The Effect of Personalization Provider Characteristics on Privacy Attitudes and Behaviors: An Elaboration Likelihood Model Approach

Alfred Kobsa
Department of Informatics, University of California, Irvine, CA 92697-3440. E-mail: Kobsa@uci.edu

Hichang Cho
Department of Communications and New Media, National University of Singapore, 10 Kent Ridge Crescent, Singapore 119260. E-mail: hichang_cho@nus.edu.sg

Bart P. Knijnenburg
School of Computing, Clemson University, 100 McAdams Hall, Clemson, SC 29634. E-mail: bartk@clemson.edu

Many computer users today value personalization but perceive it in conflict with their desire for privacy. They therefore tend not to disclose data that would be useful for personalization. We investigate how characteristics of the personalization provider influence users' attitudes towards personalization and their resulting disclosure behavior. We propose an integrative model that links these characteristics via privacy attitudes to actual disclosure behavior. Using the Elaboration Likelihood Model, we discuss in what way the influence of the manipulated provider characteristics is different for users engaging in different levels of elaboration (represented by the user characteristics of privacy concerns and self-efficacy). We find particularly that (a) reputation management is effective when users predominantly use the peripheral route (i.e., a low level of elaboration), but much less so when they predominantly use the central route (i.e., a high level of elaboration); (b) client-side personalization has a positive impact when users use either route; and (c) personalization in the cloud does not work well in either route. Managers and designers can use our results to instill more favorable privacy attitudes and increase disclosure, using different techniques that depend on each user's levels of privacy concerns and privacy self-efficacy.

Introduction

Since the mid-1990s, personalization has experienced tremendous growth in e-business. Consumers today expect personalized interaction and a personalized relationship with online retailers (Ernst & Young, 2012; Humfries, 2012). A personalized shopping experience to meet their needs and preferences has been the top priority in year-on-year consumer surveys (Schaefer, 2011). 35% of Amazon purchases and 75% of Netflix views are currently a result of personalized recommendations (MacKenzie, Meyer, & Noble, 2013).

A necessary prerequisite of personalization is that personalized systems collect a significant amount of personal data (Kobsa, 1990, 2007; Riedl, 2001; Toch, Wang, & Cranor, 2012). This is often in conflict with today's Internet users' high levels of privacy concerns and their strong proclivity towards privacy-preserving behavior (Cho, Rivera-Sánchez, & Lim, 2009; TRUSTe, 2014). To reconcile personalization with privacy, various privacy-enhancing tools and methods have been proposed (for surveys see Kobsa, 2007; Toch et al., 2012; Wang & Kobsa, 2008).

Although from a technical and conceptual point of view privacy-enhancing technologies do protect personal data, a theoretical model is still lacking that links this technology via attitudes to behaviors and that is validated in a realistic setting. This article investigates the effect of different characteristics of a personalization provider on users' perceptions, how dispositional differences moderate these effects, and how these perceptions in turn influence their behavioral reactions. Specifically, this article makes the following contributions:

1. It investigates three techniques for the provider to *present* itself to the user (reputation management, client-side

personalization, cloud-based personalization) to obtain favorable privacy-related attitudes.

2. It develops an *integrated* causal model that explains how these privacy-related attitudes influence information disclosure in the context of personalization.

3. Using the Elaboration Likelihood Model (ELM), it identifies how the relative success of the presentation techniques depends on two personal variables (namely privacy self-efficacy beliefs and general online privacy concerns).

4. It validates the research model in an experiment ($n = 390$), focusing on actual information disclosure behavior as its outcome.

From a managerial perspective, the results of our study allow us to answer the question: What can a personalization provider do to instill more favorable privacy-related attitudes, and thereby increase users' disclosure? Should it try to improve its reputation, offer privacy-preserving client-side personalization, or should it rather portray personalization as being carried out in the cloud? Our study shows that the answer to this question depends on users' general privacy concerns and privacy self-efficacy: Reputation managements seems to work best for users with lower concerns and self-efficacy, whereas client-side personalization seems to work best for users with higher concerns and self-efficacy.

Related Work

This study builds on existing research in privacy, personalization, and the ELM. At the same time, it addresses a number of theoretical, methodological, and practical gaps in prior privacy research. In the next section, we introduce related work and identify gaps that we subsequently cover in our study.

Privacy and Personalization

Personalization necessitates that at least rudimentary data about each user are available to the personalization system. In general, the quality of personalization improves with the amount of personal data available. It was recognized early on that users may, however, not agree with this data-collection (Kobsa, 1990; Riedl, 2001). Personalization is thus in conflict with users' potential privacy concerns, and there are limits to what users are willing to disclose to a personalization provider, something Awad and Krishnan (2006) call the "personalization-privacy paradox." This paradox has since been the topic of extensive research (Awad & Krishnan, 2006; Chellappa & Sin, 2005; FTC, 2012; Sheng, Nah, & Siau, 2008; Sutanto, Palme, Tan, & Phang, 2013). Particularly, researchers in information systems (IS) and human-computer interaction (HCI) have investigated key antecedents of information disclosure in personalized systems. These are:

- The perceived value of personalization (Brodie, Karat, & Karat, 2004; Chellappa & Sin, 2005; Ho & Kwok, 2002; Li & Unger, 2012),

- Users' trust in (i.e., the reputation of) the personalization provider (Andrade, Kaltcheva, & Weitz, 2002; Briggs, Simpson, & Angeli, 2004; Komiak & Benbasat, 2006; Li, 2014), and

- Antecedents of trust, such as control (Sundar & Marathe, 2010; Taylor, Davis, & Jillapalli, 2009) and transparency (Awad & Krishnan, 2006; Kobsa & Teltzrow, 2005).

Research in information technology, in contrast, focuses on ways to protect users' privacy with technical means, such as:

- Preserving users' anonymity while maintaining the same quality of personalization (Ishitani, Almeida, & Wagner, 2003; Kobsa & Schreck, 2003),

- Changing some personal data randomly or in a user-controlled manner to allow plausible deniability (Berkovsky, Eytani, Kuflik, & Ricci, 2007; Chow, Jin, Knijnenburg, & Saldamli, 2013), and

- Keeping all personal data on the users' local device and performing all personalization on this device rather than sending personal data to a remote site where personalization is performed ("client-side personalization," Cassel & Wolz, 2001; Juels, 2001; Mulligan & Schwartz, 2000; Newman & Enscoe, 2000).

Empirical studies on users' privacy-related attitudes and behaviors with regard to these technical means for privacy protection are still largely missing. For example, from a technical point of view, by performing the personalization locally and not sending any personal data to the provider, client-side personalization enhances the privacy of the user (Solove, 2006). However, it is unclear whether users will indeed *perceive* client-side personalization to be privacy-friendly, that is, whether users' attitudes and disclosure behaviors are different towards a provider that uses client-side personalization versus a provider that uses traditional remote personalization practices.

Another very recent form of personalization is cloud-based personalization, or "cloud personalization" for short (Guo, Chen, Wu, & Wang, 2009; Hsueh et al., 2010; Jalali, Bouyer, Arasteh, & Moloudi, 2013; López-Nores, Blanco-Fernández, & Pazos-Arias, 2013). In contrast to client-side personalization, cloud personalization may *increase* rather than decrease privacy concerns. For example, many respondents in Ion, Sachdeva, Kumaraguru, and Capkun (2011) indicate they do not entrust the cloud with sensitive data. Moreover, the notion of "personalization in the cloud" detaches personalization from a specific provider and from any associations users may have with that provider, (i.e., its reputation).

Our research aims to close the research gap between the technical implications of these personalization techniques and the way users perceive them. We treat these technical aspects as cues for the formation of privacy attitudes, and compare the effect of these cues against reputation-related cues traditionally found in IS and HCI research: Is client-side personalization more effective than reputation

management? Will a less reputable provider benefit by highlighting that its personalization is carried out in the cloud? We are the first to answer these questions by comparing the influence of reputation management and personalization techniques on users' privacy attitudes and information disclosure.

Elaboration Likelihood Model for Privacy Decision Making

Before we can argue about the effect of different presentation techniques on users' information disclosure decisions, we first need to theoretically unpack the means by which users reconcile the "personalization-privacy paradox" and decide whether or not to disclose a certain piece of information. Two competing views on *privacy decision making* exist: the "privacy calculus view" and the "heuristic shortcuts view."

Many researchers argue that people employ a privacy calculus when making privacy-related decisions (Culnan & Armstrong, 1999; Culnan & Bies, 2003; Dinev & Hart, 2006; Hann, Hui, Lee, & Png, 2007; Li, Sarathy, & Xu, 2010; Milne & Gordon, 1993; Min & Kim, 2014; Petronio, 2002; Wilson & Valacich, 2012; Xu, Luo, Carroll, & Rosson, 2011; Xu, Teo, Tan, & Agarwal, 2009). Central to the idea of privacy calculus is that when people have to decide whether or not to perform a privacy-relevant action (such as disclosing certain personal data), they weigh the anticipated benefits of this action against its perceived privacy risks. The survey articles by Pavlou (2011), Smith, Dinev, and Xu (2011) and Li (2012) demonstrate that the notion of privacy calculus has become a well-established concept in privacy research.

On the other hand though, many recent experiments have shown that people's privacy decision making often cannot be explained well by assuming they trade off perceived benefits and risks in an entirely rational manner. Rather, their disclosure intentions and actual disclosure of personal information is influenced by various heuristics, such as:

- Information on others' willingness to disclose this information (i.e., "social proof heuristic," cf. Acquisti, John, & Loewenstein, 2012),
- The order of sensitivity in which items are being asked (i.e., "foot in the door" and "door in the face" technique, cf. Acquisti et al., 2012),
- The overall professionalism of the user interface design (i.e., "affect heuristic," cf. John, Acquisti, & Loewenstein, 2011),
- the available options to choose from (i.e., "decision context noninvariance," cf. Knijnenburg, Kobsa, & Jin, 2013b), and
- What the default is and how one asks (i.e., "default" and "framing" effects, cf. Knijnenburg & Kobsa, 2014; Lai & Hui, 2006).

Some privacy researchers therefore hypothesize that people predominantly use shortcuts for making privacy decisions (Acquisti & Grossklags, 2008; Acquisti, Adjerid, &

Brandimarte, 2013; Adjerid, Acquisti, Brandimarte, & Loewenstein, 2013; Cho, Lee, & Chung, 2010; LaRose & Rifon, 2006; Lowry et al., 2012).

How can we reconcile the "rational" privacy calculus view on privacy decision making with the alternative heuristic view? The Elaboration Likelihood Model (ELM) is a "dual process theory" of attitude formation and decision making that integrates decision-making processes with different degrees of elaboration (Petty & Cacioppo, 1986; Petty & Wegener, 1999). According to the ELM, people use two processing routes to a varying extent: a central route (high elaboration) and a peripheral route (low elaboration). When predominantly taking the central route, people engage in a more effortful elaboration process (Zhang, 1996) and form their attitudes about a product based on a more careful assessment of the most relevant available information, such as: argument quality (Cacioppo, Petty, Kao, & Rodriguez, 1986), messages that highlight the superiority of the product (Petty, Cacioppo, & Schumann, 1983), or distinctive features of the product (Lord, Lee, & Sauer, 1995). This central route is much in line with the privacy calculus. When predominantly taking the peripheral route, people instead perform a more heuristic evaluation, which often relies on superficial but easily accessible cues, such as their mood or general feelings, consensus heuristics (Slovic, Finucane, Peters, & MacGregor, 2004), the credibility and attractiveness of the message source (Petty & Cacioppo, 1986), or famous endorsers (Petty et al., 1983). Online, typical peripheral cues are website reputation (Shamdasani, Stanaland, & Tan, 2001), and design quality (Bansal, Zahedi, & Gefen, 2008), which is in line with the heuristic accounts of privacy decision making (Andrade et al., 2002; John et al., 2011; Li, 2014).

How do we know under what circumstances participants use the central route, and under what circumstances the peripheral route? ELM research rarely endeavors to measure the amount of elaboration directly, but instead specifies two important variables that determine the extent to which someone uses the central or peripheral route: motivation and ability. Motivation for engaging in elaboration is in turn affected by personal or dispositional characteristics (e.g., need for cognition) or situational characteristics (e.g., personal relevance, involvement). Similarly, the ability to process presented information can be affected by personal characteristics (e.g., prior knowledge, expertise in subject matter) or situational factors (e.g., sufficient time, or lack of distraction) (Cacioppo et al., 1986; Petty et al., 1983). In privacy research, the idea that elaboration depends on motivation and ability is corroborated by privacy scholars who have argued that people use shortcuts and heuristics because they are incapable (Liu, Gummadi, Krishnamurthy, & Mislove, 2011; Madejski, Johnson, & Bellovin, 2012) or not motivated (Compañó & Lusoli, 2010) to perform an elaborate privacy calculus.

