Inc.
- Madden, M. 2012. "Privacy management on social media sites," Washington, DC: Pew Internet & American Life Project, Pew Research Center.
- Madejski, M., Johnson, M., and Bellovin, S. M. 2012. "A study of privacy settings errors in an online social network," in *2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)* Lugano, Switzerland, pp. 340–345.
- Malhotra, N. K., Kim, S. S., and Agarwal, J. 2004. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Nomological Framework," *Information Systems Research* (15:4), pp. 336–355.
- McDonald, A., Reeder, R., Kelley, P., and Cranor, L. 2009. "A Comparative Study of Online Privacy Policies and Formats," in *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, (Vol. 5672) Springer Berlin / Heidelberg, pp. 37–55–55.
- McGinty, L., and Smyth, B. 2006. "Adaptive Selection: An Analysis of Critiquing and Preference-Based Feedback in Conversational Recommender Systems," *International Journal of Electronic Commerce* (11:2), pp. 35–57.
- McGuire, W. J. 1974. "Psychological motives and communication gratification," *The uses of mass communications: Current perspectives on gratifications research* (3), pp. 167–196.
- McKenzie, C. R. M., Liersch, M. J., and Finkelstein, S. R. 2006. "Recommendations Implicit in Policy Defaults," *Psychological Science* (17:5), pp. 414–420.
- McNee, S. M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S. K., Rashid, A. M., Konstan, J. A., and Riedl, J. 2002. "On the recommending of citations for research papers," in *Proceedings of the 2002 ACM conference on Computer supported cooperative work* New Orleans, LA, pp. 116–125.
- McNee, S. M., Lam, S. K., Konstan, J. A., and Riedl, J. 2003. "Interfaces for Eliciting New User Preferences in Recommender Systems," in *User Modeling 2003*, P. Brusilovsky, A. Corbett, and F. Rosis (eds.), (Vol. 2702) Berlin: Springer Heidelberg, pp. 178–187.

- McNee, S. M., Riedl, J., and Konstan, J. A. 2006. "Being accurate is not enough: how accuracy metrics have hurt recommender systems," in *Extended abstracts on Human factors in computing systems* Montreal, Quebec, Canada, pp. 1097–1101.
- Meloy, M. G., and Russo, J. E. 2004. "Binary choice under instructions to select versus reject," *Organizational Behavior and Human Decision Processes* (93:2), pp. 114–128.
- Metzger, M. J. 2004. "Privacy, Trust, and Disclosure: Exploring Barriers to Electronic Commerce," *Journal of Computer-Mediated Communication* (9:4).
- Metzger, M. J. 2006. "Effects of Site, Vendor, and Consumer Characteristics on Web Site Trust and Disclosure," *Communication Research* (33:3), pp. 155–179.
- Metzger, M. J. 2007. "Communication Privacy Management in Electronic Commerce," *Journal of Computer-Mediated Communication* (12:2), pp. 335–361.
- Milne, G. R., and Culnan, M. J. 2004. "Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don't Read) Online Privacy Notices," *Journal of Interactive Marketing* (18:3), pp. 15–29.
- Milne, G. R., Culnan, M. J., and Greene, H. 2006. "A Longitudinal Assessment of Online Privacy Notice Readability," *Journal of Public Policy & Marketing* (25:2), pp. 238–249.
- Milne, G. R., and Gordon, M. E. 1993. "Direct Mail Privacy-Efficiency Trade-offs within an Implied Social Contract Framework," *Journal of Public Policy & Marketing* (12:2), pp. 206–215.
- Mirzadeh, N., Ricci, F., and Bansal, M. 2005. "Feature selection methods for conversational recommender systems," in *Proceedings of the 2005 IEEE International Conference on e-Technology, e-Commerce and e-Service*, pp. 772–777.
- Mohamed, N., and Ahmad, I. H. 2012. "Information privacy concerns, antecedents and privacy measure use in social networking sites: Evidence from Malaysia," *Computers in Human Behavior* (28:6), pp. 2366–2375.
