Investigating Privacy Concerns Related to Mobile Augmented Reality Apps – A Vignette Based Online Experiment

Abstract:

Augmented reality (AR) gained much public attention after the success of Pokémon Go in 2016, and has found application in online games, social media, interior design, and other services since then. AR is highly dependent on various different sensors gathering real time context-specific personal information about the users causing more severe and new privacy threats compared to other technologies. These threats have to be investigated as long as AR is still shapeable in order to ensure users' privacy and foster market adoption of privacy-friendly AR systems.

To provide viable recommendations regarding the design of privacy-friendly AR systems, we follow a user-centric approach and investigate the role and causes of privacy concerns within the context of mobile AR (MAR) apps. We design a vignette-based online experiment adapting ideas from the framework of contextual integrity to analyze drivers of privacy concerns related to MAR apps, such as characteristics of permissions, trust-evoking signals, and AR-related contextual factors. The results of the large-scale experiment with 1,100 participants indicate that privacy concerns are mainly determined by the sensitivity of app permissions (i.e., whether sensitive resources on the smartphone are accessed) and the number of prior app downloads. Furthermore, we devise detailed practical and theoretical implications for developers, regulatory authorities and future research.

Keywords: Mobile augmented reality, pervasive systems, privacy, vignette-based experiment, app transparency, app permissions

Investigating Privacy Concerns Related to Mobile Augmented Reality Apps - A Vignette Based Online Experiment

1. INTRODUCTION

The release of Pokémon Go in 2016 increased public awareness about augmented reality (AR) (Nicas & Zakrzewski, 2016). Big technology companies engage heavily in acquisitions of AR companies (10.5 billion and 18.8 billion dollar in investment in 2019 and 2020, respectively (Rossolillo, 2020)) and there are projections which state that there will be 83.1 million people in 2020 in the US who use AR at least once a month on any kind of device (Petrock, 2020). AR is defined as a technology which "[...] combines real and virtual objects in a real environment; runs interactively, and in real time; and registers (aligns) real and virtual objects with each other" (Azuma et al., 2001, p. 34).

The two main types of AR are considered to be smart glasses and mobile AR (MAR) apps. AR glasses like the Microsoft Hololens (Microsoft, 2017) are currently not mature enough (regarding the size, price and usability) for the end consumer market. This type of AR is primarily used in the business-to-business (B2B) environment in which AR successfully showed that it can save time and costs (Kohn & Harborth, 2018). Another way of presenting AR to the user is via smartphones or tablets (MAR). Pokémon Go is the most widely known example for this category. When Apple (ARKit) and Google (ARCore) released AR development kits in 2017, AR features like object tracking started to become better (Nellis, 2017) and many new MAR applications (apps) diffused into the consumer market which is why we analyze this type of AR technology in this article.

Several privacy issues related to MAR emerged in the past. In order to reach their full potential, MAR apps require massive amounts of data from a variety of sensors, i.e., access to the smartphone camera. Users of MAR apps with such capabilities are exposed to more severe and new types of privacy risks compared to the ones related to regular (non-MAR) smartphones

apps. Literature suggests five major risks which distinguish MAR apps from non-MAR apps (de Guzman et al., 2018; Harborth et al., 2019):

1. Risks which relate to the MAR app input due to limited feedback regarding what data is captured by the app's camera. Problems arise when privacy-sensitive information about the user itself is gathered aside from fulfilling the app's primary task (the user cannot know what is captured in which context when using the app). Often users are also not aware which information (e.g., location or other persons) could be inferred from their camera.
2. Risks which relate to the MAR app output due to malicious apps which could alter the digital objects and information presented to the user.
3. Increasing data aggregation capabilities of MAR apps due to a simultaneous employment of multiple privacy-sensitive sensors (e.g., location, camera, accelerometer data, etc.). The main issue evolves not primarily because of a single dangerous permission associated with AR, like the camera, but because of the opacity of potential privacy risks when combining and analyzing these different data types of multiple permissions of MAR apps.
4. Risks which relate to privacy breaches in collaborative and shared AR environments when two or more users work with separate AR devices on the same digital objects (Lebeck et al., 2018). It is required to ensure that collaborative spaces and the respective digital information of users are protected against third-party attacks.
5. Risks for bystanders of AR systems who are in the field of view and get filmed by the systems without awareness or possibility to control (Denning et al., 2014).

Privacy issues arise with MAR apps (e.g., for Pokémon Go, see Peterson, 2016) and research indicates that individuals are concerned about their privacy when using AR. These concerns are about being filmed by AR devices (as bystanders), distributing data involuntarily and being surveilled due to using the devices (Dacko, 2017; Harborth, 2019; Harborth & Pape, 2018; Rauschnabel et al., 2018). However, there is still a lack in current research which investigates

how these privacy concerns of end users influence the use of MAR (Harborth, 2017). Addressing this research gap is especially important since context-specific privacy concerns can differ greatly from general privacy concerns and MAR represents a new technology with the aforementioned risks (Ackerman & Mainwaring, 2005; Acquisti et al., 2015; Nissenbaum, 2010). Thus, it is crucial for research to provide insights for a privacy-friendly design and potential regulation of such systems since there are several useful AR applications which could provide much value to users (e.g., for medical purposes like helping Parkinson's disease patients, see McNaney et al., 2014; van der Meulen et al., 2016).

We develop a research model tailored to MAR – a technology relying on highly contextualized information about the end user – based on the so-called "antecedents – privacy concerns – outcomes" (APCO) model and the framework of contextual integrity (CI) (Nissenbaum, 2010; Smith et al., 2011). We base our research on the APCO model since it is derived from a large basis of privacy research and it shows the general relations of privacy concerns with different antecedents and outcome variables. We merge it with the framework of CI, since this framework provides detailed insights into the contextual nature of privacy. CI aims at explaining data sharing decisions of individuals by assessing the perceptions of individuals about the appropriateness of a respective data flow (Nissenbaum, 2010). However, our research has a different goal compared to CI. We do not want to assess whether individuals perceive a data flow as appropriate, but rather what factors contribute to their privacy concerns in the first place within the context of a specific MAR app download scenario. CI provides a starting point, rather than a complete theory to rely on, to systematically approach the concept of privacy as a context-dependent construct and analyze the different factors which influence users' privacy expectations and concerns. Due to this different goal, we merge it with the APCO model and augment the antecedents emerging from CI with relevant factors which were found to influence privacy concerns in the mobile app context in prior literature like the price of an app (Bamberger et al., 2020) and AR-related informational cues.

Based on this, we investigate two research questions:

RQ1. What factors contribute to users' privacy concerns with respect to MAR apps?

RQ2. How do privacy concerns affect the intentions to download MAR apps?