Several existing studies have applied the ELM to examine the formation of privacy-related attitudes and behavioral intentions (Angst & Agarwal, 2009; Lowry et al.,

2012; Yang, Hung, Sung, & Farn, 2006; Zhou, 2012). Their findings suggest that the degree of elaboration indeed influences the relative impact of different types of privacy-related cues. For instance, people who have higher levels of privacy concern (i.e., higher motivation) and high privacy self-efficacy or low trait-anxiety (i.e., higher ability) are likely to engage in deep information processing. Consequently, their privacy-related attitudes or intentions are affected by central rather than peripheral cues. More specifically, Yang et al. (2006) examine the effect of objective information (a central cue) and third-party seals (a peripheral cue) on privacy assurance perceptions and, ultimately, trust. They find that the peripheral route is more prominent among people with low involvement or high trait-anxiety, whereas the central route is more prominent among people with high involvement and low trait-anxiety. When it comes to trust formation about mobile banking (Zhou, 2012) and websites (Bansal et al., 2008), central cues such as content-based arguments and information quality are more important for people with high levels of privacy concern and privacy self-efficacy, whereas superficial information such as structural assurances, design, company information and reputation are more important for people with low levels of privacy concern or privacy self-efficacy. In Angst and Agarwal's (2009) study about users' attitudes towards electronic health records, the strength of arguments are more likely to affect people with high privacy concerns and high involvement than those with low concerns and involvement.

The techniques to improve privacy-related attitudes that our study investigates range from ostensive yet superficial (e.g., reputation management) to technical (e.g., client-side personalization). Based on the extensive literature described earlier, we will use the ELM to study how users' level of elaboration influences their evaluation of these techniques. In line with existing work, we argue that people with a high level of general privacy concerns (cf. motivation) and self-efficacy beliefs (cf. ability) are more likely to use the central route, whereas people with low levels of general privacy concerns and self-efficacy beliefs use the peripheral route. Consequently, we argue that users who predominantly use the peripheral route are likely to be convinced by ostensive yet superficial techniques, whereas users who predominantly use the central route are not likely convinced by superficial cues but rather by technical solutions.

Shortcomings of Existing Privacy Research

There already exists a large body of IS research on privacy, and several studies have already looked into privacy of personalization (Awad & Krishnan, 2006; Chellappa & Sin, 2005; Sheng et al., 2008; Sutanto et al., 2013) or ELM in the context of privacy (Angst & Agarwal, 2009; Lowry et al., 2012; Yang et al., 2006; Zhou, 2012). We are arguably the first to combine these two topics in an empirical study. Moreover, our study setup allows us to address two additional shortcomings of existing privacy research, including

privacy research using the ELM: the lack of *integration* and a certain lack of *realism*.

Lack of integration. As Smith et al. (2011) point out, most existing privacy studies cover only narrow subsections of the field of privacy, and there is a lack of integration. Our research is to our best knowledge one of the first works (Knijnenburg & Kobsa, 2013 are a notable exception) on personalization and privacy to develop a causal model that *integrates*:

- Provider-controlled variables or features, such as its method of presentation to the user;
- Personal traits, such as general privacy concerns and privacy self-efficacy beliefs;
- Privacy-related attitudes, such as system-specific privacy concerns and perceived security;
- Outcome expectations, such as self-anticipated satisfaction; and
- Privacy-related behaviors, such as information disclosure.

Such a causal model will serve as a useful conceptual tool for specifying theoretical relationships among key determinants of user attitudes and behavior in the context of personalization.

Lack of realism. From a methodological perspective, our study is one of a few privacy studies that tests an integrative causal model specifying relationships between personal traits, attitudes, and *actual* behavior in a realistic experiment. Previous research largely tested the role of privacy-related attitudes in generic surveys, or conducted experiments in hypothetical situations, focusing on behavioral *intention* rather than actual behavior. Numerous studies have demonstrated that people may have low intentions to disclose their personal information because of privacy concerns, but in reality often end up sharing their personal information anyway (Berendt, Günther, & Spiekermann, 2005; Norberg, Horne, & Horne, 2007; Spiekermann, Grossklags, & Berendt, 2001). To the extent that this so-called "privacy paradox"[1] holds, associations between privacy concerns, attitudes, and behavioral intention may not be reflective of actual behavior (Smith et al., 2011). In fact, although it is well established that privacy concerns and attitudes influence intention, less evidence exists regarding the effect of privacy concerns and attitudes on actual behavior (Norberg et al., 2007). Though several studies have measured privacy-related attitudes and actual behaviors in realistic settings, they seldom integrate both into a comprehensive model and framework. Our study is a behavioral experiment in which 390 participants downloaded and installed a prototype of a smartphone-based personalization app and *actually* disclosed personal data to the system.

---

[1]The privacy paradox should not be confused with the personalization-privacy paradox discussed previously.
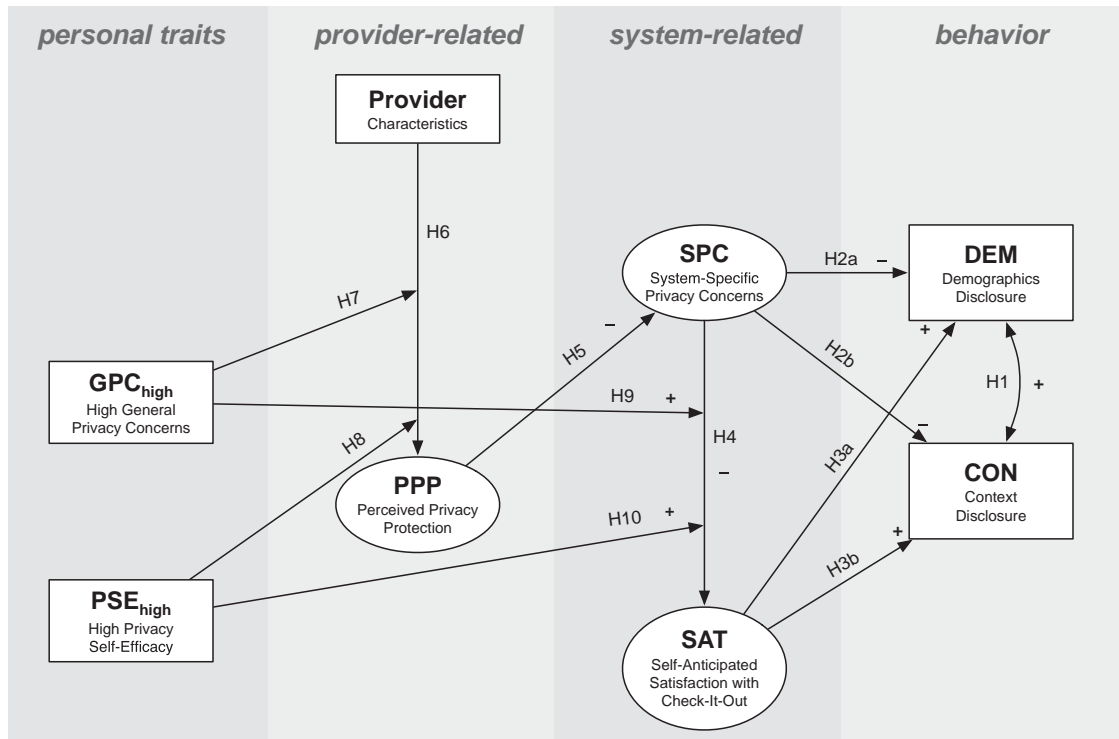
FIG. 1. Hypothesized effects.

Within this realistic (rather than hypothetical) user context, we tested the effects of personalization provider, personal traits, and privacy-related attitudes on actual disclosure behavior. This is arguably more valuable from a managerial perspective than studying the effects on behavioral intentions.

Research Model and Hypothesis Development

Our study investigates the perceptions participants have of the different characteristics of the personalization provider presented to them, and how these perceptions in turn influence their behavioral reactions. The results of our study allow us to answer the questions: (a) what can a personalization provider do to instill more favorable privacy-related attitudes, and to increase users' disclosure? and (b) under what conditions do different provider strategies have more or less desired effects on users' attitudes and behaviors? We consider an unknown personalization provider as our baseline condition, and test the effects of three different arrangements that arguably influence disclosure. Specifically, we manipulate the characteristics of the personalization provider by telling participants that the recommendations they receive are provided by:

- American Personalization (baseline, a fictitious company with neutral reputation)
- Amazon (a well-known and highly reputable company, to test the effect of reputation)

- The cloud (a nameless entity, to test the effect of removing focus from a particular provider)[2]
- Client-side personalization (a technique that divulges no data to the provider, to test the effect of an actual privacy-preserving technology)

Drawing on the ELM and previous research on privacy and personalization, we developed a conceptual model (Figure 1) that specifies how this Provider manipulation ("Provider") influences users' attitudes and behaviors, and how these influences differ between users who predominantly use the central route and those who predominantly use the peripheral route.

Information disclosure was chosen as a main dependent variable, as it is the key behavioral outcome in the context of personalization and privacy (Hui, Teo, & Lee, 2007; Knijnenburg & Kobsa, 2013). As with most research on privacy and personalization (Chellappa & Sin, 2005; Kobsa & Teltzrow, 2005; Taylor et al., 2009), we predicted that disclosure is determined directly by two salient factors, namely System-specific Privacy Concerns (SPC) and anticipated benefits (operationalized here as Satisfaction [SAT]).

---

[2]In the interview study on cloud services by Uusitalo, Karppinen, Juhola, and Savola (2010, p. 71) it was found that "Brand, Reputation, Image, History, and Name were seen as the most important aspects [for judging the cloud service], raised by half of the interviewees." We withhold this information in our provider-less cloud condition to study the effects on the privacy decision-making process.

Together with Perceived Privacy Protection (PPP), these variables mediate the effects of Provider on Disclosure.

Consistent with the ELM, we argue that Provider-related reputation management (i.e., the Amazon condition) would make ostensive yet superficial improvements over the baseline condition (i.e., American Personalization), susceptible to peripheral-route attitude formation only, whereas privacy-preserving technology (i.e., client-side personalization) would *additionally* make an *actual* improvement, susceptible to both peripheral-route and central-route processing. Regarding personalization in the Cloud, we are less optimistic: Although referring to the personalization as happening in the cloud removes focus from the provider itself, studies have shown that users of cloud-based personalization do not entrust the cloud with sensitive data (Ion et al., 2011) and take protective actions when given this opportunity (Clark, Snyder, McCoy, & Kanich, 2015). In line with the ELM, we argue that "the cloud" is still an unfamiliar concept to most users (thereby reducing its peripheral-route effectiveness), and that deferring to the cloud still leaves the question of which entity actually manages the personal information unanswered (thereby reducing its central-route effectiveness).

Finally, based on existing privacy research that uses ELM, we argue that two privacy-related personal traits (i.e., general privacy concerns [GPC] and privacy self-efficacy [PSE]) determine users' elaboration likelihood. In line with this, we specified that GPC and PSE moderate the relative effect of different Provider characteristics on users' privacy attitudes and the relative importance of different attitudes in influencing users' disclosure behavior.

Figure 1 depicts these causal relationships between variables and hypotheses. We will discuss them in more detail here.

Demographics Disclosure (DEM) and Context Disclosure (CON)

As our study concerns a realistic experiment, users' actual information disclosure is the main dependent variable. Though current privacy literature commonly treats information disclosure as a one-dimensional behavior, recent studies have revealed that it is rather a multidimensional construct, in the sense that different *types* of information can be distinguished, which are disclosed to a different extent by different people (Knijnenburg, Kobsa, & Jin, 2013a). As such, our study considers two previously distinguished factors of information disclosure (Knijnenburg & Kobsa, 2013): demographics (DEM; self-reported personal information) and context (CON; system-tracked smartphone usage data). We postulate that the determinants of information disclosure behavior may influence each factor to a different extent, and we therefore formulate separate hypotheses for each factor. However, we predict that the two factors correlate with each other because they refer to the same higher-level construct, information disclosure. Hence, we postulate the following hypothesis:

H1: *Demographics disclosure (DEM) and context disclosure (CON) are two separate but correlated factors.*

System-Specific Privacy Concerns (SPC) and Satisfaction (SAT)

Numerous privacy studies have revealed that *privacy concern* and *satisfaction with the system* are central determinants for information disclosure in personalization. People are likely to withhold information when they are concerned about their privacy, but are at the same time likely to disclose personal data and forgo privacy in return for anticipated benefits from personalization (Chellappa & Sin, 2005; Kobsa & Teltzrow, 2005; Taylor et al., 2009; Xu et al., 2011). As such, we identified these two factors (i.e., System-specific Privacy Concerns [SPC] and Satisfaction [SAT]) as the main determinants of information disclosure behavior, and hypothesize them to have a direct impact on disclosure behavior.