- Mothersbaugh, D. L., Foxx, W. K., Beatty, S. E., and Wang, S. 2012. "Disclosure Antecedents in an Online Service Context The Role of Sensitivity of Information," *Journal of Service Research* (15:1), pp. 76–98.
- Muthén, B. 2007. "Latent variable hybrids: Overview of old and new models," in *Advances in latent variable mixture models*, G. R. Hancock and K. M. Samuelsen (eds.), Information Age Publishing, Inc.
- Nehf, J. P. 2005. "Shopping for Privacy Online: Consumer Decision-Making Strategies and the Emerging Market for Information Privacy," *University of Illinois Journal of Law, Technology & Policy* (2005), p. 1.
- Nissenbaum, H. 2004. "Privacy as Contextual Integrity," *Washington Law Review* (79), pp. 119–157.
- Nissenbaum, H. 2011. "A Contextual Approach to Privacy Online," *Daedalus* (140:4), pp. 32–48.
- Nissenbaum, H. F. 2009. *Privacy in context : technology, policy, and the integrity of social life*, Stanford, CA: Stanford Law Books.
- Norberg, P. A., Horne, D. R., and Horne, D. A. 2007. "The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors," *Journal of Consumer Affairs* (41:1), pp. 100–126.
- Nosko, A., Wood, E., and Molema, S. 2010. "All about me: Disclosure in online social networking profiles: The case of FACEBOOK," *Computers in Human Behavior* (26:3), pp. 406–418.
- Oh, J.-M., and Moon, N. 2012. "User-selectable interactive recommendation system in mobile environment," *Multimedia Tools and Applications* (57:2), pp. 295–313.
- Olivero, N., and Lunt, P. 2004. "Privacy versus willingness to disclose in e-commerce exchanges: The effect of risk awareness on the relative role of trust and control," *Journal of Economic Psychology* (25:2), pp. 243–262.

- Olson, E. L., and Widing, R. E. 2002. "Are interactive decision aids better than passive decision aids? A comparison with implications for information providers on the internet," *Journal of Interactive Marketing* (16:2), pp. 22–33.
- Olson, J. S., Grudin, J., and Horvitz, E. 2005. "A study of preferences for sharing and privacy," in *CHI '05 Extended Abstracts* Portland, OR: ACM, pp. 1985–1988.
- Oremus, W. 2014. "Facebook's Privacy Dinosaur Wants to Make Sure You're Not Oversharing," *Slate*.
- Page, X., Knijnenburg, B. P., and Kobsa, A. 2013. "What a Tangled Web We Weave: Lying in Location-Sharing Social Media," in *Proc. CSCW 2013*.
- Page, X., Kobsa, A., and Knijnenburg, B. P. 2012. "Don't Disturb My Circles! Boundary Preservation Is at the Center of Location-Sharing Concerns," in *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* Dublin, Ireland, May 20, pp. 266–273.
- Pallapa, G., Das, S. K., Di Francesco, M., and Aura, T. 2014. "Adaptive and context-aware privacy preservation exploiting user interactions in smart environments," *Pervasive and Mobile Computing* (12), pp. 232–243.
- Pan, Y., and Zinkhan, G. M. 2006. "Exploring the impact of online privacy disclosures on consumer trust," *Journal of Retailing* (82:4), pp. 331–338.
- Pariser, E. 2012. *The filter bubble: how the new personalized Web is changing what we read and how we think*, New York, N.Y.: Penguin Books.
- Pathak, B., Garfinkel, R., Gopal, R., Venkatesan, R., and Yin, F. 2010. "Empirical Analysis of the Impact of Recommender Systems on Sales," *Journal of Management Information Systems* (27:2), pp. 159–188.
- Patil, S., and Kobsa, A. 2005. "Uncovering Privacy Attitudes in Instant Messaging," in *Proceedings of the 5th ACM Conference on Supporting Group Work* Sanibel Island, FL, pp. 101–104.