We deliberately focus on download intentions instead of other known target variables such as use intentions (Kim et al., 2016) since the initial decision to download an app is associated with far-reaching implications for users' privacy since they oftentimes need to grant permission (i.e., allow the app to access certain resources of the device) either after installing the app or during the runtime phase. Furthermore, the number of downloads of an app is an economically relevant number for app developers, app operators and app stores since it determines a certain share of their profit (besides in-app purchases) (Gu et al., 2017). Therefore, it is important to investigate the factors which influence privacy concerns and, in turn, affect the intentions to download MAR apps in order to derive managerial recommendations to address users' privacy concerns and increase potential downloads. Due to this importance, download intentions are a common endogenous variable studied in the literature (Kang et al., 2015; Wu et al., 2016).

We use a vignette-based design to present 1,100 participants a highly specified context (in our case the context of downloading a hypothetical MAR app represented by a high-fidelity mockup of an app store website) in order to evaluate our structural equation model.

Our results contribute to the current body of knowledge on context-dependency of privacy concerns in general, users' privacy perceptions in the app ecosystem as well as privacy issues related to MAR apps and technologies. From a practical point of view, we derive actionable insights aiming to protect users' privacy which can be considered by MAR developers and policy makers alike.

2. USER PRIVACY IN THE SMARTPHONE ECOSYSTEM

There is a plethora of research on privacy threats and concerns related to the smartphone ecosystem. Research can be categorized along the user journey related to mobile apps (Figure

1). In the first stage, the user has to download the app. Before that, users need to get aware of the app and decide if they want to install it. In the second stage, the user installs the app. Although, the relevant decision process was already finished before, users can still decide to cancel the installation. Several types of static information, like the phone number or the IMEI, may be already transferred to the app when the installation is completed. The last stage is the actual use of the app in which further information may be transferred which allows apps to profile users. Users may decide to uninstall the app, which will stop further profiling, but cannot withdraw the leakage of static information. Due to the economic relevance of download intentions (see Section 1), we locate our research in the first stage, i.e., analyzing factors driving privacy concerns and download intentions in the download stage.

Research analyzing the download stage is relatively rare compared to the other stages (Gu et al., 2017). Findings from prior work show that users' privacy concerns in the download stage are alleviated by the popularity of the app and by the existence of permission justifications (explaining users why apps need certain permissions), whereas privacy concerns increase if apps require more sensitive permissions (Gu et al., 2017). Other research shows that the demand is lower for apps with sensitive permissions for a given price and functionality. However, the strength of this relationship depends on contextual factors (Kummer & Schulte, 2019). Since these variables are important contextual concepts, we adapt them at a later stage for our study.

Figure 1

App Usage Lifecycle

Other studies investigate the effect of presenting permissions clearly to the users in the download stage. They find that making permissions of an app clear and apparent helps users become aware of these permissions and they show that users would like to better understand why applications need certain information, whereas permission justifications were not included in these study designs (Kelley et al., 2013). In addition, research shows that privacy notices shown at the download stage are not as effective as shown during the use of the app (Balebako et al., 2015). When considering permissions and justifications it is relevant to note that it is hard for layman users to identify the reason an app uses a specific resource. However, users' expectations and the purpose why sensitive resources are used have a major impact on users' trust (Lin et al., 2012).

2.1 Privacy Concerns Related to AR

There is a large body of technical research about privacy and security in augmented reality technologies (de Guzman et al., 2018). However, there is little research on end users' perceptions and privacy concerns regarding AR technologies (Harborth, 2017). The limited empirical evidence which exists, suggests that AR raises privacy concerns among users, for instance, about being filmed by AR devices as bystanders (Denning et al., 2014), distributing data involuntarily and being surveilled due to using the devices (Dacko, 2017; Harborth, 2019; Rauschnabel et al., 2018). However, none of this research investigates the drivers of privacy concerns regarding AR, but rather uses privacy concerns as an antecedent to explain other phenomena (e.g., explaining AR use behavior). In addition, we could see that context-specific privacy concerns can differ greatly from general privacy concerns and privacy concerns in the context of other technologies (Acquisti et al., 2015). Thus, it is required to specifically investigate MAR and derive a tailored model which aims at explaining the contextual factors which influence privacy concerns regarding MAR apps. In this study, we attempt to close this gap by developing and evaluating such a model.

3. RESEARCH METHOD AND RESEARCH MODEL

There are several models and concepts describing antecedents and outcome variables related to privacy concerns. The APCO model is one of the most established models showing the general relations between antecedents such as demographic variables, privacy concerns and outcome variables (e.g., use intentions) (Smith et al., 2011). We use the APCO model as an overarching model to develop potential antecedents of privacy concerns as well as the outcome variables such as trust and download intentions as well as the role of the privacy calculus. However, the APCO model only includes antecedents which are not specific to a technology. In addition, we argue that previous literature on AR and privacy has rarely, if, indeed, at all, covered the multitude of factors influencing privacy concerns in the context of MAR apps (see Section 2.1). Thus, we augmented the general APCO model with the framework of contextual integrity (CI) (Nissenbaum, 2010) and, thereby, introduce contextually relevant antecedents to explain privacy concerns related to MAR apps.

3.1 Framework of Contextual Integrity

CI aims at providing a systematic account for understanding user expectations regarding privacy and perceived violations of individuals' (Nissenbaum, 2010). A privacy violation occurs when the information practice does not correspond to users' expectations in a given context. CI theorizes that privacy is evaluated by individuals for each specific context in which they are confronted with respective privacy-related decisions. Each context is guided by certain informational norms which can be seen as “[...] juggling balls in the air, moving in sync: subjects, senders, receivers, information types, and transmission principles” (Nissenbaum, 2010, p. 145). Thus, one needs to define each of these parameters for a specific context when evaluating privacy perceptions of individuals. However, the framework of CI has several shortcomings when it comes to implementing the theoretical aspects to a practically applicable research model. In addition, as elaborated before, the goal of the theoretical framework is another one

than ours. CI has the goal to assess the perceived appropriateness of personal information flows. In contrast, our study investigates the relevant factors influencing the privacy concerns of users and consequent download intentions.

Empirical studies applying CI are rare since it is a rather theoretical and abstract framework. Several articles investigating CI use simplifying assumptions to bypass the difficulty of working out a fully elaborated norm for the specific research context. Such simplifications include the assumption that norms are solely defined and based on the type of information (Barth, Datta, Mitchell, & Nissenbaum, 2006) or the restriction of the analysis to a subset of the original factors like type of information, receivers fitted to specific contexts, and use of the information (Martin & Nissenbaum, 2016).

Most research applying the framework of CI is located in the computer science discipline. For example, research on user expectations regarding smartphone permissions finds that 80% of the participants are confronted with at least one permissions request which they perceive as inappropriate, thus, violating their expectations regarding what the app should be allowed to do (Wijesekera et al., 2015). Other results show that users focus more on how and for what (purpose) the data is used than on what actual data is flowing (Martin & Nissenbaum, 2016). Furthermore, users care to whom (receiver) data flows (Martin & Nissenbaum, 2019).