H2: *System-specific Privacy Concerns (SPC) have a negative impact on demographics disclosure (DEM; H2a) and context disclosure (CON; H2b).*

H3: *Satisfaction (SAT) has a positive impact on demographics disclosure (DEM; H3a) and context disclosure (CON; H3b).*

We also posit that SPC has a negative impact on SAT. Privacy concerns are salient elements affecting the user experience of personalization systems (Teltzrow & Kobsa, 2004). Specifically, people may have certain expectations about privacy, and they may use these expectations as key evaluation standards in determining their satisfaction with online services (Chen, Huang, & Muzzerall, 2012; Fox et al., 2000). Hence, people who feel that personal data collection is intrusive or uncomfortable (i.e., people with high SPC) are less likely to be satisfied with a personalized service (Lukaszewski, Stone, & Stone-Romero, 2008). Therefore, we advance the following hypothesis:

H4: *System-specific Privacy Concerns (SPC) have a negative impact on Satisfaction (SAT).*

Perceived Privacy Protection (PPP) and Provider Characteristics

We define Perceived Privacy Protection (PPP) as the degree to which people believe that the personalization provider protects their personal information. This construct is essential in measuring participants' evaluation of the Provider-related conditions we introduce. The construct is related to perceived security (Chellappa, 2008; Shin, 2010), privacy protection belief (Li et al., 2010), perceived (website) privacy protection (Metzger, 2004), and trust in an organization/provider (Bhattacherjee, 2002; Zimmer, Arsal, Al-Marzouq, & Grover, 2010), and it is inverse to perceived risk (Norberg et al., 2007). Given that privacy is considered a problem of uncertainty or risk (Acquisti & Grossklags, 2008), perceived protection is likely to play a key role in

determining user attitudes and behavior. Conceptually, PPP and System-specific Privacy Concerns (SPC) should be interrelated to the extent that perceived protection of personal information influences how concerned an individual is about the information collection: The safer an individual believes the information is, the less s/he minds it being collected. Empirically, previous research on privacy has demonstrated the interrelatedness among those factors (Chen et al., 2012; Dinev & Hart, 2004, 2006). Hence, we predict the following:

H5: *Perceived Privacy Protection (PPP) has a negative impact on System-specific Privacy Concerns (SPC).*

In our study we tested the effect of different techniques that a budding personalization provider (operationalized as American Personalization) can use to instill more favorable privacy related attitudes: increasing its reputation (i.e., becoming as reputable as Amazon), hiding its brand by referring to the personalization as happening in "the cloud," or implementing a privacy-enhancing technology (i.e., client-side personalization). Users observe these techniques as Provider Characteristics, and we hypothesize that these characteristics influence users' Perceived Privacy Protection (PPP):

H6: *On average, the manipulated Provider characteristics significantly influence Perceived Privacy Protection (PPP).*

As mentioned earlier, though, we argue that the effects of these different characteristics depend on the user's level of elaboration. The following section defines the specific hypotheses regarding these effects.

The Moderating Role of General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE)

One of the primary aims of our study is to examine how the effects of our provider characteristics change under the central vs. the peripheral route of privacy decision making. We operationalize the concepts of *motivation* and *ability* (which are known to determine users' level of elaboration) by General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE), respectively, and test their moderating effect on the relationship between privacy-related cues and users' attitudes.

Moderating the effect of Provider on Perceived Privacy Protection (PPP). As Figure 1 shows, we predict that General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE) moderate the relationship between Provider and Perceived Privacy Protection (PPP). GPC is defined as a privacy-related personal trait, shaped through an individual's overall online experiences (Miyazaki & Fernandez, 2001). Previous studies on ELM in privacy have used GPC as a measure of one's motivation to engage in issue-relevant thinking and cognitive elaboration when making privacy-related decisions (Bansal et al., 2008; Zhou, 2012). Privacy issues are of central importance to people with high levels of GPC, and those individuals will thus be more motivated to undertake closer scrutiny of the features of the

personalization provider, making systematic use of issue-relevant cues and information. People with low levels of GPC, on the other hand, will be less motivated to scrutinize the features of the personalization provider, and are thus more likely to use ostensive yet superficial cues in their evaluation process. As such, we predict that GPC plays a moderating role in determining the relative impact of Provider on the formation of provider-related privacy attitudes. Specifically, people with high levels of GPC will perceive less protection, unless the provider relies on privacy-enhancing techniques such as client-side personalization:

H7: *General Privacy Concerns (GPC) moderate the relationship between Provider and Perceived Privacy Protection (PPP). Specifically, PPP is likely to be lower for people with high levels of GPC, except for those who use client-side personalization (an actual privacy-enhancing technology).*

In a similar vein, Privacy Self-Efficacy (PSE) is related to one's ability to engage in cognitive elaboration of privacy-related cues and information (Bansal et al., 2008; Zhou, 2012). PSE refers to a person's belief in her capabilities and cognitive resources required to cope with privacy-related problems (LaRose & Rifon, 2007). The ELM suggests that people who are equipped with more knowledge and resources (e.g., high PSE) are more able to engage in extensive elaboration. In contrast, those with low ability (e.g., low PSE) will elaborate less and are more likely to rely on decisional shortcuts (Slovic et al., 2004)—cues that help them decide without needing to engage in cognitively elaborate processes. As such, we predict that PSE plays a moderating role in determining the relative impact of Provider on the formation of provider-related privacy attitudes. Specifically, people with high levels of PSE will perceive less protection, unless the provider relies on privacy-enhancing techniques:

H8: *Privacy Self-Efficacy (PSE) moderates the relationship between Provider and Perceived Privacy Protection (PPP). Specifically, PPP is likely to be lower for people with high levels of PSE, except for those who use client-side personalization.*

To sum up, the effect of the different provider characteristics depends on the user's elaboration likelihood. Specifically:

- Generally, people with higher motivation and ability will be more skeptical about the privacy protection offered by the provider. In the "American Personalization" condition (the baseline), users who use more central-route processing (i.e., who have a high level of General Privacy Concerns [GPC] and Privacy Self-Efficacy [PSE]) therefore have lower levels of Perceived Privacy Protection (PPP) than users who use more peripheral route processing.
- Reputation management may increase perceived protection in the peripheral route. Participants who predominantly use the peripheral route will therefore have higher levels of PPP in the Amazon condition than in the American Personalization condition. However, reputation management may only work in the peripheral route, so users in the Amazon condition who

use more central route processing will have lower levels of PPP than those who use more peripheral route processing.

- Privacy-enhancing techniques, on the other hand, do result in a perception of adequate protection both in the central route and the peripheral route. Therefore, we predict that users of client-side personalization will have similar levels of PPP in the central route and in the peripheral route. In other words, unlike the other conditions, PPP in the client-side condition will *not* be lower for users who use more central route processing; client-side personalization thus has an advantage over other methods for these users.

- Finally, although referring to the personalization as happening in the cloud removes focus from the provider itself, "the cloud" is still an unfamiliar concept, and deferring to the cloud still leaves the question of which entity actually manages the personal information unanswered. Therefore, cloud-based personalization results in lower levels of PPP for both peripheral and central route users.

Moderating the effect of System-specific Privacy Concerns (SPC) on Satisfaction (SAT). We also posit that General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE) moderate the relationship between System-specific Privacy Concerns (SPC) and Satisfaction (SAT). In general, privacy concerns are salient elements affecting the user experience of personalization systems (Teltzrow & Kobsa, 2004). Specifically, people may have certain expectations about privacy, and they may use these expectations as key evaluation standards in determining their satisfaction with online services (Chen et al., 2012; Fox et al., 2000).

Extending the ELM perspective, we posit that GPC and PSE will moderate the impact of SPC on SAT. In brief, privacy-related attitudes are more salient and have more impact on overall satisfaction for those who are motivated (GPC) and able (PSE) to reflect upon these attitudes. Hence, high GPC and PSE will lead to a stronger impact of SPC on SAT. Therefore, we advance the following hypotheses:

> **H9**: *General Privacy Concerns (GPC) moderate the relationship between System-specific Privacy Concerns (SPC) and Satisfaction (SAT). Specifically, the effect of SPC on SAT will be stronger for people with high levels of GPC.*

> **H10**: *Privacy Self-Efficacy (PSE) moderates the relationship between SPC and SAT. Specifically, the effect of SPC on SAT will be stronger for people with high levels of PSE.*

Online Experiment

Procedure

We test the developed hypotheses in an online experiment with a smartphone app that gives personal recommendations on a wide variety of topics. We summarize this experiment here, whereas Appendix 1 provides more details including screen shots.

The experiment started with an online survey at the survey website "Instantly," which introduced participants to an Android app named Check-it-Out (CiO) that purportedly tracks and analyzes their smartphone use as input for its recommendations. It then presents the personalization provider (consistent with the condition to which a participant was assigned; see below), as well as three examples of the personalized services of Check-it-Out.

To verify participants' comprehension (which we deemed important for the validity of the experiment), the survey asks 15 questions about the personalized services that Check-it-Out provides, the whereabouts of participants' collected personal data, and the location at which the personalization logic is performed. Participants then install the Check-it-Out app on their own smartphones, and answer the demographics and context data requests under the pretense that Check-it-Out would give better recommendations if it has additional information about them.

Once participants install the "Check-it-Out" app, the introduction screen explains once what happens with the personal information they are about to submit (consistent with the assigned experimental condition; see below). The app subsequently asks participants for various kinds of demographic information, alternating with requests for permission to track various aspects of their smartphone usage context. Table 1 in the Results section lists these questions and the order in which they appear. Participants are free to decline the answer or permission for any request.

After answering all demographics and context data requests, participants return to the online survey and answer the survey items that are listed in Table 2 of the Results section.

Experimental Conditions

In the experiment, we manipulate the personalization provider in a way that allows us to argue about the effectiveness of different strategies to instill more favorable privacy-related attitudes and increase users' disclosure. We consider an unknown personalization provider as our baseline condition, and test the effects of three methods to increase disclosure. Each method has implications for the whereabouts of participants' collected personal data, which are described as follows:

- American Personalization: "all data are sent to American Personalization"
- Amazon: "all data are sent to Amazon"
- The cloud: "all data are sent to the Cloud"
- Client-side personalization: "all data remain on your smartphone"

The introductory survey explains these implications to participants in detail, and verifies their understanding through comprehension questions. Moreover, the introductory screen of the CiO app explains to users what will happen with their personal data: In the client-side condition, the app informs participants that all entered data will remain on their phone, and participants are encouraged to turn off their network connection. In the three other conditions, participants are told that the data they enter will be sent to

TABLE 1. Items requested by Check-it-out.

Type of data	Seq. #	Item	Disclosure	Factor loading
Demographics	1	Phone data plan	94.9%	
Alpha: 0.84	3	Household composition	87.4%	
AVE: 0.694	5	Field of work	91.5%	
Factor correlation: 0.478	7	Housing situation	85.9%	0.790
	9	Relationship status	93.6%	0.825
	11	Children	90.0%	0.966
	13	Household income	80.8%	0.872
	15	Household savings	66.7%	0.848
	17	Household debt	68.5%	0.790
	19	Race	93.1%	0.822
	21	Political preferences	82.8%	0.771
	23	Workout routine	85.1%	0.798
Context	2	Recommendation browsing	79.5%	0.790
Alpha: 0.88	4	Location	50.5%	0.805
AVE: 0.681	6	App usage	72.8%	0.864
Factor correlation: 0.478	8	App usage location	56.2%	0.902
	10	App usage time	70.5%	0.949
	12	Web browsing	56.9%	0.824
	14	Calendar data	49.7%	0.753
	16	E-mail messages	16.4%	0.771
	18	Phone model	83.3%	0.819
	20	Accelerometer data	58.2%	0.811
	22	Microphone	17.9%	0.769
	24	Credit card purchases	0.0%	

TABLE 2. Survey items.