- Patil, S., and Lai, J. 2005. "Who Gets to Know What when: Configuring Privacy Permissions in an Awareness Application," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* Portland, OR: ACM, pp. 101–110.
- Patil, S., Page, X., and Kobsa, A. 2011. "With a little help from my friends: can social navigation inform interpersonal privacy preferences?," in *Proceedings of the ACM 2011 conference on Computer supported cooperative work, CSCW '11*, Hangzhou, China: ACM Press, March, pp. 391–394.
- Pattaraintakorn, P., Zaverucha, G. M., and Cercone, N. 2007. "Web Based Health Recommender System Using Rough Sets, Survival Analysis and Rule-Based Expert Systems," in *Rough Sets, Fuzzy Sets, Data Mining and Granular Computing*, Lecture Notes in Computer Science, A. An, J. Stefanowski, S. Ramanna, C. J. Butz, W. Pedrycz, and G. Wang (eds.), Springer Berlin Heidelberg, pp. 491–499.
- Pavlou, P. A. 2003. "Consumer Acceptance of Electronic Commerce: Integrating Trust and Risk with the Technology Acceptance Model," *International Journal of Electronic Commerce* (7:3), pp. 101–134.
- Pavlou, P. A. 2011. "State of the Information Privacy Literature: Where Are We Now and Where Should We Go," *MIS Quarterly* (35:4), pp. 977–988.
- Pazzani, M. J., and Billsus, D. 2002. "Adaptive Web Site Agents," *Autonomous Agents and Multi-Agent Systems* (5:2), pp. 205–218.
- PCAST. 2014. "Big Data and Privacy: A Technological Perspective," Report to the President, Washington, D.C.: President's Council of Advisors on Science and Technology.
- Pedersen, P. E. 2000. "Behavioral Effects of Using Software Agents for Product and Merchant Brokering: An Experimental Study of Consumer Decision-Making," *International Journal of Electronic Commerce* (5:1), pp. 125–141.
- Petronio, S. 1991. "Communication Boundary Management: A Theoretical Model of Managing Disclosure of Private Information Between Marital Couples," *Communication Theory* (1:4), pp. 311–335.

- Petronio, S. 2002. *Boundaries of Privacy: Dialectics of Disclosure*, Albany, NY: State University of New York Press.
- Petronio, S. 2010. "Communication Privacy Management Theory: What Do We Know About Family Privacy Regulation?," *Journal of Family Theory & Review* (2:3), pp. 175–196.
- Phelps, J. E., Souza, G. D', and Nowak, G. J. 2001. "Antecedents and consequences of consumer privacy concerns; an empirical investigation," *Journal of Interactive Marketing (John Wiley & Sons)* (15:4), pp. 2–17.
- Phelps, J., Nowak, G., and Ferrell, E. 2000. "Privacy Concerns and Consumer Willingness to Provide Personal Information," *Journal of Public Policy & Marketing* (19:1), pp. 27–41.
- Pollach, I. 2007. "What's Wrong with Online Privacy Policies?," *Communications of the ACM* (50:9), pp. 103–108.
- Posner, R. A. 1981. "The Economics of Privacy," *The American Economic Review* (71:2), pp. 405–409.
- Prasad, A. 2012. "Exposing Privacy Concerns in mHealth Data Sharing," M.S. Thesis (TR2012-711), Hanover, NH: Dartmouth College.
- Preibusch, S., Krol, K., and Beresford, A. R. 2012. "The Privacy Economics of Voluntary Over-disclosure in Web Forms," in *10th Annual Workshop on the Economics of Information Security* Berlin, Germany.
- Randall, T., Terwiesch, C., and Ulrich, K. T. 2007. "User Design of Customized Products," *Marketing Science* (26:2), pp. 268–280.
- Rashid, A. M., Albert, I., Cosley, D., Lam, S. K., McNee, S. M., Konstan, J. A., and Riedl, J. 2002. "Getting to know you: learning new user preferences in recommender systems," in *Proceedings of the 7th international conference on Intelligent user interfaces*, IUI '02, San Francisco, CA: ACM, pp. 127–134.