3.2 Operationalization of Contextual Integrity and Treatment Variables

For developing the research model of our study, we need to start by operationalizing the factors which influence users' privacy concerns about MAR apps according to the framework of CI. The information flow within the context of a hypothetical MAR app download scenario is – to the best of our knowledge – not defined in past research. The same holds for a co-constitutive norm governing the smartphone ecosystem and especially the download of an app. Prior work manipulates *context* in a way that it differentiates between a contextual use and, for example, a commercial use (Martin & Nissenbaum, 2016). We assume that the download of an MAR app is

always a commercial context since users are most likely in a situation in which their data is gathered and processed for commercial purposes. This assumption is underlined by prior work finding that it does not even play a role for users' privacy whether apps cost money or whether they are free. Both types of apps gather similar amounts and types of personal information (Han et al., 2019).

In addition, we could see a general shift in the economic imperative of technology companies in the last decade from the “behavioral value reinvestment cycle”, trying to gather and analyze data to improve services, to the idea of “behavioral surplus” where the same data are primarily used as “raw material” to generate new products based on user data (Zuboff, 2019). Based on this logic, the user in the smartphone ecosystem is most likely, knowingly or not, part of this new paradigm. Consequently, context in our research model is held constant and we assume a commercial context in which apps gather personal information as a primary source of revenue. The *subject and the sender* are the same; in our case the user. The *receiver* is the app, i.e., the app developer or operator. All actors are held constant in our treatments.

However, we manipulate the factor *app popularity* which specifically influences *trust in the app*, i.e., the receiver of the information flow. This treatment is manipulated by the number of downloads and the ranking of an app (in the app's respective category, in our case utility). The treatment is either 300,000 downloads and rank 2 of 100 or 800 downloads and rank 90 of 100. This choice of download numbers represents the idea that popular apps are perceived as trustworthier compared to unpopular apps (Duan et al., 2009). Thus, we chose these extremely opposing values as treatments for this variable. Furthermore, our pretest as well as prior work shows that such a manipulation causes the desired effect (Gu et al., 2017).

We decided against introducing a manipulation of a star-based rating since this would have introduced another dimension which is directly related to the quality of the app (besides the pure number of reviews in this case). This would have led to the necessity of introducing another

treatment dimension resulting in 64 groups to evaluate. Thus, we decided against this due to limited financial resources.

The *information type* is manipulated by altering the permissions the app requests with the variable *permission sensitivity*. The treatment compares the effects of including a set of three "normal" permissions vs. a set of three "normal" and three "dangerous" permissions. We differentiate these permissions according to the Android developer guide about different permission types (Android Developers, 2019). Exemplary normal permissions request access to the storage of the smartphone or the network. In contrast, dangerous permissions request access to the camera, the contacts or the microphone of a smartphone. Relevant to note is that since our hypothetical app is an MAR app way, we included the camera permission in every manifestation. We also analyzed to what extent our chosen permissions are requested by real-life MAR applications. We collected the data in the permission manifests of 198 MAR apps and compared this with 25,975 non-MAR (NMAR) apps. We included all MAR apps from the Google Play Store that we could find in through the Play Store search with the keyword "Augmented Reality". As expected, all MAR apps require camera access and a larger share of MAR apps need access to the device's USB storage. The differences between MAR and NMAR apps are smaller for other permissions (see Figure 3 in Appendix B).

The *transmission principle* aspect represents a broad variety of possible factors (Nissenbaum, 2010). We manipulated the transmission principle by introducing *permission justification* as a way to increase transparency and provide a valid purpose for collecting data of users. By that, we manipulate the "condition[s] under which such transfers [of information] ought (or ought not) to occur" (Nissenbaum 2010, p. 145). The respective treatment tests the effect of including a permission request justification in the vignette scenario, aiming at increasing transparency, compared to the scenario in which no such justification is provided. Transparent information about privacy can help users to select less problematic apps in an app download context (e.g., Bal (2014). In addition, prior research shows that the mere presence of a text that resembles a

valid justification but does not have actual explanatory connotation, may affect users' preference (Tan et al., 2014). In contrast, prior work on CI actually shows that an explanation for certain information flows (use of data in a specific context) does result in high degrees of individuals' privacy expectations being met. Additionally, it is shown that users focus more on how and for what the data is used than on what actual data is flowing (Martin & Nissenbaum, 2016). This indicates that a permission justification serves as a valid indicator providing a reason why data are collected by the MAR app. However, since prior results are ambiguous, we implicitly assume for our design that such justifications will be subject to close scrutiny by the app store providers (i.e., Google or Apple) in order to ensure correctness of the justifications. Such scrutiny is partially already in place since Google checks whether apps request only permissions which are relevant for the proper functionality of the app (Google Play Developer Policy Center, 2020). Our assumption regarding the justification assessment can therefore be viewed as realistic considering such processes being already in place.

We introduce two additional contextual variables based on prior empirical findings and AR-related contextual information which are not covered by the framework of CI. We manipulate the *price of the app* (free versus lump-sum cost) in order to compare the effects of the app's price on user perceptions. Participants either see the MAR app as free to download or a lump sum cost of $6.99. The second treatment, *AR label*, compares an app which is clearly labeled and described as an AR app versus an app with no such descriptions and labels. We discuss the reasons for including these variables in the next section.

3.3 Experimental Treatments and Related Research Hypotheses

Our model consists of three main parts which are derived from the APCO model. First, we have the contextual variables (antecedents) which influence privacy concerns. Second, we have the outcome variables like trust, download intentions and the variables related to the privacy calculus which we hypothesize to be in place for our model (please note that we consider trust

to be an outcome variable in the terminology of the APCO model although trust influences privacy concerns in our model). The privacy calculus is a core element in the APCO model and shows how risks (in the form of privacy concerns) and benefits are weighed up by individuals to make privacy-related decision (Dinev & Hart, 2006). The privacy calculus determines download intentions in our research model.

App popularity represents a heuristic for individuals indicating that a product or service is trustworthy due to many prior adopters (Duan, Gu, & Whinton, 2009). This information serves as a trust cue and therefore, increases users' trust in the app, due to belief that many other users tried it and had no obvious issues (Duan et al., 2009):

H1a: Perceived app popularity positively influences the trust in the app.

Furthermore, *download intentions* are influenced by a trade-off between benefits and costs which individuals face (privacy calculus) (Dinev & Hart, 2006). Besides being an indicator of trustworthiness, *app popularity* can also indicate product attractiveness (Duan et al., 2009) and, therefore, can be seen as a benefit of downloading an app:

H1b: Perceived app popularity positively influences the intention to download the app.