Subjective construct	Items	Factor loading
Self-anticipated satisfaction with Check-it-Out (SAT)	Check-it-Out is useful	0.898
	Using Check-it-Out makes me happy	0.885
Alpha: 0.92	Using Check-it-Out is annoying	−0.703
AVE: 0.751	Overall, I am satisfied with Check-it-Out	0.925
	I would recommend Check-it-Out to others	0.903
	I would quickly abandon using this system	
System-specific privacy concerns (SPC)	Check-it-Out has too much information about me	0.756
	Check-it-Out does not know anything I would be uncomfortable sharing with it	
Alpha: 0.76	I felt tricked into disclosing more information than I wanted	
AVE: 0.660	I find the questions intrusive that Check-it-Out asks me	0.847
	I'm afraid Check-it-Out discloses information about me to third parties	
Perceived privacy protection (PPP)	I feel my personal data are safe [on my smartphone / at American Personalization / at Amazon / in the Cloud].	0.917
Alpha: 0.95	I feel [my smartphone / American Personalization / Amazon / the Cloud] will not share my personal data with anyone.	0.954
AVE: 0.887	I feel my interests will be protected when my personal data are [on my smartphone / with American Personalization / with Amazon / in the Cloud].	0.953
General online privacy concerns (GPC)	It usually bothers me when online companies ask me for personal information	0.939
	It bothers me to give personal information to so many online companies	0.957
Alpha: 0.89	Online companies may collect any information about me because I have nothing to hide	
AVE: 0.806	I am concerned that online companies are collecting too much personal information about me	0.787
	I am not bothered by data collection, because my personal information is publicly available anyway	
Privacy self-efficacy (PSE)	It's easy to figure out which sites you can trust on the Internet	
Alpha: 0.85	I am confident I know how to protect my credit card information online	
AVE: 0.656	I know how to identify sites with secure servers	0.800
	I know how to evaluate online privacy policies	0.751
	It's easy to set up dummy email account to shield my identity	
	I know how to change the security settings of my browser to increase privacy	0.867
	I know how to use a virus scanning program	0.824
	I am able to protect myself against the release of personal information	
	I know how to block unwanted E-mails	0.804
	Overall, I am confident that I can protect my privacy online	

American Personalization / Amazon / the Cloud, respectively. If their network connection is disabled, participants are asked to turn it on; the app does not proceed otherwise.

Subjects

We recruited study participants through the crowdsourcing platform Mechanical Turk (MTurk) and a similar corporate service, and posted recruitment ads on Craigslist in eight metropolitan areas across the United States. Because MTurk's General Policies do not allow HITs that require workers to download software, we followed Kanich, Checkoway, and Mowery (2011) and gave participants a choice between the full study including app download, or merely completing the survey part for a much lower reward. Of those who chose this latter option, 63.5% indicated that they did not own an Android phone. We found no significant differences in privacy concerns between those who chose to download the CiO app and this "control group" that did not, allaying fears that only the less privacy concerned would self-select to download the app.

The data of 390 subjects[3] were used in the statistical analysis. Nine subjects with a very short interaction time (less than 9 minutes 40 seconds) were removed from the analysis. Participants' average score on the 15-item comprehension test was 13.4 out of 15 (with only eight subjects lower than 10), which we regarded as a very satisfactory level of comprehension of the selected personalization condition and its privacy implications.

Measurements

Behavioral measure: information disclosure. We tracked participants' interactions with the CiO app to measure their disclosure behavior. In line with Knijnenburg et al. (2013a), we treat demographics and context disclosure as two distinct behavioral factors. The 12 demographics items and 12 context items are taken from Knijnenburg and Kobsa (2013). Table 1 displays the results of the confirmatory factor analysis (CFA) for the 24 disclosures. Because the indicators are dichotomous, we used a weighted least squares estimator that treats the items as ordered-categorical, thereby not assuming multivariate normality.

Among the observed behaviors, not a single participant agreed to disclose their credit card purchases, so this item had to be removed from the analysis because of zero

variance. Moreover, the initial scale-refinement process suggested three additional items ("phone data plan," "household composition," and "field of work") should be removed because of low communality (R^2-values of 0.467, 0.451, and 0.419, respectively, with the next-lowest R^2 being 0.568). As shown in Table 1, the two factors had a low to moderate positive correlation, good convergent validity (average variance extracted [AVE] > 0.50, Cronbach's alpha > 0.70) and good discriminant validity (square root of AVE larger than the factor correlation). This confirms H1, which states that Demographics and Context disclosure are two separate but positively correlated factors.

Postexperimental questionnaire. In the postexperimental questionnaire, we measured five subjective factors using 29 statements for which users were asked to state their agreement or disagreement on a seven-point scale:

Self-anticipated satisfaction with Check-it-Out (SAT): participants' anticipated outcomes of, or experience with, the personalization system. Note that users do not actually get to use the system, hence the anticipated nature[4] of this construct. SAT is similar to the "preference for benefits" constructs in Hui, Tan, and Goh (2006), "disclosure-privacy benefits" in Xu et al. (2009) and "perceived benefits of info disclosure" in (Xu et al., 2011). The items are taken from Knijnenburg and Kobsa (2013), and were originally developed for the framework of user-centric evaluation of recommender systems by Knijnenburg, Kobsa, and Saldamli (2012) and Knijnenburg, Willemsen, Gantner, Soncu, and Newell (2012).

System-specific privacy concerns (SPC): Participants' concerns with the amount and the sensitivity of the information that CiO collects. This factor is a system-specific analogy to the "collection" factor of the Internet Users Information Privacy Concerns scale (Malhotra, Kim, & Agarwal, 2004). Its items are taken from the "perceived privacy threats" factor of Knijnenburg and Kobsa (2013).

Perceived privacy protection (PPP): Users' perception of protection (or, inversely, the fear of absence of that protection) by the provider against unintended use of the data (either by the provider or a third party). This factor is a provider-specific analogy to the "unauthorized secondary use" and "improper access" factors of the Concern For Information Privacy scale (Smith, Milberg, & Burke, 1996) and borrows from the benevolence subconstruct of trust (Mayer, Davis, & Schoorman, 1995). Its item definitions were informed by Bhattacherjee (2002), Chellappa (2008), Pirim, James, Boswell, Reithel, and Barkhi (2008), and Shin (2010).

General online privacy concerns (GPC): Users' concern about the collection and misuse of personal data by online

---

[3]Two hundred fifty-one participants originated from MTurk, 105 from Craigslist, and 34 from the corporate recruitment platform. As part of the standard procedures for statistical evaluation, we tested whether our results were invariant with regard to different recruitment sources, and found this to be the case: There were no statistically significant differences in the outcomes between recruitment sources (i.e., Source → GPC, PSE, SPC, PPP, SAT, DEM, CON were all nonsignificant), and no difference in effects of the conditions on relevant outcomes for the different recruitment sources (Source × Provider → SPC, PPP, SAT, DEM, CON were all nonsignificant). This invariance with regard to recruitment source increases the robustness of our results.

[4]We specifically opted to measure anticipated satisfaction, because users of a personalized app typically have to rely on their self-anticipated (and not postusage) satisfaction when making disclosure decisions. Anticipated satisfaction is thus a more realistic predictor of users' disclosure behavior than postusage satisfaction.

	sqrt(AVE)	SPC	PPP	SAT	GPC	PSE
SPC	0.812		−0.451	−0.710	0.612	−0.025
PPP	0.942	−0.451		0.504	−0.285	0.172
SAT	0.866	−0.710	0.504		−0.319	0.112
GPC	0.898	0.612	−0.285	−0.319		0.115
PSE	0.810	−0.025	0.172	0.112	0.115	

companies. GPC can be seen as a personal trait; a general feeling of concern that is not directed to any specific system or company. The items for this construct are taken from the "collection concerns" construct in Knijnenburg & Kobsa (2013), which is an expansion of the "collection" factor or 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).

Privacy self-efficacy (PSE): Users' feeling of empowerment to engage in privacy protection behaviors, such as the use of privacy enhancing technologies. We used the "privacy self-efficacy" scale developed by LaRose and Rifon (2007).

Table 2 displays the results of the confirmatory factor analysis (CFA) for the 29 items, and Table 3 shows the correlations between factors. The initial scale-refinement process suggested 11 items should be removed (marked out in grey):

- Six items had low communalities (SPC2, $R^2 = 0.324$; SPC3, $R^2 = 0.318$; GPC3, $R^2 = 0.173$; GPC5, $R^2 = 0.212$; PSE1, $R^2 = 0.397$; PSE5, $R^2 = 0.437$; the next-lowest $R^2 = 0.547$).
- Three items were subsequently removed because of high residual cross-loadings with other factors (PSE10 with SPC, $\chi^2 = 47.29$; PSE8 with PPP, $\chi^2 = 30.76$; and PSE2 with PPP, $\chi^2 = 41.28$).
- Two items were removed because of theoretical misfit with the defined construct (SAT6 measured usage intentions rather than satisfaction, SPC5 measured disclosure risk rather than concerns about data collection).

As shown, the factors had a good convergent and discriminant validity.

To evaluate the difference in effects when users use the peripheral versus the central route, we classify participants into different groups based on their level of General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE). We employ a mixture factor analysis (MFA) with a varying number of classes to find the optimal classification of users according to their level of GPC and PSE. We find that a four-class solution is the best (i.e., a Lo-Mendel-Rubin adjusted LRT test shows that the five-class solution does not fit significantly better; $p = 0.077$), resulting in the following classes:

1. Low concerns and efficacy [GPC_{low}, PSE_{low}]: 85 participants

2. Low concerns but high efficacy [GPC_{low}, PSE_{high}]: 132 participants
3. High concerns but low efficacy [GPC_{high}, PSE_{low}]: 116 participants
4. High concerns and efficacy [GPC_{high}, PSE_{high}]: 57 participants

We use these classes to dichotomize GPC and PSE (i.e., we assign GPC = low [0] to classes 1 and 2, GPC = high [1] to classes 3 and 4, PSE = low [0] to classes 1 and 3, and PSE = high [1] to classes 2 and 4).

Results

We tested our hypotheses using structural equation modeling (SEM). To overcome identification problems, we modeled demographics and context disclosure (DEM and CON) as a set of dichotomous repeated measures rather than a set of latent factors. Moreover, we dichotomized General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE) using the clustering results presented earlier.

We tested our hypotheses in four iterative steps, improving the model based on empirical evidence in each step. First, to estimate all main effects and check for misspecifications of these effects, we tested a model without moderating effects of GPC and PSE. Inspecting the coefficients and modification indices of this model, we made the following changes:[5]

- Remove the effect of System-specific Privacy Concerns on context disclosure (SPC → CON; H2b), not significant, $p = 0.126$
- Remove the effect of Satisfaction on demographics disclosure (SAT → DEM; H3a), not significant, $p = 0.139$
- Add a direct effect of provider on System-specific Privacy Concerns (Provider → SPC), large modification index, $p < 0.001$
- Add a direct effect of Perceived Privacy Protection on Satisfaction (PPP → SAT), large modification index, $p < 0.001$

We then introduced the moderating effects of General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE). Note that (per convention) this model also includes main effects of GPC and PSE on Perceived Privacy Protection (PPP) and Satisfaction (SAT). Inspecting the coefficients and modification indices of this model, we made the following final changes:

- Remove the moderating effect of Privacy Self-Efficacy on the effect of provider on Perceived Privacy Protection (PSE × Provider → PPP), not significant, $p = 0.255$
- Add a main effect of General Privacy Concerns on System-specific Privacy Concerns (GPC → SPC), large modification index, $p < 0.001$
- Add a main effect of General Privacy Concerns on context disclosure (GPC → CON), large modification index, $p = 0.001$