- Ravichandran, R., Benisch, M., Kelley, P., and Sadeh, N. 2009. "Capturing Social Networking Privacy Preferences:," in *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, I. Goldberg and M. Atallah (eds.), (Vol. 5672) Springer Berlin / Heidelberg, pp. 1–18.
- Resnick, P., and Varian, H. R. 1997. "Recommender Systems," *Communications of the ACM* (40:3), pp. 56–58.
- Ricci, F. 2011. "Mobile Recommender Systems," *Information Technology & Tourism* (12:3), pp. 205–231.
- Rifon, N. J., LaRose, R., and Choi, S. M. 2005. "Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures," *Journal of Consumer Affairs* (39:2), pp. 339–360.
- Roy, A., Mackin, P., Wallenius, J., Corner, J., Keith, M., Schymik, G., and Arora, H. 2008. "An Interactive Search Method Based on User Preferences," *Decision Analysis* (5:4), pp. 203–229.
- Rubin, A. M. 2002. "The uses-and-gratifications perspective of media effects," in *Media effects: Advances in theory and research (2nd ed.)*, LEA's communication series., J. Bryant and D. Zillmann (eds.), Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers, pp. 525–548.
- Rust, R. T., Kannan, P. K., and Peng, N. 2002. "The Customer Economics of Internet Privacy," *Journal of the Academy of Marketing Science* (30:4), pp. 455–464.
- Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M., and Rao, J. 2009. "Understanding and capturing people's privacy policies in a mobile social networking application," *Personal and Ubiquitous Computing* (13:6), pp. 401–412.
- Samuelson, W., and Zeckhauser, R. 1988. "Status quo bias in decision making," *Journal of Risk and Uncertainty* (1:1), pp. 7–59.

- Schein, A. I., Popescul, A., Ungar, L. H., and Pennock, D. M. 2002. "Methods and Metrics for Cold-start Recommendations," in *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '02, Tampere, Finland: ACM, pp. 253–260.
- Schrammel, J., Köffel, C., and Tscheligi, M. 2009. "How much do you tell?: information disclosure behaviour indifferent types of online communities," in *Proceedings of the fourth international conference on Communities and technologies*, C&T '09, New York, NY, USA: ACM, pp. 275–284.
- Sezgin, E., and Ozkan, S. 2013. "A systematic literature review on Health Recommender Systems," in *E-Health and Bioengineering Conference (EHB)*, 2013, November, pp. 1–4.
- Shanteau, J. 1988. "Psychological characteristics and strategies of expert decision makers," *Acta Psychologica* (68:1–3), pp. 203–215.
- Sheehan, K. B., and Hoy, M. G. 2000. "Dimensions of Privacy Concern Among Online Consumers," *Journal of Public Policy & Marketing* (19:1), pp. 62–73.
- Sheng, H., Nah, F. F.-H., and Siau, K. 2008. "An Experimental Study on Ubiquitous commerce Adoption: Impact of Personalization and Privacy Concerns," *Journal of the Association for Information Systems* (9:6).
- Simon, H. A. 1959. "Theories of Decision-Making in Economics and Behavioral Science," *The American Economic Review* (49:3), pp. 253–283.
- Simonson, I. 1989. "Choice Based on Reasons: The Case of Attraction and Compromise Effects," *Journal of Consumer Research* (16:2), pp. 158–174.
- Singleton, S. M., and Harper, J. 2002. "With A Grain of Salt: What Consumer Privacy Surveys Don't Tell Us," SSRN Scholarly Paper No. ID 299930, Rochester, NY: Social Science Research Network.
- Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., and Cranor, L. F. 2013. "The post that wasn't: exploring self-censorship on facebook," in *Proceedings of the 2013 conference on Computer supported cooperative work* San Antonio, TX: ACM, pp. 793–802.
- Van Slyke, C., Shim, J. T., Johnson, R., and Jiang, J. J. 2006. "Concern for Information Privacy and Online Consumer Purchasing," *Journal of the Association for Information Systems* (7:1).