App price affects users' privacy perceptions in a way that users who buy an app think that the app developer or operator generates revenues with their monetary payment and not through selling their personal data. This is in contrast to free apps, for which users have the expectation that the use is most likely paid via their personal data (Bamberger et al., 2020). These expectations are problematic since prior work finds that both types of apps, free and paid ones, collect similar amounts and types of data (Han et al., 2019). Therefore, we include this treatment since it is associated with relevant expectations of users in the context of the download phase and hypothesize:

H2a: Users are more concerned about privacy in free MAR apps than in paid MAR apps.

In addition, *app price* is assumed to have a direct effect on *download intentions* since mobile app users prefer free apps over paid ones:

H2b: Users are more likely to download free MAR apps than paid MAR apps.

Permission justifications were designed in a way that they explain even the dangerous and unnecessary permission requests of the MAR app. For example, one could argue that a measurement app does not need access to the microphone. Here, we developed a justification stating that users can add audio notes to each measurement file and, therefore, easily annotate it. Thus, we hypothesize that such explanations increase users' *trust in the MAR app* and alleviate the *privacy concerns* related to the app. This is in line with prior results suggesting that users' expectations and the purpose of why sensitive resources are accessed influence their trust (Lin et al., 2012):

H3a: Permission justifications positively influence trust in the app.

H3b: Permission justifications negatively influence privacy concerns related to the app.

Permission sensitivity is particularly important since layman users have a hard time identifying the reason an app accesses a resource. Thus, the information sensitivity may concern users independently of the given justification (Lin et al., 2012). The idea of including this concept stems from past research indicating that privacy concerns are influenced by information sensitivity (Bansal et al., 2010). Furthermore, recent research finds that users concerns about app permissions have a strong positive influence on mobile users' information privacy concerns (Degirmenci, 2020). Thus, we hypothesize:

H4: Perceived permission sensitivity positively influences privacy concerns related to the app.

The ***AR label*** treatment is supposed to uncover certain pre-existing privacy concerns or other (negative) attitudes towards AR. Prior work on AR shows that individuals have such pre-existing attitudes (Harborth & Kreuz, 2020) and privacy concerns regarding AR (Dacko, 2017; Harborth, 2019). Thus, it is important to include such an informational cue related to AR in order to account for the fact that certain participants might not be directly aware about the app being an

MAR app as well as to evaluate whether the mere mentioning of AR has any effect on their privacy concerns. Therefore, we hypothesize:

H5a: The AR label of the MAR app positively influences privacy concerns related to the app.

Furthermore, we argue that the fact that it is clearly labeled as "augmented reality" might positively influence certain users regarding their download intentions since AR might be seen as a new and innovative feature which promises added value. Thus, we hypothesize:

H5b: The AR label of the MAR app positively influences the intention to download the app.

3.4 Further Determinants of Privacy Concerns and Download Intentions

Trust is one of variables which are found to be relevant within the context of the APCO model (Smith et al., 2011, p. 998). However, the associated literature review shows that research results are inconclusive with respect to the direction of the relationship between trust and privacy concerns. Based on prior research, we argue that trust in a service or technology can alleviate associated risks and privacy concerns (Pavlou, 2003) as well as positively influence behavioral intentions (McKnight et al., 2011). Thus, we hypothesize:

H6a: Trust in the MAR app negatively influences privacy concerns related to the app.

H6b: Trust in the MAR app positively influences the intention to download the app.

Based on the privacy calculus (Dinev & Hart, 2006), we hypothesize that users weigh up risks and benefits when making the choice to use MAR apps. Privacy risks can be seen as the associated costs with a negative influence within the privacy calculus (Keith et al. 2013). Thus, such risks and related concerns negatively affect the download intentions:

H7: Privacy concerns related to the MAR app negatively influence the intention to download the app.

Benefits positively impact the calculus. The respective variables in our case are *perceived usefulness*, *app popularity* (see H1b) and *AR label* (see H5b) as possible benefits of using the MAR app. Findings related to technology acceptance indicate that *perceived usefulness* is an

important driver of intentions to use technologies (Venkatesh et al., 2012) and therefore, constitutes an important benefit to consider in the analysis. Thus, we hypothesize:

H8: The perceived usefulness of the MAR app positively influences the intention to download the app.

3.5 Hypotheses Related to the Control Variables

We control for common demographic variables and privacy victim experiences which were shown to have an influence on privacy concerns in the APCO model (Smith et al., 2011) as well as variables reflecting experience with technologies. We decided to control for the effects of these variables on trust in the MAR app as well. Furthermore, we add AR-specific variables as the *number of installed MAR apps*, *attitudes towards AR in general* or *AR knowledge*. We use this set of variables for all three endogenous variables and add one specific non-context-specific control variable for privacy concerns and trust, respectively. We control for *general privacy concerns related to MAR apps* when assessing the impact of the treatment variables on the privacy concerns related to the specific MAR app from the mockup. Furthermore, research indicates that there are differences in institutional trust beliefs in the internet between more and less concerned internet users (Martin & Nissenbaum, 2016). Thus, we control for this potentially confounding factor influencing users' *trust in the MAR app* by introducing the general construct *trust in the app store* as a control variable to account for these differences in institutional trust beliefs. We also checked whether there are statistically significant differences in the trust levels in the app store between iOS (240 participants) and Android (841 participants) users since prior work finds that Apple users trust their devices more compared to users of other devices (Frik et al., 2019). However, we did not find a statistically significant difference in the levels of *trust in the app store* between both groups (the remaining 19 participants stated to use another mobile operating system). Figure 2 shows the resulting research model with the hypotheses.

Figure 2

Research Model (paid app=1, permission justification given=1, AR label given=1)