---

[5]In general, models can be trimmed or built based on theoretical and empirical standards (Kline, 2004).
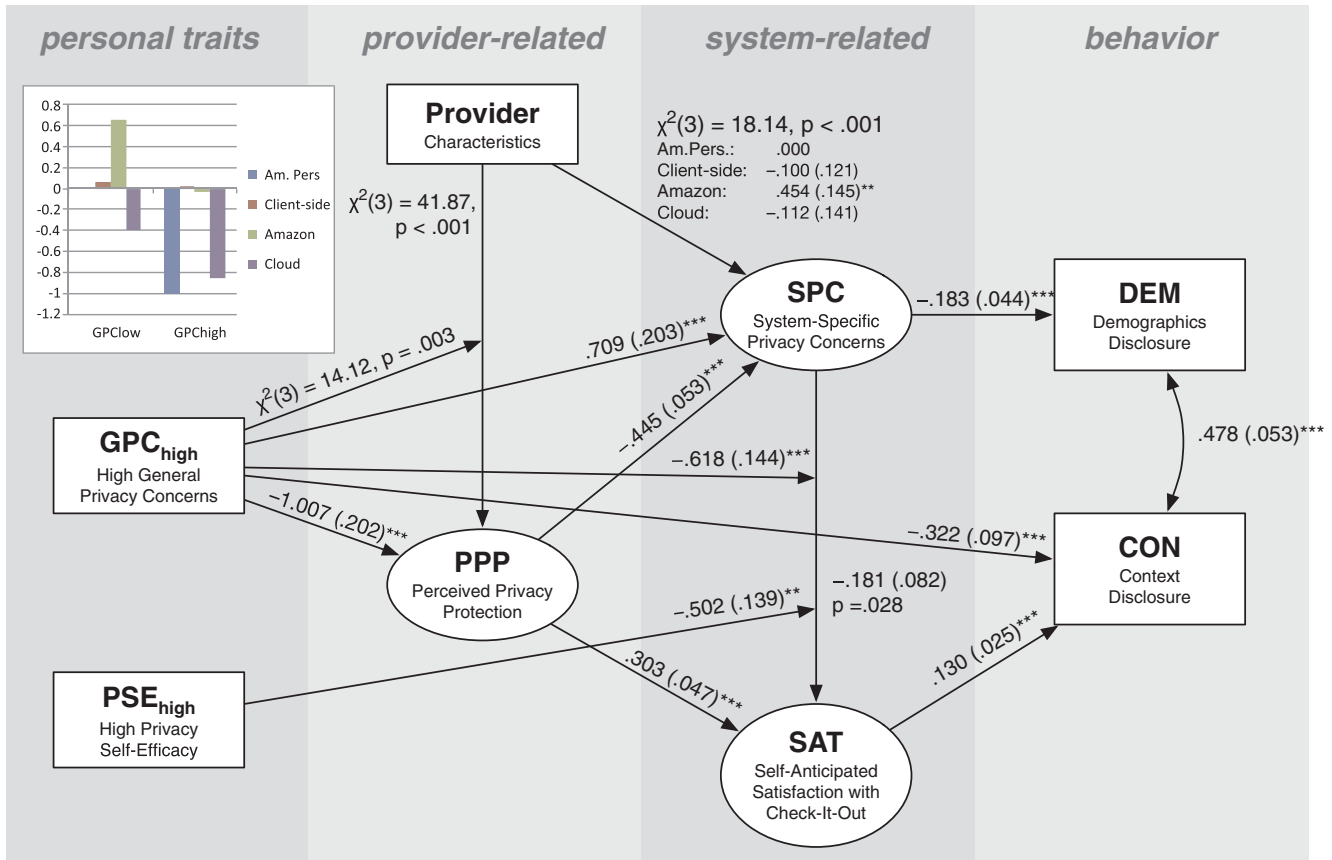
FIG. 2. Final model and estimates of its path coefficients. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

The final model had an excellent fit ($\chi 2(749) = 715.89$, $p = 0.80$; RMSEA = .000, 90% CI: [.000, .006]; CFI = 1.00). The model and the estimates of its path coefficients are presented in Figure 2. We see how users in different elaboration modes evaluate the perceived privacy protection of the provider, which influences their system-specific privacy concerns, their self-anticipated satisfaction, and eventually their disclosure.

The most interesting aspect of our model is the variation in the effects of Provider on Perceived Privacy Protection (PPP), and of System-specific Privacy Concerns (SPC) on Satisfaction (SAT) for the different elaboration modes. The overall effect of Provider on PPP is significant, $\chi 2(3) = 72.449$, $p < 0.001$, and the effect differs significantly, $p = 0.003$, for people with different levels of General Privacy Concerns (GPC). Assuming that people with low GPC are more inclined to use the peripheral route and people with high GPC the central route, we find that:

- For American Personalization, users who predominantly use the central route feel less protected than users who predominantly use the peripheral route, $\beta_{\text{difference}}$ is -1.007, $p < 0.001$.
- Client-side personalization is not seen as more protective than American Personalization in the peripheral route, $\beta = 0.067$, $p = 0.64$. However, in contrast to American Personalization, the privacy protection afforded by client-side personalization is the same in both routes, $\beta_{\text{difference}} = -0.053$, $p = 0.76$.

Participants who predominantly use the peripheral route think that Amazon protects them more, $\beta = 0.649$, $p < 0.001$, whereas this effect is significantly lower when they predominantly use the central route, $\beta_{\text{difference}} = -0.685$, $p < 0.001$.
- Participants who predominantly use the peripheral route think that personalization in the Cloud is marginally worse in terms of protection than at American Personalization, $\beta = -0.394$, $p = 0.067$, and they think it is even worse when they predominantly use the central route, $\beta_{\text{difference}} = -0.465$, $p = 0.022$.

The overall effect of System-specific Privacy Concerns (SPC) on Satisfaction (SAT) is significant, $\beta = -0.511$, $p < 0.001$, but as Figure 2 shows, General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE) each moderate the effect. The moderation effects are additive, that is, the two-way interactions of $\text{GPC} \times \text{SPC} \rightarrow \text{SAT}$ and $\text{PSE} \times \text{SPC} \rightarrow \text{SAT}$ were significant whereas the three-way interaction of $\text{GPC} \times \text{PSE} \times \text{SPC} \rightarrow \text{SAT}$ was not. The moderations result in the following effects:

- For participants with low GPC and PSE, there is only a small effect of SPC on SAT: $\beta = -0.181$, $p = 0.028$.
- For participants with low GPC but high PSE, there is an effect: $\beta = -0.683$, $p < 0.001$.
- For participants with high GPC but low PSE, there is also an effect: $\beta = -0.799$, $p < 0.001$.
- The largest effect is for participants with high GPC and PSE (the effects are additive): $\beta = -1.301$, $p < 0.001$.
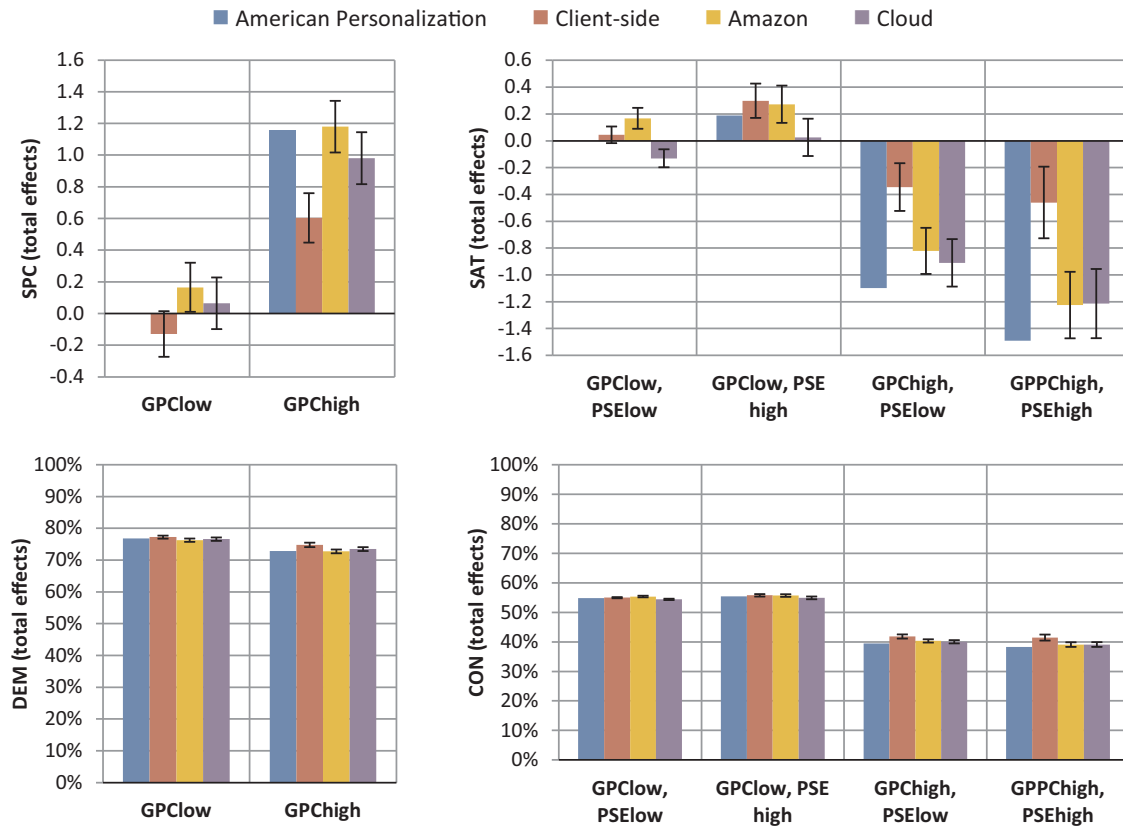
FIG. 3. Total effects of provider on SPC, SAT, and disclosure (DEM and CON). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Given that people with high GPC and high PSE are more inclined to use the central route than people with low GPC and low PSE, we argue that central route processing leads to a stronger effect of SPC on SAT than peripheral route processing.

The effects presented in the model create an interesting interplay between direct and mediated effects, some moderated by General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE), and some not. This leads to the total effects of Provider on System-specific Privacy Concerns (SPC), Satisfaction (SAT), and disclosure (DEM and CON) displayed in Figure 3. These graphs show that although low-GPC users in the Amazon condition report the highest Perceived Privacy Protection (PPP; which in turn reduces their SPC), the Amazon condition also has a direct positive effect on SPC, largely negating the benefit it receives from perceived protection (see Figure 3, top left). As a result, only the users with low GPC and PSE perceive a higher SAT in the Amazon condition; users with either high GPC or high PSE (or both) perceive more satisfaction in the Client-side condition (see Figure 3, top right). Ultimately, this leads to a very small increase in demographics (DEM) and context (CON) disclosure in the Client-side condition (a 2.7% and 8.5% increase, respectively).

Discussion

This study investigates techniques to instill more favorable attitudes towards personalization and increase disclosure by manipulating characteristics of the personalization provider. Using the ELM, we also specify conditions under which the manipulated provider characteristics have more or less impact on users' beliefs and attitudes. Finally, we propose an integrative research model that specifies the relationships between these manipulations, privacy attitudes, and actual disclosure behavior.

Table 4 summarizes the results of our hypothesis testing. Overall, the findings provide support for most hypotheses: people distinguish between demographics (DEM) and context (CON) disclosure (H1); disclosure is affected primarily by System-specific Privacy Concerns (SPC) (H2a) and self-anticipated satisfaction (SAT) (H3b); and significant relationships exist between SPC and SAT (H4), Perceived Privacy Protection (PPP) and SPC (H5), and provider characteristics and PPP (H6).

More importantly, the findings highlight the alternative processes or mechanisms through which people evaluate privacy protection (PPP) and outcomes (SAT) when using a personalization system. Specifically, in line with H7 and the

TABLE 4. Summary of hypothesis testing.

Hypothesis	Effect	Confirmed
H1	DEM <-> CON	Yes
H2a	SPC → DEM	Yes
H2b	SPC → CON	No
H3a	SAT → DEM	No
H3b	SAT → CON	Yes
H4	SPC → SAT	Yes
H5	PPP → SPC	Yes
H6	Provider → PPP	Yes
H7	GPC × Provider → SPC	Yes
H8	PSE × Provider → SPC	No
H9	GPC × SPC → SAT	Yes
H10	PSE × SPC → SAT	Yes

ELM, the effect of Provider on PPP indeed differed for participants with low versus high levels of General Privacy Concerns (GPC). Participants with high GPC perceived less protection *except* when they happened to use the client-side personalization version of Check-it-Out. This confirms that participants who predominantly used the peripheral route as well as those who predominantly used the central route perceived this technique as adequate privacy protection. On the other hand, reputation management (i.e., becoming as reputable as Amazon) was perceived as adequate protection only by those participants who predominantly used the peripheral route. Referring to the personalization as happening in "the cloud" resulted in a lower perceived protection in either route.

Similarly, consistent with H9 and H10, the degree to which System-specific Privacy Concerns (SPC) influence Satisfaction (SAT) depends on individual differences in users' motivation and ability to attend to privacy-related factors. Specifically, SPC has a stronger impact on the satisfaction of participants who predominantly used the central route (i.e., high General Privacy Concerns [GPC] and Privacy Self-Efficacy [PSE]), compared to participants who predominantly used the peripheral route (i.e., low GPC and PSE).

Our confirmation of H9 is in line with Joinson, Reips, Buchanan, and Schofield (2010), who find that trust mediates the effects of privacy concerns, and that trust can even *compensate* for high levels of privacy concerns. We consider SPC to be a subcomponent of trust, and we can therefore confirm that trust mediates the effect of privacy concerns on satisfaction and disclosure (GPC → SPC → DEM and GPC → SPC → SAT → CON). Moreover, for participants with high concerns (GPC$_{high}$), trust (SPC) has a relatively stronger effect on Satisfaction, which implies that higher levels of trust can indeed compensate the effects of high levels of privacy. This is clearly visible in the graphs for Satisfaction in Figure 3: The effect of privacy concerns (i.e., the difference between GPC$_{low}$ and GPC$_{high}$) is weakest for the condition that receives the most trust (i.e., the lowest SPC), namely client-side personalization.