- Smith, H. J., Dinev, T., and Xu, H. 2011. "Information Privacy Research: An Interdisciplinary Review," *MIS Quarterly* (35:4), pp. 989–1016.
- Smith, H. J., Milberg, S. J., and Burke, S. J. 1996. "Information Privacy: Measuring Individuals' Concerns about Organizational Practices," *MIS Quarterly* (20:2), pp. 167–196.
- Smith, N. C., Goldstein, D. G., and Johnson, E. J. 2013. "Choice Without Awareness: Ethical and Policy Implications of Defaults," *Journal of Public Policy & Marketing* (32:2), pp. 159–172.
- Smyth, B. 2007. "Case-Based Recommendation," in *The Adaptive Web: Methods and Strategies of Web Personalization*, Lecture Notes in Computer Science, P. Brusilovsky, A. Kobsa, and W. Nejdl (eds.), (Vol. 4321) Berlin: Springer Verlag, pp. 342–376.
- Solove, D. J. 2013. "Privacy Self-Management and the Consent Dilemma," *Harvard Law Review* (126), pp. 1880–1903.
- De Souza, Z., and Dick, G. N. 2009. "Disclosure of information by children in social networking—Not just a case of 'you show me yours and I'll show you mine,'" *International Journal of Information Management* (29:4), pp. 255–261.
- Sparling, E. I., and Sen, S. 2011. "Rating: How Difficult is It?," in *Proceedings of the Fifth ACM Conference on Recommender Systems* Chicago, IL: ACM, pp. 149–156.
- Spiekermann, S., and Berthold, O. 2005. "Maintaining Privacy in RFID Enabled Environments," in *Privacy, Security and Trust within the Context of Pervasive Computing*, P. Robinson, H. Vogt, and W. Wagealla (eds.), (Vol. 780) Boston: Kluwer Academic Publishers, pp. 137–146.

- Spiekermann, S., Grossklags, J., and Berendt, B. 2001a. "E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior," in *Proceedings of the 3rd ACM conference on Electronic Commerce*Tampa, FL, pp. 38–47.
- Spiekermann, S., Grossklags, J., and Berendt, B. 2001b. "E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior," in *Proceedings of the 3rd ACM conference on Electronic Commerce*Tampa, FL: ACM Press, October, pp. 38–47.
- Spiekermann, S., and Paraschiv, C. 2002. "Motivating human-agent interaction: Transferring insights from behavioral marketing to interface design," *Electronic Commerce Research* (2:3), pp. 255–285.
- Srinivasan, V. 1988. "A Conjunctive-Compensatory Approach to the Self-Explication of Multiattributed Preferences," *Decision Sciences* (19:2), pp. 295–305.
- Staddon, J., Huffaker, D., Brown, L., and Sedley, A. 2012. "Are privacy concerns a turn-off?: engagement and privacy in social networks," in *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, New York, NY, USA: ACM, pp. 10:1–10:13.
- Stafford, T. F., Stafford, M. R., and Schkade, L. L. 2004. "Determining Uses and Gratifications for the Internet," *Decision Sciences* (35:2), pp. 259–288.
- Stigler, G. J. 1980. "An Introduction to Privacy in Economics and Politics," *The Journal of Legal Studies* (9:4), pp. 623–644.
- Stone, D. L. 1981. "The effects of the valence of outcomes for providing data and the perceived relevance of the data requested on privacy-related behaviors, beliefs, and attitudes," PhD Thesis, Purdue University.
- Stone, E. F., and Stone, D. L. 1990. "Privacy in Organizations: Theoretical Issues, Research Findings, and Protection Mechanisms," *Research in Personnel and Human Resources Management* (8), pp. 349–411.
- Strater, K., and Lipford, H. R. 2008. "Strategies and struggles with privacy in an online social networking community," in *Proc. of the 22nd British HCI Group Annual Conference on People and Computers*Swinton, UK: British Computer Society, pp. 111–119.