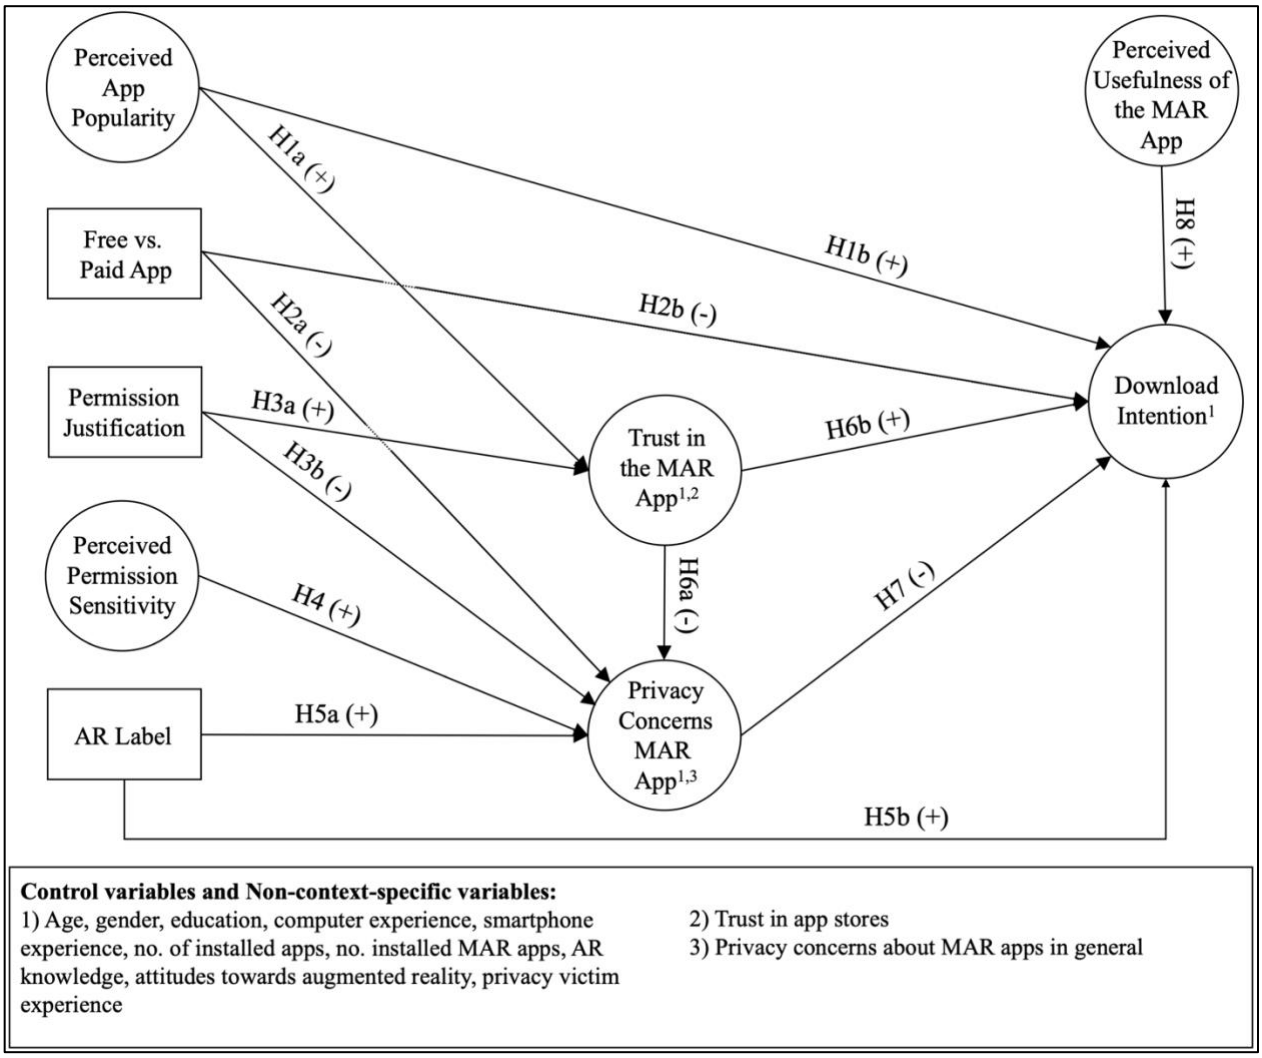

3.6 Experiment Design

We consequently address our research hypotheses with a 2x2x2x2x2 between-subjects vignette-based online experiment with the variables *app popularity*, *app price*, *permission justification*, *permission sensitivity* and *AR label* as described in Section 3.3. Participants were randomly assigned uniformly to one of the 32 different vignettes containing the respective combinations of information about a hypothetical MAR measurement app which allows users to

measure distances between objects in the real environment. This category of apps is popular (e.g., apps like “CamToPlan” or “AR Ruler” app have more than 1,000,000 downloads each) and has reliable performance (e.g., Apple integrated such an app in iOS12 (Apple, 2019)). Participants easily understand the use case scenarios of such an app, and it is likely to be used in potentially privacy sensitive environments capturing users’ homes and family members. This is in contrast to other potential hypothetical MAR apps like games which have much more complex operating principles. After seeing one mockup (examples in Appendix B), participants continued by answering the survey instrument. All questions for the research model were randomized. We approached the ethics board of the university and our study was judged to be exempt of a detailed ethics review since we did not collect and save any personal information and we did not expose the participants to false or misleading situations.

3.7 Data

The constructs of the questionnaire are adapted from previous literature except for the construct *privacy concerns about MAR apps in general* which is based on privacy concerns which were found in previous literature on AR and privacy (Harborth, 2019; Rauschnabel et al., 2018). A list of all items can be found in Appendix A. Since we conducted the study with a German panel, the items needed to be translated into German. To ensure the validity and reliability of the translated constructs, we conducted a pretest with the German version of the questionnaire with 91 participants. We collected the data with the help of a market research institute certified following ISO 26362 in order to get a high-quality data set.

For the number of participants for our 2x2x2x2x2 between-subjects vignette-based online experiment we considered the sample size of multifactorial analysis of variances according to Döring and Bortz (2016, p. 846): The number of necessary participants for each cell is calculated by (n-1)(df+1) / <number of cells> +1 where n is looked up in a table (Döring & Bortz, 2016, p. 844) and depends on the degree of freedom, the potential effect size and the desired

significance level, and df is the degree of freedom. Since each dimension is binary, the degree of freedom (df) for each of the cells is 1. Looking up n for df=1, small effect sizes and a significance level of .05 results in n=393. Therefore, the number of necessary participants for each cell is :(393-1)(1+1) / 32 + 1 = 25.50

With 32 cells and 25.50 participants per cell, 816 participants were needed. In order to have a margin for faulty answers and to be able to exclude answers should we spot flaws in the answers, we planned with 1,100 participants, which is roughly 32 participants per cell.

5,566 participants started the online survey and 1,100 remained after filtering out participants who answered either one or more of five test questions about the mockup or one attention question incorrectly. The descriptive statistics for the final sample with 1,100 participants are shown in Table 1. The median age of the participants in our sample is 43 years, the gender distribution is almost uniform and the median level of education is equal to the A levels degree (qualification for a Bachelor's program). The median experience with personal computers and smartphones is 21 years and 8 years, respectively. The median number of apps and AR apps is 22 and 0, respectively. We also tested the participants' knowledge about AR by asking a multiple-choice question about the correct definition of AR. This question was correctly answered by 682 participants (62%). We will discuss implications of this result in Section 4.

Table 1

Descriptive Statistics for the Demographics (N=1,100)

Variable \ Statistics	Mean	Median	Std. Dev.	Minimum	Maximum
Age	42.985	43	12.213	18	66
Gender (0=female, 1=male)	0.499	0	0.500	0	1
Degree (1=no degree, 7=PhD)	4.065	4	1.289	1	7
Experience PC	18.290	21	3.992	2	> 20 years
Experience smartphone	8.108	8	2.620	0	> 10 years
Number of apps	36.482	22	41.605	0	365
Number of AR apps	0.297	0	1.057	0	20

The number of participants for each group is shown in Table 6 (Appendix C). The treatments (*app popularity* and *permission sensitivity* are included in the model as latent variables. The manipulation checks show that both manipulations yield the desired outcome by only significantly impacting the respective latent constructs. We use the mean sum scores of the latent variables to calculate the descriptive statistics and differences between treatment groups (Table 2). We will discuss further information shown in this table in the robustness section.