Though the findings support most hypotheses in this study, we failed to find support for three. The nonsignificant relationships between System-specific Privacy Concerns (SPC) and Context disclosure (CON; H2b) and Satisfaction (SAT) and Demographics disclosure (DEM; H3a) can be explained by the multidimensional nature of information disclosure behaviors as described previously. Further, it is worthwhile to note that the moderating effect of Privacy Self-Efficacy (PSE) on the relationship between Provider and Perceived Privacy Protection (PPP; H8) was not significant. A possible reason for this is that provider characteristics are beyond the users' control; therefore, users' self-efficacy is arguably neither a limiting nor a motivating factor in elaborating on the implications of these provider characteristics. In this specific case it thus makes sense that users' use of the peripheral versus the central processing route is primarily governed by motivation rather than ability.

Limitations and Future Work

This study has some limitations that point out directions for future research. First, although we aimed to develop an integrative theoretical model, we nevertheless needed to focus on a relatively small number of central constructs in order to make the model as parsimonious and comprehensible as possible. We therefore disregarded some potentially significant factors, such as trust whose role has been documented in prior privacy and personalization research (Briggs et al., 2004; Chellappa & Sin, 2005; Komiak & Benbasat, 2006). We decided to exclude the construct of trust in favor of some of the *subconstructs* of trust that are represented by other variables in our model: Our notion of perceived protection (PPP) is related to the subconstruct of "benevolence," and privacy concerns (SPC) is related to "integrity." However, we suggest that our research model should be further validated and extended by exploring the role of other important factors, such as trust and prior experience.

Second, although we designed the experiment in such a way that participants disclosed their personal information to an *actual* Android app, the app itself did not give any real personalized recommendations in return. Beyond some generic introductory examples, users were thus left in the dark regarding the quality of the recommendations that would result from their disclosed information. This experimental design mimics real life usage of personalized apps, where permissions have to be given before the system can be used. Users have to rely on their self-anticipated rather than postusage attitudes when making disclosure decisions.

A personalized system could however also operate "conversationally" and already give users recommendations based on the first few pieces of disclosed information, further improving the personalization with additional disclosures. Such a system, which we currently develop, would allow for a more direct measurement of the privacy-personalization paradox because users could trade off the

privacy sensitivity of each item with the actually observed (rather than anticipated) benefits in terms of personalization quality improvements. Moreover, such system could collect *actual* context data, and a study with this system could run for an extended time period in order to see if users change the context data collection settings over time.

Finally, as we made a first foray into user perceptions of client-side and cloud-based personalization, these experimental conditions do not probe into the different mechanisms that could be used to implement and support these types of personalization. For example, for client-side personalization one could further investigate the following questions: How does a system "prove" that data are not being transmitted (without telling users to just switch off their Internet connection)? How are client-side data stored securely on the device? How can these data be transferred to a new device? Similarly, for cloud-based personalization one could investigate the following questions: Where does the data reside? Who has access to the data? Who is responsible for the security of the data? In the current article we investigated users' initial perceptions only; future work could further unpack the implications of specific implementations of client-side and cloud-based personalization.

Managerial Implications

Managers and personalization providers will benefit from our integrative theoretical approach focusing on actual behavior. For instance, if observed user behaviors do not meet their expectations (e.g., low levels of information disclosure), the integrative causal model can allow them to identify those user attitudes that they will need to improve to bring about the desired behavior. To do this, they can use—but are not limited to—the techniques presented in this article: reputation management, "cloud branding," or client-side personalization. Importantly, they can use our findings from the ELM to argue which market segment is most likely to be susceptible to these techniques: people with low or rather high privacy concerns and privacy self-efficacy. Our findings show that people place varying degrees of importance on different types of provider characteristics when forming judgments on privacy protection and (ultimately) information disclosure. When catering to an audience with a wide range of privacy preferences, a mixed strategy of reputation management and privacy-preserving personalization techniques therefore seems to be the best solution.[6] However, because privacy has its strongest impact on satisfaction for users who predominantly use the central route, techniques that cater to the central route (e.g., privacy-preserving personalization techniques such as client-side personalization) will result in the highest overall satisfaction.

The results also show that the effect of provider characteristics on disclosure behavior was small (see Figure 3),

---

[6]Using a mixed strategy of reputation management and privacy-preserving personalization techniques also ensures that perceived privacy and actual privacy improvements occur simultaneously.

and mediated by other intervening factors such as System-specific Privacy Concerns (SPC), Perceived Privacy Protection (PPP), and Satisfaction (SAT). The findings suggest that merely deploying a privacy-enhancing feature may not result in substantial changes in user behavior. In order to elicit desired user behavior, developers and managers of privacy-enhanced personalization should also make it clear to users that an IT design like client-side personalization has tangible benefits in terms of protecting personal information in a more effective, secure, or easier way. In Kobsa, Knijnenburg, and Livshits (2014), we use additional data to elaborate on this idea for the special case of client-side personalization.

It is important to note that participants generally perceived Amazon to be more privacy-protecting, but that they also had higher system-specific privacy concerns regarding Amazon, which, overall, more than offset this positive effect. Participants also generally perceived the cloud to be less privacy-protecting. The latter means that the strategy of hiding the brand by deferring to the cloud does not work. Arguably, participants perceived even less reason to trust "the cloud" than American Personalization; those who predominantly used the peripheral route may have found "the cloud" even less familiar-sounding, and those who predominantly used the central route may have had even more concerns about the actual privacy protection offered by the cloud.

Contributions to Research and Theory Development

Our work contributes to personalization versus privacy research, by integrating and thus reconciling the "privacy calculus view," and the "heuristic shortcuts view" using the Elaboration Likelihood Model. Most research on personalization versus privacy assumed that users' attitudes and behavior are either determined primarily by instrumental beliefs constructed through deliberative cognitive processes (Awad & Krishnan, 2006; Chellappa & Sin, 2005; Kobsa & Teltzrow, 2005; Li & Unger, 2012; Sutanto et al., 2013; Xu et al., 2011), or by heuristic shortcuts (Acquisti & Grossklags, 2008; Acquisti et al., 2013; Adjerid et al., 2013; Cho et al., 2010; LaRose & Rifon, 2006; Lowry et al., 2012). These assumptions result in wildly differing recommendations for improving privacy-related attitudes and behavior; that is, from increasing transparency and control (which only works when users are deliberate in their decision process) to using privacy "nudges" to subtly influence users' privacy decisions (which primarily works when users use heuristic shortcuts). Our findings suggest that techniques to instill more favorable privacy-related attitudes towards personalization should be sensitive to the fact that users come to judgments and decisions through multiple routes ranging from cognitively elaborative processes to decisional shortcuts: When elaboration likelihood is low, peripheral cues such as reputation management take on more importance, leading to higher levels of Perceived Privacy Protection (PPP). When elaboration is high, however, such

convenient yet superficial privacy cues have virtually no effect on PPP. Similarly, when elaboration likelihood is low, System-specific Privacy Concerns (SPC) play a much smaller role in determining users' anticipated Satisfaction with a system (SAT) than when elaboration likelihood is high. The findings thus show that people place varying degrees of importance on different types of cues when forming judgments on privacy protection and (ultimately) information disclosure. Research should specify boundary conditions under which different factors or techniques play a more or less significant role, and the present study revealed at least two important individual-difference factors governing these processes, namely General Privacy Concerns (GPC) and Privacy Self-Efficacy (PSE).

The findings also make an important contribution to privacy research by highlighting the multidimensional nature of information disclosure behavior. Knijnenburg and Kobsa (2013) find that treating these behaviors as multidimensional can attain more powerful models of information disclosure behaviors. Confirming these findings, the present study shows that System-specific Privacy Concerns (SPC) has a direct effect on Demographic information disclosure (DEM) but not on Context information disclosure (CON), whereas Satisfaction (SAT) has an effect on CON but not on DEM. This suggests that users mainly consider unintended and unauthorized use of their personal information when deciding whether to disclose their demographics. Yet, when deciding whether to grant an app access to their context information, users rather consider how well the app can satisfy their needs. This makes intuitive sense: context information is often more ambiguous than demographics information, and hence systems will need to take an additional interpretative step in order to use context information as input for personalization. Users will therefore disclose their context information only if they are confident that the system is competent enough to correctly perform this interpretation and provide accurate personalized results.

More generally, the integrative research model tested in our study provides a useful conceptual framework to reveal complex relationships among personal traits, attitudes and behaviors related to privacy and personalization. The model also allows us to make predictions about how changes in one factor (e.g., the provision of a privacy-enhancing feature) lead to changes in outcomes such as users' attitudes and behavior. Several studies have tested privacy-enhancing interventions, only to find disappointing results (Brandimarte, Acquisti, & Loewenstein, 2010; Knijnenburg et al., 2012; Patil & Kobsa, 2009). Our results suggest that taking mediating and moderating constructs into account (e.g., GPC, PSE, SPC, SAT) may either increase the statistical robustness of the effects of privacy-enhancing interventions on disclosure behavior, or provide a detailed explanation why such an effect does not exist. Because there is little research that examines those factors and aspects simultaneously, our findings contribute to the development of a theoretical foundation from which issues of privacy can be further explored.

References

Acquisti, A., & Grossklags, J. (2008). What can behavioral economics teach us about privacy? In A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, & C. Lambrinoudakis (Eds.), Digital privacy: Theory, technologies, and practices (pp. 363–377). New York/London: Auerbach Publications.

Acquisti, A., John, L.K., & Loewenstein, G. (2012). The impact of relative standards on the propensity to disclose. Journal of Marketing Research, 49(2), 160–174. http://doi.org/10.1509/jmr.09.0215

Acquisti, A., Adjerid, I., & Brandimarte, L. (2013). Gone in 15 seconds: The limits of privacy transparency and control. IEEE Security & Privacy, 11(4), 72–74. http://doi.org/10.1109/MSP.2013.86

Adjerid, I., Acquisti, A., Brandimarte, L., & Loewenstein, G. (2013). Sleights of privacy: Framing, disclosures, and the limits of transparency. In Proceedings of the Ninth Symposium on Usable Privacy and Security (pp. 9:1–9:11). New York: ACM. http://doi.org/10.1145/2501604.2501613

Andrade, E.B., Kaltcheva, V., & Weitz, B. (2002). Self-disclosure on the web: The impact of privacy policy, reward, and company reputation. In S.M. Broniarczyk & K. Nakamoto (Eds.), Advances in consumer research (pp. 350–353). Valdosta, GA: Association for Consumer Research.

Angst, C.M., & Agarwal, R. (2009). Adoption of electronic health records in the presence of privacy concerns: The elaboration likelihood model and individual persuasion. MIS Quarterly, 33(2), 339–370. Retrieved from http://dl.acm.org/citation.cfm?id=2017424.2017430

Awad, N.F., & 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), 13–28.

Bansal, G., Zahedi, F., & Gefen, D. (2008). The moderating influence of privacy concern on the efficacy of privacy assurance mechanisms for building trust: A multiple-context investigation. In ICIS 2008 Proceedings Paper 7. Paris, France. Retrieved from http://aisel.aisnet.org/icis2008/7

Berendt, B., Günther, O., & Spiekermann, S. (2005). Privacy in e-commerce: Stated preferences vs. actual behavior. Communications of the ACM, 48(4), 101–106. http://doi.org/10.1145/1053291.1053295

Berkovsky, S., Eytani, Y., Kuflik, T., & Ricci, F. (2007). Enhancing privacy and preserving accuracy of a distributed collaborative filtering. In Proceedings of the 2007 ACM Conference on Recommender Systems Paper 7. Minneapolis, MN: ACM. http://doi.org/10.1145/1297231.1297234

Bhattacherjee, A. (2002). Individual trust in online firms: Scale development and initial test. Journal of Management Information Systems, 19(1), 211–241. http://dx.doi.org/10.1080/07421222.2002.11045715

Brandimarte, L., Acquisti, A., & Loewenstein, G. (2010). Misplaced confidences: Privacy and the control paradox. In Ninth Workshop on the Economics of Information Security (WEIS). Cambridge MA. Retrieved from http://weis2010.econinfosec.org/papers/session2/weis2010_brandimarte.pdf

Briggs, P., Simpson, B., & Angeli, A.D. (2004). Personalisation and trust: A reciprocal relationship? In C.-M. Karat, J.O. Blom, & J. Karat (Eds.), Designing personalized user experiences in eCommerce (pp. 39–55). Dordrecht, Netherlands: Kluwer Academic Publishers. http://dx.doi.org/10.1007/1-4020-2148-8_4

Brodie, C., Karat, C.-M., & Karat, J. (2004). Creating an E-Commerce Environment Where Consumers Are Willing to Share Personal Information. In C.-M. Karat, J.O. Blom, and J. Karat (Eds.), Designing personalized user experience for eCommerce (pp. 185–206). Dordrecht, Netherlands: Kluwer Academic Publishers.