- Strater, K., and Richter, H. 2007. "Examining privacy and disclosure in a social networking community," in *Proceedings of the 3rd symposium on Usable privacy and security*Pittsburgh, Pennsylvania: ACM, pp. 157–158.
- Stutzman, F., and Kramer-Duffield, J. 2010. "Friends only: examining a privacy-enhancing behavior in facebook," in *Proceedings of the 28th international conference on Human factors in computing systems*Atlanta, Georgia, USA: ACM, pp. 1553–1562.
- Sutanto, J., Palme, E., Tan, C.-H., and Phang, C. W. 2013. "Addressing the Personalization-Privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users," *MIS Quarterly* (37:4), pp. 1141–1164.
- Tang, K., Hong, J., and Siewiorek, D. 2012. "The implications of offering more disclosure choices for social location sharing," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*Austin, TX: ACM Press, pp. 391–394.
- Tang, K., Lin, J., Hong, J., Siewiorek, D., and Sadeh, N. 2010. "Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing," in *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing*Copenhagen, Denmark, pp. 85–04.
- Taylor, D., Davis, D., and Jillapalli, R. 2009. "Privacy concern and online personalization: The moderating effects of information control and compensation," *Electronic Commerce Research* (9:3), pp. 203–223.
- Thaler, R. 1980. "Toward a positive theory of consumer choice," *Journal of Economic Behavior & Organization* (1:1), pp. 39–60.

- Thaler, R. H., and Sunstein, C. 2008. *Nudge : improving decisions about health, wealth, and happiness*, New Haven, NJ & London, U.K.: Yale University Press.
- Thambusamy, R., Church, M., Nemati, H., and Barrick, J. 2010. "Socially exchanging privacy for pleasure: Hedonic use of computer-mediated social networks," *ICIS 2010 Proceedings*.
- Tintarev, N., and Masthoff, J. 2011. "Designing and Evaluating Explanations for Recommender Systems," in *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor (eds.), Boston, MA: Springer US, pp. 479–510.
- Toch, E., Cranshaw, J., Drielsma, P. H., Tsai, J. Y., Kelley, P. G., Springfield, J., Cranor, L., Hong, J., and Sadeh, N. 2010. "Empirical models of privacy in location sharing," in *Proc. of the 12th ACM intl. conference on Ubiquitous computing* Copenhagen, Denmark: ACM Press, pp. 129–138.
- Treiblmaier, H., and Pollach, I. 2007. "Users' Perceptions of Benefits and Costs of Personalization," in *ICIS 2007 Proceedings*.
- Trewin, S. 2006. "Physical usability and the mobile web," in *Proceedings of the 2006 international cross-disciplinary workshop on Web accessibility (W4A): Building the mobile web: rediscovering accessibility* Edinburgh, U.K.: ACM, pp. 109–112.
- Tsai, J. Y., Egelman, S., Cranor, L., and Acquisti, A. 2010. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research*.
- Tsai, J. Y., Kelley, P., Drielsma, P., Cranor, L. F., Hong, J., and Sadeh, N. 2009. "Who's viewed you?: the impact of feedback in a mobile location-sharing application," in *Proceedings of the 27th international conference on Human factors in computing systems* Boston, MA, USA: ACM, pp. 2003–2012.
- Tufekci, Z. 2008. "Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites," *Bulletin of Science, Technology & Society* (28:1), pp. 20–36.
- Turner, M. A., and Varghese, R. 2002. "Making sense of the privacy debate: a comparative analysis of leading consumer privacy surveys," *Privacy & American Business*.
- Turow, J., Feldman, L., and Meltzer, K. 2005. "Open to Exploitation: America's Shoppers Online and Offline," Departmental Papers (ASC), Annenberg Public Policy Center, University of Pennsylvania.
- Tversky, A. 1972. "Elimination by aspects: A theory of choice," *Psychological Review* (79:4), pp. 281–299.