Table 2

Descriptive Statistics of the Latent Variables

Variables	Treatment group	N	Mean	Median	Std. Dev.
	Treatment variables				
PAP	High popularity	548	5.004	5	1.245
	Low popularity	552	2.986	3	1.476
PPS	Sensitive Permissions	548	5.794	6	1.202
	Less Sensitive Permissions	552	4.513	4.667	1.430
	Context-specific variables				
PC	PC (full sample)	1,100	4.647	4.875	1.618
	Sensitive Permissions	548	5.073	5.25	1.578
	Less Sensitive Permissions	552	4.223	4.25	1.545
TRUST	TRUST (full sample)	1,100	4.113	4	1.279
	High popularity	548	4.110	4	1.330
	Low popularity	552	4.117	4	1.226
DI	DI (full sample)	1,100	3.651	4	1.770
	High popularity	548	3.707	4	1.788
	Low popularity	552	3.595	4	1.752
DI	Cost: free	548	3.844	4	1.787
	Cost: lump-sum	552	3.460	3.667	1.735
DI	AR Label	548	3.591	4	1.808
	No AR Label	552	3.710	4	1.732
PU		1,100	3.961	4.25	1.718
	Non-context-specific variables				
VIC		1,100	2.115	2	1.302
ATT		1,100	4.843	5	1.372
PC_{MAR}		1,100	4.87	5	1.345
$TRUST_{AS}$		1,100	4.281	4	1.263

4. DATA ANALYSIS

Since our research is exploratory with respect to the development of a new model to predict the target construct *download intentions* while maximizing the explained variance, we use partial least squares structural equation modeling (PLS-SEM) for our analysis (Hair et al., 2011, 2017). We created one structural equation model including the whole sample. We tested our research model using SmartPLS version 3.2.8 (Ringle et al., 2015). We first discuss the measurement model and check for reliability and validity of our results. For the PLS algorithm, we chose the path weighting scheme with a maximum of 300 iterations and a stop criterion of 10^{-7}. We used 5,000 bootstrap subsamples and no sign changes as the method for handling sign changes for the bootstrapping iterations.

4.1 Measurement Model Assessment

For the reflective measurement model, we evaluate the internal consistency reliability (ICR), convergent validity and discriminant validity (Hair et al., 2017). The values for Cronbach's alpha and the composite reliability are all well above 0.7 (Appendix D, Table 7). Convergent validity (based on the assessment of outer loadings and the AVE) is given since all loadings are higher than 0.7. The AVE values of the constructs are also well above 0.5, demonstrating convergent validity. Discriminant validity is also given (all outer loadings of our analyzed constructs are larger than their cross-loadings with other constructs) (Appendix D, Table 7). Second, the square roots of the AVEs of all single constructs are larger than the correlation with other constructs (Fornell-Larcker criterion) (Appendix D, Table 8). Third, we used the HTMT (heterotrait-monotrait ratio of correlations) criterion to assess discriminant validity. Discriminant validity can be assumed if the HTMT value is below 0.85 (Henseler et al., 2015) which is given for our model (Appendix D, Table 9). We also evaluate whether the HTMT statistics are significantly different from 1 based on a bootstrapping procedure with 5,000 subsamples. No single 95% confidence interval for two constructs contains the value 1 in our model. Thus,

discriminant validity is established for our model. We test for common method bias (CMB) since our data was gathered with a self-reported survey at one point in time in one questionnaire. An unrotated principal component factor analysis is performed to conduct a Harman's single-factor (Podsakoff et al., 2003). The test shows that twelve factors have eigenvalues larger than 1 which account for 77.67% of the total variance. The first factor explains 29.75% of the total variance. Based on these values, we argue that CMB is not likely to be an issue in the data set.

4.2 Structural Model Assessment

Collinearity is present if two predictor variables are highly correlated with each other. We assess the inner variance inflation factor (VIF) which should not be above 5. For our model, the highest VIF is 2.196. Thus, collinearity is not an issue. Table 3 shows the results of the path estimates and the R^2-values of the endogenous variables TRUST, PC and DI. The R^2-values for PC and DI are excellent with 70.2% and 61.2%, respectively. The adjusted R^2-value for TRUST is medium-sized with 38.1%. Effect sizes of path coefficients are interpreted relative to each other within the nomological net. Values of f^2 show the impact of a construct on the explained variance for an endogenous variable by omitting it from the analysis and assessing the resulting change in the R^2-value. The values are assessed based on thresholds by Cohen (1988), who defines effects as small, medium and large for values of 0.02, 0.15 and 0.35, respectively. Furthermore, we calculate the predictive relevance Q^2, which indicates the out-of-sample predictive relevance of the structural model with regard to the endogenous latent variables based on a blindfolding procedure with an omission distance d equal to seven (sample size divided by d should not be an integer which is avoided in our case of N=1,100 and d=7) (Hair et al., 2017). The reported Q^2-values are based on the cross-validated redundancy approach, since this approach is based on both, the results of the measurement model as well as of the structural model (Chin, 1998). Q^2-values above 0 indicate that the model has the property of predictive relevance. In our case, the Q^2-values for PC, TRUST and DI are equal to 0.624,

0.333 and 0.546, respectively. Thus, predictive relevance is established. Values of q^2 are calculated by deleting the respective relation between the exogenous and endogenous variable, while keeping the latent variable in cases of multiple relations of that exogenous variables to other endogenous variables (e.g., ATT influences all three endogenous variables). q^2 shows the predictive power of the respective exogenous variables by omitting their relation to the endogenous variable and comparing the change in the values of Q^2 for the respective endogenous variable (Hair et al., 2017). Table 4 shows the indirect effects and the total effects on the endogenous variables privacy concerns and download intentions. We find that all shown context-specific and non-context-specific variables have a statistically significant effect on *privacy concerns* and *download intentions* except for *permission justifications* as well as *permission justifications* and *prior privacy victim experience*, respectively.