Cacioppo, J.T., Petty, R.E., Kao, C.F., & Rodriguez, R. (1986). Central and peripheral routes to persuasion: An individual difference perspective. Journal of Personality and Social Psychology, 51(5), 1032–1043. http://dx.doi.org/10.1037/0022-3514.51.5.1032

Cassel, L.N., & Wolz, U. (2001). Client side personalization. In DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries (pp. 8–12). Dublin, Ireland. Retrieved from http://www.ercim.eu/publication/ws-proceedings/DelNoe02/CasselWolz.pdf

Chellappa, R.K. (2008). Consumers' trust in electronic commerce transactions: The role of perceived privacy and perceived security. Atlanta, GA: Emory University. (unpublished paper).

Chellappa, R.K., & Sin, R. (2005). Personalization versus privacy: An empirical examination of the online consumer's dilemma. Information Technology and Management, 6(2–3), 181–202. http://doi.org/10.1007/s10799-005-5879-y

Chen, J.-J.V., Huang, A.H., & Muzzerall, A. (2012). Privacy concerns and expectation of control. Human Systems Management, 31(2), 123–131. http://doi.org/10.3233/HSM-2012-0764

Cho, H., Rivera-Sánchez, M., & Lim, S.S. (2009). A multinational study on online privacy: Global concerns and local responses. New Media & Society, 11(3), 395–416. http://doi.org/10.1177/1461444808101618

Cho, H., Lee, J.-S., & Chung, S. (2010). Optimistic bias about online privacy risks: Testing the moderating effects of perceived controllability and prior experience. Computers in Human Behavior, 26(5), 987–995. http://doi.org/10.1016/j.chb.2010.02.012

Chow, R., Jin, H., Knijnenburg, B.P., & Saldamli, G. (2013). Differential Data Analysis for Recommender Systems. Proceedings of the 7th ACM Conference on Recommender Systems. (pp. 323–326). New York: ACM. http://doi.org/10.1145/2507157.2507190

Clark, J.W., Snyder, P., McCoy, D., & Kanich, C. (2015). "I saw images I didn't even know I had": Understanding user perceptions of cloud storage privacy. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (pp. 1641–1644). New York, NY, USA: ACM. http://doi.org/10.1145/2702123.2702535

Compañó, R., & Lusoli, W. (2010). The policy maker's anguish: Regulating personal data behavior between paradoxes and dilemmas. In T. Moore, D. Pym, & C. Ioannidis (Eds.), Economics of information security and privacy (pp. 169–185). New York Dordrecht Heidelberg London: Springer US. doi:10.1007/978-1-4419-6967-5_9

Culnan, M.J., & Armstrong, P.K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. Organization Science, 10(1), 104–115. http://doi.org/10.1287/orsc.10.1.104

Culnan, M.J., & Bies, R.J. (2003). Consumer privacy: Balancing economic and justice considerations. Journal of Social Issues, 59(2), 323–342. http://doi.org/10.1111/1540-4560.00067

Dinev, T., & Hart, P. (2004). Internet privacy concerns and their antecedents: Measurement validity and a regression model. Behaviour & Information Technology, 23(6), 413–422. http://doi.org/10.1080/01449290410001715723

Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. Information Systems Research, 17(1), 61–80. http://doi.org/10.1287/isre.1060.0080

Ernst&Young (2012). Voice of the customer—Time for insurers to rethink their relationships. Global Consumer Insurance Survey 2012. Retrieved from http://www.ey.com/GL/en/Industries/Financial-Services/Insurance/Global-Consumer-Insurance-Survey-2012

Fox, S., Rainie, L., Horrigan, J., Lenhart, A., Spooner, T., & Carter, C. (2000). Trust and Privacy Online: Why Americans Want to Rewrite the Rules. The Pew Internet & American Life Project. Retrieved from http://www.pewinternet.org/2000/08/20/trust-and-privacy-online/

FTC. (2012). Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers. Federal Trade Commission. Retrieved from http://www.ftc.gov/reports/protecting-consumer-privacy-era-rapid-change-recommendations-businesses-policymakers

Guo, H., Chen, J., Wu, W., & Wang, W. (2009). Personalization as a service: The architecture and a case study. In Proceedings of the First International Workshop on Cloud Data Management (pp. 1–8). New York: ACM. Retrieved from http://dl.acm.org/citation.cfm?id=1651265

Hann, I.-H., Hui, K.-L., Lee, S.-Y., & Png, I. (2007). Overcoming online information privacy concerns: An information-processing theory approach. Journal of Management Information Systems, 24(2), 13–42. http://doi.org/10.2753/MIS0742-1222240202

Ho, S.Y., & Kwok, S.H. (2002). The attraction of personalized service for users in mobile commerce: An empirical study. ACM SIGecom Exchanges—Mobile Commerce, 3(4), 10–18. http://doi.org/10.1145/844351.844354

Hsueh, P., Lin, R.J., Hsiao, M.J., Zeng, L., Ramakrishnan, S., & Chang, H. (2010). Cloud-based platform for personalization in a wellness management ecosystem: Why, what, and how. In Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2010 6th International Conference on (pp. 1–8). IEEE. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5767045

Hui, K.-L., Tan, B.C.Y., & Goh, C.-Y. (2006). Online information disclosure: Motivators and measurements. ACM Transactions on Internet Technology, 6(4), 415–441. http://doi.org/10.1145/1183463.1183467

Hui, K.-L., Teo, H.H., & Lee, S.-Y.T. (2007). The value of privacy assurance: An exploratory field experiment. MIS Quarterly, 31(1), 19–33.

Humfries, D. (2012). Smarter consumer products marketing: Understanding consumers, building brands. Somers, NY: IBM Global Business Services.

Ion, I., Sachdeva, N., Kumaraguru, P., & Čapkun, S. (2011). Home is safer than the cloud!: Privacy concerns for consumer cloud storage. In Proceedings of the Seventh Symposium on Usable Privacy and Security (pp. 13:1–13:20). Pittsburgh, PA: ACM. http://doi.org/10.1145/2078827.2078845

Ishitani, L., Almeida, V., & Wagner, M. (2003). Masks: Bringing anonymity and personalization together. IEEE Security & Privacy Magazine, 1(3), 18–23. http://doi.org/10.1109/MSECP.2003.1203218

Jalali, M., Bouyer, A., Arasteh, B., & Moloudi, M. (2013). The effect of cloud computing technology in personalization and education improvements and its challenges. Procedia-Social and Behavioral Sciences, 83, 655–658. http://dx.doi.org/10.1016/j.sbspro.2013.06.124

John, L.K., Acquisti, A., & Loewenstein, G. (2011). Strangers on a plane: Context-dependent willingness to divulge sensitive information. Journal of Consumer Research, 37(5), 858–873. http://doi.org/10.1086/656423

Joinson, A.N., Reips, U.-D., Buchanan, T., & Schofield, C.B.P. (2010). Privacy, trust, and self-disclosure online. Human–Computer Interaction, 25(1), 1–24. http://doi.org/10.1080/07370020903586662

Juels, A. (2001). Targeted advertising . . . and privacy too. In D. Naccache (Ed.), Topics in cryptology—CT-RSA 2001 (pp. 408–424). Berlin/Heidelberg: Springer. Retrieved from http://link.springer.com/chapter/10.1007/3-540-45353-9_30

Kanich, C., Checkoway, S., & Mowery, K. (2011). Putting out a HIT: Crowdsourcing malware installs. In Proceedings of the 5th USENIX Conference on Offensive Technologies (pp. 9:1–9:10). Berkeley, CA, USA: USENIX Association. Retrieved from https://www.usenix.org/conference/woot11/putting-out-hit-crowdsourcing-malware-installs

Kline, R.B. (2004). Beyond significance testing: Reforming data analysis methods in behavioral research. Washington, DC: American Psychological Association.

Knijnenburg, B.P., & Kobsa, A. (2013). Making decisions about privacy: Information disclosure in context-aware recommender systems. ACM

Transactions on Interactive Intelligent Systems, 3(3), 20:1–20:23. http://doi.org/10.1145/2499670

Knijnenburg, B.P., & Kobsa, A. (2014). Increasing sharing tendency without reducing satisfaction: Finding the best privacy-settings user interface for social networks. In ICIS 2014 Proceedings Article 11. Auckland, New Zealand. Retrieved from http://aisel.aisnet.org/icis2014/proceedings/ISSecurity/4/

Knijnenburg, B.P., Kobsa, A., & Saldamli, G. (2012). Privacy in mobile personalized systems: The effect of disclosure justifications. In Proceedings of the SOUPS 2012 Workshop on Usable Privacy & Security for Mobile Devices (pp. 11:1–11:5). Washington, DC: ACM.

Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H., & Newell, C. (2012). Explaining the user experience of recommender systems. User Modeling and User-adapted Interaction, 22(4–5), 441–504. http://doi.org/10.1007/s11257-011-9118-4

Knijnenburg, B.P., Kobsa, A., & Jin, H. (2013a). Dimensionality of information disclosure behavior. International Journal of Human-Computer Studies, 71(12), 1144–1162. http://doi.org/10.1016/j.ijhcs.2013.06.003

Knijnenburg, B.P., Kobsa, A., & Jin, H. (2013b). Preference-based location sharing: Are more privacy options really better? In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 2667–2676). Paris, France: ACM. http://doi.org/10.1145/2470654.2481369

Kobsa, A. (1990). User modeling in dialog systems: Potentials and hazards. AI & Society, 4(3), 214–240. http://doi.org/10.1007/BF01889941

Kobsa, A. (2007). Privacy-enhanced web personalization. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), The adaptive web: Methods and strategies of web personalization (Vol. 4321, pp. 628–670). Berlin/Heidelberg/New York: Springer Verlag. doi:10.1007/978-3-540-72079-9_21

Kobsa, A., & Schreck, J. (2003). Privacy through pseudonymity in user-adaptive systems. ACM Transactions on Internet Technology, 3(2), 149–183. http://doi.org/10.1145/767193.767196

Kobsa, A., & Teltzrow, M. (2005). Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing behavior. In D. Martin & A. Serjantov (Eds.), Privacy Enhancing Technologies: Fourth International Workshop, PET 2004, Toronto, Canada (Vol. 3424, pp. 329–343). Heidelberg, Germany: Springer Verlag. http://dx.doi.org/10.1007/11423409_21

Kobsa, A., Knijnenburg, B.P., & Livshits, B. (2006). Let's Do It at My Place Instead?: Attitudinal and Behavioral Study of Privacy in Client-side Personalization. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 81–90). Toronto, Canada: ACM. http://doi.org/10.1145/2556288.2557102

Komiak, S.Y.X., & Benbasat, I. (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. MIS Quarterly, 30(4), 941–960. Retrieved from http://www.jstor.org/stable/25148760

Lai, Y.-L., & Hui, K.-L. (2006). Internet Opt-In and Opt-Out: Investigating the roles of frames, defaults and privacy concerns. In Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research (pp. 253–263). Claremont, CA: ACM. http://doi.org/10.1145/1125170.1125230

LaRose, R., & Rifon, N.J. (2006). Your privacy is assured—Of being disturbed: Comparing web sites with and without privacy seals. New Media & Society, 8(6), 1009–1029.

LaRose, R., & 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), 127–149. http://doi.org/10.1111/j.1745-6606.2006.00071.x

Li, H., Sarathy, R., & Xu, H. (2010). Understanding situational online information disclosure as a privacy calculus. Journal of Computer Information Systems, 51(1), 62–71.