- Venkatesh, V., Morris, M. G., Gordon B. Davis, and Davis, F. D. 2003. "User Acceptance of Information Technology: Toward a Unified View," *MIS Quarterly* (27:3), pp. 425–478.
- Vijayasarathy, L. R., and Jones, J. M. 2001. "Do Internet Shopping Aids Make a Difference? An Empirical Investigation," *Electronic Markets* (11:1), pp. 75–83.
- Vozalis, M. G., and Margaritis, K. G. 2007. "Using SVD and demographic data for the enhancement of generalized Collaborative Filtering," *Information Sciences* (177:15), pp. 3017–3037.
- Wang, W., and Benbasat, I. 2007. "Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs," *Journal of Management Information Systems* (23:4), pp. 217–246.
- Wang, Y., and Kobsa, A. 2007. "Respecting Users' Individual Privacy Constraints in Web Personalization," in *User Modeling 2007*, Lecture Notes in Computer Science, C. Conati, K. McCoy, and G. Paliouras (eds.), Corfu, Greece: Springer Berlin / Heidelberg, pp. 157–166.
- Wang, Y., Leon, P. G., Acquisti, A., Cranor, L. F., Forget, A., and Sadeh, N. 2014. "A Field Trial of Privacy Nudges for Facebook," in *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* Toronto, Canada: ACM, pp. 2367–2376.
- Wang, Y., Leon, P. G., Scott, K., Chen, X., Acquisti, A., and Cranor, L. F. 2013. "Privacy Nudges for Social Media: An Exploratory Facebook Study," in *Second International Workshop on Privacy and Security in Online Social Media* Rio De Janeiro, Brazil.

- Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P. G., and Cranor, L. F. 2011a. "I regretted the minute I pressed share': a qualitative study of regrets on Facebook," in *Proceedings of the Seventh Symposium on Usable Privacy and Security* Pittsburgh, PA: ACM, pp. 10:1–10:16.
- Wang, Y., Norice, G., and Cranor, L. 2011b. "Who Is Concerned about What? A Study of American, Chinese and Indian Users' Privacy Concerns on Social Network Sites," in *TRUST*, Lecture Notes in Computer Science, J. McCune, B. Balacheff, A. Perrig, A.-R. Sadeghi, A. Sasse, and Y. Beres (eds.), (Vol. 6740) Pittsburgh, PA: Springer, pp. 146–153.
- Watson, J., Besmer, A., and Lipford, H. R. 2012. "+Your circles: sharing behavior on Google+," in *Proc. of the 8th Symposium on Usable Privacy and Security* Pittsburgh, PA: ACM.
- Wedell, D. H. 1997. "Another look at reasons for choosing and rejecting," *Memory & Cognition* (25:6), pp. 873–887.
- Weinstein, N. D. 1989. "Optimistic biases about personal risks," *Science* (246:4935), pp. 1232–1233.
- Wenning, R., and Schunter, M. 2006. "The Platform for Privacy Preferences 1.1 (P3P1.1) Specification," W3C Working Group Note.
- Westin, A. F., Harris, L., and associates. 1981. *The Dimensions of privacy : a national opinion research survey of attitudes toward privacy*, New York: Garland Publishing.
- Westin, A. F., and Maurici, D. 1998. "E-Commerce & Privacy: What the Net Users Want," Privacy & American Business, and PricewaterhouseCoopers LLP.
- White House. 2012. "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Economy," Washington, D.C.: White House.
- White, T. B. 2004. "Consumer Disclosure and Disclosure Avoidance: A Motivational Framework," *Journal of Consumer Psychology* (14:1&2), pp. 41–51.
- Willemsen, M. C., Knijnenburg, B. P., Graus, M. P., Velter-Bremmers, L. C. M., and Fu, K. 2011. "Using Latent Features Diversification to Reduce Choice Difficulty in Recommendation Lists," in *RecSys'11 Workshop on Human Decision Making in Recommender Systems, CEUR-WS, vol. 811* Chicago, IL, pp. 14–20.