Table 3

*Results of the Structural Model (***, **, * Asterisks Indicate Statistical Significance at the 0.001, 0.01 or 0.05 Level, Respectively. Small or Medium Effect Sizes for f^2 and q^2 are Indicated in Italic or Bold Font, Respectively.)*

DV: Privacy Concerns (PC)	Path Coefficient	Effect Size f^2	Effect Size q^2
Adjusted R^2	0.702		
App Price (1 if not free)	-0.001	0.000	0.000
Permission Justification (PJ)	-0.016	0.001	0.000
Perceived Permission Sensitivity (PPS)	0.378***	**0.294**	**0.205**
AR Label	0.000	0.000	0.000
Trust in the MAR App (TRUST)	-0.249***	*0.129*	*0.088*
Age	0.014		
Gender	0.022		
Education	-0.005		
Computer Exp.	0.000		
Smartphone Exp.	0.004		
No. of installed apps	-0.008		
No. of installed MAR apps	0.035		
AR knowledge	-0.004		
Attitudes towards AR in general (ATT)	0.029		
Prior privacy victim experience (VIC)	0.083***	*0.022*	0.016
Privacy Concerns related to MAR Apps in general (PC_{MAR})	0.388***	**0.306**	**0.213**

DV: Trust in the MAR App (TRUST)	Path Coefficient	Effect Size f^2	Effect Size q^2
Adjusted R^2	0.381		
Perceived App Popularity (PAP)	0.191***	*0.053*	*0.042*
Permission Justification (PJ)	0.027	0.001	0.000
Age	-0.088**	0.009	0.006
Gender	-0.046		
Education	-0.048		
Computer Exp.	0.022		
Smartphone Exp.	-0.028		
No. of installed apps	-0.027		
No. of installed MAR apps	0.026		
AR knowledge	-0.041		
Attitudes towards AR in general (ATT)	0.293***	*0.107*	*0.084*
Prior privacy victim experience (VIC)	-0.061*	0.006	0.003
Trust in the App Store ($TRUST_{AS}$)	0.312***	*0.120*	*0.094*
DV: Download Intention (DI)	Path Coefficient	Effect Size f^2	Effect Size q^2
Adjusted R^2	0.612		
Perceived App Popularity (PAP)	0.056*	0.007	*0.046*
App Price (1 if not free)	-0.081***	0.017	*0.053*
AR Label	-0.042*	0.004	*0.044*
Trust in the MAR App (TRUST)	0.308***	*0.113*	*0.128*
PC	-0.130***	*0.028*	*0.062*
PU	0.410***	**0.256**	**0.241**
Age	0.009		
Gender	0.016		
Education	-0.022		
Computer Exp.	-0.010		
Smartphone Exp.	0.024		
No. of installed apps	0.038		
No. of installed MAR apps	-0.007		
AR knowledge	-0.021		
Attitudes towards AR in general (ATT)	0.090**	0.014	*0.050*
Prior privacy victim experience (VIC)	0.049*	0.006	*0.046*

Table 4

*Total Effects (***, **, * Asterisks Indicate Statistical Significance at the 0.001, 0.01 or 0.05 Level, Respectively.)*

DV: Privacy Concerns	Total Effect Size
Perceived App Popularity (PAP)	-0.048***
Permission Justification (PJ)	-0.023
Attitudes towards AR in general (ATT)	-0.044*
Trust in the App Store ($TRUST_{AS}$)	-0.078***

DV: Download Intention (DI)	Total Effect Size
Perceived App Popularity (PAP)	0.121***
App Price (1 if not free)	-0.081***
Permission Justification (PJ)	0.011
Perceived Permission Sensitivity (PPS)	-0.049***
AR Label	-0.042*
Trust in the MAR App (TRUST)	0.341***
Privacy Concerns MAR App (PC)	-0.130***
Perceived Usefulness (PU)	0.410***
Privacy Concerns related to MAR Apps in general (PC_{MAR})	-0.051***
Attitudes towards AR in general (ATT)	0.186***
Trust in the App Store ($TRUST_{AS}$)	0.106***
Prior privacy victim experience (VIC)	0.017

4.3 Differences in AR Knowledge Are Negligible

The correct definition of AR was given by 682 participants which corresponds to 62% of the sample. Therefore, it is essential for the meaningfulness of our results to check whether this finding has an influence on the results. We did so by excluding *AR knowledge* from the model as a control variable and conducting a multigroup analysis. The sample is consequently divided into two sub-samples according to whether participants answered the question correctly (group 1) or not (group 2). The PLS-MGA result indicates that the overall effect of *AR knowledge* is not problematic for our results since there are only four relationships between variables which are significantly different between the two groups. First, the effect of the *number of installed apps* on *privacy concerns* about the MAR app is larger (path coefficient difference equals 0.098) for the participants who correctly answered the AR knowledge question (group 1). Second, *trust in app stores* has a smaller effect on *trust in the MAR app* for group 1 (difference equals -0.147). Third, *perceived app popularity* has a larger effect on *trust in the MAR app* for group 1 (difference equals 0.143). Fourth, *privacy concerns about MAR apps in general* have a smaller effect on privacy concerns about the MAR app for group 1 (difference equals -0.114). All of these effects are relatively small and do not change the main outcomes of the research model.

4.4 Further Robustness Checks to Confirm Existing Results

Our results suggest that the treatments *permission justifications*, *app price* and *AR label* do not have an effect on the respective endogenous variables. We conduct multigroup analyses for each of the treatments to check for other undetected effects of these treatments in the nomological net. We find that there are mostly small differences in relationships between control variables and endogenous variables which do not change the significance of results related to our hypotheses. For example, *privacy concerns* had a smaller (negative) effect on the *download intentions* for participants with the treatment *app price*. A reason for that could be that if people paid for the app, they assumed that the app provider has a business model build on the payment of the app and thus their privacy concerns did not influence the *download intention* as much as for non-paid apps (Kummer & Schulte, 2019).

We augment our knowledge from the SEM with the descriptive statistics of the latent variables (Table 2) in order to have a more detailed overview about the effects in the model. *Privacy concerns regarding the MAR* app are significantly larger for the sensitive permissions group compared to the less sensitive permissions group. Overall, users are concerned about the hypothetical MAR app (median equals 4.875) as well as MAR apps in general (median equals 5). Users are indifferent regarding *trust in the MAR app* and *trust in the app store* (median values of 4 on the 7-point Likert scale indicate the statement “Neither agree or disagree”). Interestingly, the popularity treatment seems to have no effect on the absolute agreement rates of *trust in the MAR app* (median equals 4 for both groups), although we find a low to medium-sized effect of *app popularity* in the SEM. We plotted the mean sum scores of the variables *app popularity* and *trust in the MAR* app with a fitted line. We observed that there is a positive relationship between popularity and trust. Thus, it appears that there is a mediation between the popularity treatment and *trust in the MAR* app with the mediator *app popularity*. The median values show that participants perceive the MAR app as useful (median equals 4.25) and are indifferent with respect to downloading the app. Privacy victim experiences are rare in our

dataset with a median value of 2 (equals the answer “Very infrequently”) and a relatively low standard deviation. *Attitudes towards AR* are positive with a median value of 5.