Li, T., & Unger, T. (2012). Willing to pay for quality personalization? Trade-off between quality and privacy. European Journal of Information Systems, 21(6), 621–642. http://doi.org/10.1057/ejis.2012.13

Li, Y. (2012). Theories in online information privacy research: A critical review and an integrated framework. Decision Support Systems, 54(1), 471–481. http://doi.org/10.1016/j.dss.2012.06.010

Li, Y. (2014). The impact of disposition to privacy, website reputation and website familiarity on information privacy concerns. Decision Support Systems, 57, 343–354. http://doi.org/10.1016/j.dss.2013.09.018

Liu, Y., Gummadi, K.P., Krishnamurthy, B., & Mislove, A. (2011). Analyzing Facebook privacy settings: User expectations vs. reality. In Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (pp. 61–70). Berlin, Germany: ACM. http://doi.org/10.1145/2068816.2068823

López-Nores, M., Blanco-Fernández, Y., & Pazos-Arias, J.J. (2013). Cloud-based personalization of new advertising and e-commerce models for video consumption. The Computer Journal, 56(5), 573–592. http://dx.doi.org/10.1093/comjnl/bxs103

Lord, K.R., Lee, M.-S., & Sauer, P.L. (1995). The combined influence hypothesis: Central and peripheral antecedents of attitude toward the ad. Journal of Advertising, 24(1), 73–85. http://dx.doi.org/10.1093/comjnl/bxs103

Lowry, P.B., Moody, G., Vance, A., Jensen, M., Jenkins, J., & 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: JASIST, 63(4), 755–776. http://doi.org/10.1002/asi.21705

Lukaszewski, K.M., Stone, D.L., & Stone-Romero, E.F. (2008). The effects of the ability to choose the type of human resources system on perceptions of invasion of privacy and system satisfaction. Journal of Business and Psychology, 23(3–4), 73–86. http://doi.org/10.1007/s10869-008-9074-0

MacKenzie, I., Meyer, C., & Noble, S. (2013). How retailers can keep up with consumers. Retrieved from http://www.mckinsey.com/insights/consumer_and_retail/how_retailers_can_keep_up_with_consumers

Madejski, M., Johnson, M., & 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) (pp. 340–345). Lugano, Switzerland: TEEE. http://doi.org/10.1109/PerComW.2012.6197507

Malhotra, N.K., Kim, S.S., & Agarwal, J. (2004). Internet Users' Information Privacy Concerns (IUIPC): The construct, the scale, and a nomological framework. Information Systems Research, 15(4), 336–355. http://doi.org/10.1287/isre.1040.0032

Mayer, R.C., Davis, J.H., & Schoorman, F.D. (1995). An integrative model of organizational trust. Academy of Management Review. Academy of Management, 20(3), 709–734. http://doi.org/10.5465/AMR.1995.9508080335

Metzger, M.J. (2004). Privacy, trust, and disclosure: Exploring barriers to electronic commerce. Journal of Computer-Mediated Communication, 9(4). http://doi.org/10.1111/j.1083-6101.2004.tb00292.x

Milne, G.R., & Gordon, M.E. (1993). Direct mail privacy-efficiency trade-offs within an implied social contract framework. Journal of Public Policy & Marketing, 12(2), 206–215. http://doi.org/10.2307/30000091

Min, J., & Kim, B. (2014). How are people enticed to disclose personal information despite privacy concerns in social network sites? The calculus between benefit and cost. Journal of the Association for Information Science and Technology, 66(4), 839–857. http://doi.org/10.1002/asi.23206

Miyazaki, A.D., & Fernandez, A. (2001). Consumer perceptions of privacy and security risks for online shopping. Journal of Consumer Affairs, 35(1), 27–44. http://doi.org/10.1111/j.1745-6606.2001.tb00101.x

Mulligan, D., & Schwartz, A. (2000). Your place or mine?: Privacy concerns and solutions for server and client-side storage of personal information. In Tenth Conference on Computers, Freedom and Privacy (pp. 81–84). Toronto, Ontario.

Newman, G.H., & Enscoe, C.J. (2000, July 4). System and method for providing client side personalization of content of web pages and the like. US patent 6085229. Retrieved from http://www.google.com/patents/US6085229

Norberg, P.A., Horne, D.R., & Horne, D.A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. Journal of Consumer Affairs, 41(1), 100–126. http://doi.org/10.1111/j.1745-6606.2006.00070.x

Patil, S., & Kobsa, A. (2009). Why is evaluating usability of privacy designs so hard? Lessons learned from a user study of PRISM. In iConference. Chappel Hill, NC. Retrieved from http://ischools.org/images/iConferences/patil-kobsa-iConference-CAMERA_READY1.pdf

Pavlou, P.A. (2011). State of the information privacy literature: Where are we now and where should we go. MIS Quarterly, 35(4), 977–988. Retrieved from http://aisel.aisnet.org/misq/vol35/iss4/10/

Petronio, S. (2002). Boundaries of privacy: Dialectics of disclosure. Albany, NY: State University of New York Press.

Petty, R.E., & Cacioppo, J.T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.), Advances in experimental social psychology (Vol. 19, pp. 123–205). Orlando, San Diego, New York, Austin, London, Montreal, Sydney, Tokyo, Toronto: Academic Press. http://dx.doi.org/10.1016/S0065-2601(08)60214-2

Petty, R.E., & Wegener, D.T. (1999). The elaboration likelihood model: Current status and controversies. In S. Chaiken & Y. Trope (Eds.), Dual-process theories in social psychology (pp. 37–72). New York, NY, US: Guilford Press.

Petty, R.E., Cacioppo, J.T., & Schumann, D. (1983). Central and peripheral routes to advertising effectiveness: The moderating role of involvement. Journal of Consumer Research, 10(2), 135–146. Retrieved from http://www.jstor.org/stable/2488919

Pirim, T., James, T., Boswell, K., Reithel, B., & Barkhi, R. (2008). An empirical investigation of an individual's perceived need for privacy and security. International Journal of Information Security and Privacy, 2(1), 42–53. http://doi.org/10.4018/jisp.2008010103

Riedl, J. (2001). Editorial: Personalization and privacy. IEEE Internet Computing, 5(6), 29–31. http://doi.org/10.1109/4236.968828

Schaefer, M. (2011). Capitalizing on the smarter consumer. Somers, NY: IBM Institute for Business Value, IBM Global Services. Retrieved from http://public.dhe.ibm.com/common/ssi/ecm/gb/en/gbe03390usen/GBE03390USEN.PDF

Schmidt, E. (2011). 36h MacTaggart Lecture. Retrieved from http://www.youtube.com/watch?v=hSzEFsfc9Ao#t=1224s

Shamdasani, P.N., Stanaland, A.J., & Tan, J. (2001). Location, location, location: Insights for advertising placement on the web. Journal of Advertising Research, 41(4), 7–21. Retrieved from http://psycnet.apa.org/psycinfo/2001-11901-001

Sheng, H., Nah, F.F.-H., & 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), Article 15. Retrieved from http://aisel.aisnet.org/jais/vol9/iss6/15

Shin, D.-H. (2010). The effects of trust, security and privacy in social networking: A security-based approach to understand the pattern of adoption. Interacting With Computers, 22(5), 428–438. http://doi.org/10.1016/j.intcom.2010.05.001

Slovic, P., Finucane, M.L., Peters, E., & MacGregor, D.G. (2004). Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. Risk Analysis, 24(2), 311–322. http://dx.doi.org/10.1111/j.0272-4332.2004.00433.x

Smith, H.J., Milberg, S.J., & Burke, S.J. (1996). Information privacy: Measuring individuals' concerns about organizational practices. MIS Quarterly, 20(2), 167–196. http://doi.org/10.2307/249477

Smith, H.J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. MIS Quarterly, 35(4), 989–1016. Retrieved from http://aisel.aisnet.org/misq/vol35/iss4/11/

Solove, D.J. (2006). A taxonomy of privacy. University of Pennsylvania Law Review, 154(3), 477–564. http://dx.doi.org/10.2307/40041279

Spiekermann, S., Grossklags, J., & Berendt, B. (2001). E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In Proceedings of the 3rd ACM Conference on Electronic Commerce (pp. 38–47). Tampa, FL: ACM.

Sundar, S.S., & Marathe, S.S. (2010). Personalization versus Customization: The importance of agency, privacy, and power usage. Human Communication Research, 36(3), 298–322. http://doi.org/10.1111/j.1468-2958.2010.01377.x

Sutanto, J., Palme, E., Tan, C.-H., & Phang, C.W. (2013). Addressing the personalization-privacy paradox: An empirical assessment from a field experiment on smartphone users. MIS Quarterly, 37(4), 1141–1164. Retrieved from http://aisel.aisnet.org/misq/vol37/iss4/9/

Taylor, D., Davis, D., & Jillapalli, R. (2009). Privacy concern and online personalization: The moderating effects of information control and compensation. Electronic Commerce Research, 9(3), 203–223. http://doi.org/10.1007/s10660-009-9036-2

Teltzrow, M., & Kobsa, A. (2004). Impacts of user privacy preferences on personalized systems: A comparative study. In C.-M. Karat, J. Blom, & J. Karat (Eds.), Designing personalized user experiences for eCommerce (pp. 315–332). Dordrecht, Netherlands: Kluwer Academic Publishers. http://dx.doi.org/10.1007/1-4020-2148-8_17

Toch, E., Wang, Y., & Cranor, L.F. (2012). Personalization and privacy: A survey of privacy risks and remedies in personalization-based systems. User Modeling and User-Adapted Interaction: The Journal of Personalization Research, 22(1–2), 203–220. http://doi.org/10.1007/s11257-011-9110-z

TRUSTe. (2014). 2014 US Consumer Confidence Privacy Report (White paper). San Francisco, CA. Retrieved from http://www.truste.com/us-consumer-confidence-index-2014/

Uusitalo, I., Karppinen, K., Juhola, A., & Savola, R. (2010). Trust and cloud services—An interview study. In 2010 IEEE Second International Conference on Cloud Computing Technology and Science (CloudCom) (pp. 712–720). Indianapolis, IN: IEEE. http://doi.org/10.1109/CloudCom.2010.41

Wang, Y., & Kobsa, A. (2008). Privacy enhancing technologies. In M. Gupta & R. Sharman (Eds.), Handbook of research on social and organizational liabilities in information security (pp. 353–376). Hershey, PA: IGI Global. http://dx.doi.org/10.4018/978-1-60566-132-2.ch013

Wilson, D., & 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. Retrieved from http://aisel.aisnet.org/icis2012/proceedings/ResearchInProgress/101

Xu, H., Teo, H.-H., Tan, B.C.Y., & 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), 135–174. http://doi.org/10.2753/MIS0742-1222260305

Xu, H., Luo, X. (Robert), Carroll, J.M., & Rosson, M.B. (2011). The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. Decision Support Systems, 51(1), 42–52. http://doi.org/10.1016/j.dss.2010.11.017

Yang, S.-C., Hung, W.-C., Sung, K., & Farn, C.-K. (2006). Investigating initial trust toward e-tailers from the elaboration likelihood model perspective. Psychology and Marketing, 23(5), 429–445. http://doi.org/10.1002/mar.20120

Zhang, Y. (1996). Responses to humorous advertising: The moderating effect of need for cognition. Journal of Advertising, 25(1), 15–32. Retrieved from http://www.tandfonline.com/doi/abs/10.1080/00913367.1996.10673493

Zhou, T. (2012). Understanding users' initial trust in mobile banking: An elaboration likelihood perspective. Computers in Human Behavior, 28(4), 1518–1525. http://doi.org/10.1016/j.chb.2012.03.021

Zimmer, J.C., Arsal, R.E., Al-Marzouq, M., & Grover, V. (2010). Investigating online information disclosure: Effects of information relevance, trust and risk. Information & Management, 47(2), 115–123. http://doi.org/10.1016/j.im.2009.12.003

Appendix: Details of the Experimental Procedures

Figure 4 shows how the survey introduced CiO to those participants who were in the American Personalization condition.

FIG. 4. Introductory screen in the American Personalization condition.

FIG. 5. Screenshot of the Check-it-Out Android app. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

The three subsequent examples of purported personalized services of Check-it-Out are the following:

1. CiO points out an upcoming U2 concert, because the user played their music and chatted with friends about them.
2. CiO points out a Sears promotion for appliances, because the user searched for dishwashers on the web.
3. CiO recommends a friend of a friend who is interested in Salsa dancing, because the user searched for a Salsa class online and bought a book on that topic.
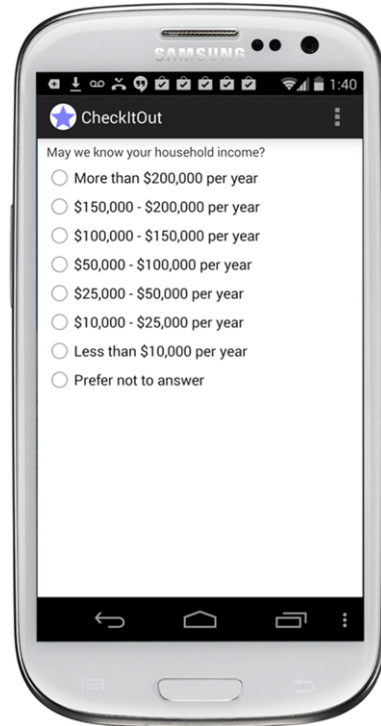
Figure 5 shows the CiO app that asks demographic and context questions. Participants can answer a demographic request by selecting an answer from a drop down list and pressing the OK button, and give permission to a context request by simply pressing the OK button. They are also free to decline the answer or permission for any request by selecting the "Prefer not to answer" option or by pressing the "Rather not" button, respectively.