- Wilson, D., and Valacich, J. 2012. "Unpacking the Privacy Paradox: Irrational Decision-Making within the Privacy Calculus," in *Proceedings of the International Conference on Information Systems* Orlando, FL, December 14.
- Wilson, S., Cranshaw, J., Sadeh, N., Acquisti, A., Cranor, L. F., Springfield, J., Jeong, S. Y., and Balasubramanian, A. 2013. "Privacy manipulation and acclimation in a location sharing application," in *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing, UbiComp '13*, Zürich, Switzerland: ACM, pp. 549–558.
- Wisniewski, P., Knijnenburg, B. P., and Richter Lipford, H. 2014. "Profiling Facebook Users' Privacy Behaviors," in *SOUPS2014 Workshop on Privacy Personas and Segmentation* Menlo Park, CA.
- Wroblewski, L. 2008. *Web Form Design: Filling in the Blanks*, Brooklyn, NY: Rosenfeld Media.
- Wu, H., Knijnenburg, B. P., and Kobsa, A. 2014. "Improving the prediction of users' disclosure behavior... by making them disclose more predictably?," in *Symposium on Usable Privacy and Security (SOUPS)*.
- Xiao, B., and Benbasat, I. 2007. "E-commerce Product Recommendation Agents: Use, Characteristics, and Impact," *Mis Quarterly* (31:1), pp. 137–209.
- Xu, H. 2007. "The effects of self-construal and perceived control on privacy concerns," in *ICIS 2007 Proceedings*, p. paper 125.
- Xu, H., Dinev, T., Smith, H. J., and Hart, P. 2008. "Examining the formation of individual's privacy concerns: Toward an integrative view," in *ICIS 2008 Proceedings* Paris, France.



- Xu, H., Luo, X. (Robert), Carroll, J. M., and Rosson, M. B. 2011. "The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing," *Decision Support Systems* (51:1), pp. 42–52.
- Xu, H., Teo, H.-H., and Tan, B. C. Y. 2005. "Predicting the Adoption of Location-Based Services: The Role of Trust and Perceived Privacy Risk," in *Proceedings of the International Conference on Information Systems* Las Vegas, NV, December 31, pp. 861–874.
- Xu, H., Teo, H.-H., Tan, B. C. Y., and Agarwal, R. 2009. "The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services," *Journal of Management Information Systems* (26:3), pp. 135–174.
- Xu, H., Wang, N., and Grossklags, J. 2012. "Privacy-by-ReDesign: Alleviating Privacy Concerns for Third-Party Applications," in *ICIS 2012 Proceedings* Orlando, FL.
- Yao, M. Z., Rice, R. E., and Wallis, K. 2007. "Predicting user concerns about online privacy," *Journal of the American Society for Information Science and Technology* (58:5), pp. 710–722.
- Young, A. L., and Quan-Haase, A. 2009. "Information Revelation and Internet Privacy Concerns on Social Network Sites: A Case Study of Facebook," in *Proceedings of the Fourth International Conference on Communities and Technologies* University Park, PA: ACM, pp. 265–274.
- Youn, S. 2009. "Determinants of Online Privacy Concern and Its Influence on Privacy Protection Behaviors Among Young Adolescents," *Journal of Consumer Affairs* (43:3), pp. 389–418.
- Zhao, L., Lu, Y., and Gupta, S. 2012. "Disclosure Intention of Location-Related Information in Location-Based Social Network Services," *International Journal of Electronic Commerce* (16:4), pp. 53–90.
- Zheng, V. W., Zheng, Y., Xie, X., and Yang, Q. 2012. "Towards mobile intelligence: Learning from GPS history data for collaborative recommendation," *Artificial Intelligence* (184–185), pp. 17–37.
- Zhou, T. 2012. "Examining Location-based Services Usage from the Perspectives of Unified Theory of Acceptance and Use of Technology and Privacy Risk," *Journal of Electronic Commerce Research* (13:2), pp. 135–144.
- Zickuhr, K. 2012. "Three-quarters of smartphone owners use location-based services," *Pew Research Center*.