5. INTERPRETATION OF THE RESULTS

Our model is able to explain a large share of the variance in *privacy concerns* (adjusted R^2 equals 70.2%) and *download intention* (adjusted R^2 equals 61.2%). This result is supported by high levels of predictive relevance Q^2 for these endogenous variables. Thus, it can be assumed that the individual drivers of *privacy concerns* (RQ1) and *download intentions* (RQ2) are relevant constructs for explaining these latent variables.

5.1 Privacy Concerns Can Be Explained by Four Variables

Regarding RQ1, we show that *privacy concerns* are mainly driven by *perceived permission sensitivity* (confirming H4), *trust in the MAR app* (confirming H6a) and non-contextual privacy related constructs (VIC and PC_{MAR}). All path coefficients are statistically significant to the 0.1% level. In addition, all variables except for *prior privacy victim experience* are also showing substantial effect sizes f^2 and q^2, indicating that these three variables contribute the most in explaining contextual privacy concerns related to the MAR app. The fact that *permission justification* does not play a role for the privacy concerns of individuals related to the MAR app indicates that they do not rely on information about why the data is collected, but that they are rather suspicious about the sole fact that certain sensitive information are requested by the app (PPS). Although, prior work finds that such justifications have an effect on privacy concerns (Tan et al., 2014), contrasting work argues that lay users have a hard time understanding the need of apps to access certain smartphone resources (Lin et al., 2012). This matches with the design of our experiment, where users either had to consider the purpose by themselves or understand the given justification. Moreover, privacy notices shown at the download stage are not as effective as shown during the use of the app (Balebako et al., 2015). In earlier studies,

the purpose was given when the app asked for the permission which is either during the install or the usage stage (Tan et al., 2014). Thus, the effect of justifications remains unclear open for future work. Besides that, we can observe that *trust in the MAR app* alleviates *privacy concerns*. We will discuss the mechanisms of how *trust in the MAR app* is built-up in the next subsection.

5.2 Trust in the App Is Largely Driven by Non-Contextual Factors

Our results show that trust in the MAR app is a concept largely based on pre-existing perceptions as well as environmental and institutional factors. *App popularity* serves as a direct trust cue (confirming H1a confirmed) and individuals rely on this information provided in app stores to assess the app itself. In addition, we see that the non-contextual *trust in app stores* is the strongest driver of the specific trust perceptions (path coefficient equals 0.312). There are no differences between users of different mobile operating systems. Thus, this result is stable across technological platforms. Furthermore, we see that *general attitudes towards augmented reality* positively influence *trust in the MAR app*. This stresses the importance of pre-existing attitudes towards technologies, in our case AR, in building specific perceptions towards instances of this technology. *Age* and *prior privacy victim experience* are statistically significant. However, they do not have substantial effect sizes f^2 and predictive relevance q^2. *Permission justifications* do not influence *trust*, indicating that users do not rely on this specific instance of transparency-enhancing information to form their trust perception.

5.3 The Privacy Calculus Determines Download Intentions

RQ2 deals with the effect of *privacy concerns* on the intentions to download the app. Our results indicate that there is a medium-sized negative effect of *privacy concerns* on *download intention* (path coefficient equals -0.130) which is statistically significant at the 0.1% level. The effect sizes f^2 and q^2 are small in size indicating that *privacy concerns* do not explain and predict as much of the variance of *download intentions* as the other significant constructs. When analyzing

these constructs, it becomes apparent that the proposed trade-off of the privacy calculus is in place. The three relevant constructs, which have at least low effect sizes f^2 and q^2, are *privacy concerns* (representing the costs in the trade-off), *perceived usefulness* (representing the benefits) and *trust* (acting as a statistically significant antecedent of *privacy concerns* (Smith et al., 2011)), confirming hypothesis 7, hypothesis 8 and hypothesis 6b, respectively.

App popularity and *app price* are both statistically significant and show predictive relevance regarding *download intentions*. However, the f^2-values are below the threshold for small effect sizes. Certain relationships could also be affected by the hypothetical scenario of our study. For example, we hypothesize that the negative effect of *app price* on *download intention* would have been larger when participants were in a situation in which they actually had to spend money.

The last contextual variable hypothesized to influence *download intentions* is *AR label*. We hypothesized that AR could be perceived as a positive feature due to the innovativeness of the technology. Based on the negative path coefficient, we have to reject the hypothesis since it indicates that users who saw the mockup with the specific AR labels and descriptions are less likely to download the app. However, as for *app popularity* and *app price*, the effect size f^2 is below 0.02 whereas q^2 is above this threshold.

The same holds for the two constructs *attitudes towards AR* and *prior privacy victim experiences* which have statistically significant path coefficients but no substantial effect sizes f^2. Table 5 summarizes the results related to our hypotheses.

Table 5

Summary of the Hypotheses

	Hypothesis	Result
H1a	Perceived app popularity positively influences the trust in the app.	Confirmed
H1b	Perceived app popularity positively influences the intention to download the app.	Confirmed
H2a	The price of the MAR app negatively influences privacy concerns related to the MAR app if the app is not free.	Not confirmed
H2b	The price of the MAR app negatively influences the intention to download the MAR app.	Confirmed
H3a	Permission justifications positively influence trust in the app.	Not confirmed
H3b	Permission justifications negatively influence privacy concerns related to the app.	Not confirmed
H4	Perceived permission sensitivity positively influences privacy concerns related to the app.	Confirmed
H5a	The AR label of the MAR app positively influences privacy concerns related to the app.	Not confirmed
H5b	The AR label of the MAR app positively influences the intention to download the app.	Rejected
H6a	Trust in the MAR app negatively influences privacy concerns related to the app.	Confirmed
H6b	Trust in the MAR app positively influences the intention to download the app.	Confirmed
H7	Privacy concerns related to the MAR app negatively influence the intention to download the app.	Confirmed
H8	The perceived usefulness of the MAR app positively influences the intention to download the app.	Confirmed

5.4 Limitations

Our study has the following limitations. First, vignette-based studies as ours can only uncover perceptions of users, and not their actual behavior. This is relevant since prior research indicates that the link between intentions and behaviors is very weak in the privacy and security context (Crossler et al., 2013). However, other research suggests that privacy perceptions in hypothetical scenarios still provide valuable insights although users might undervalue behavioral factors (Adjerid et al., 2018). A second limitation might exist because privacy perceptions differ between countries and cultures. Since our sample contains only German participants, the results can possibly differ from surveys conducted in other countries or cultural regions. Third, misunderstandings of questionnaire items or wrong answers given by the participants could

results in additional biases in the context of studies based on self-reports in online questionnaires. Such behaviors of participants can have different causes, e.g., a specific mood in which participants are when filling out the survey or the social desirability bias.

6. PRACTICAL CONTRIBUTIONS

It is crucial to consider privacy before the technologies are fully established in the market since we can easier steer technologies towards a privacy-friendly design before the economic imperative of today's internet economy captures and directs the future technological developments. This idea is also in accordance with privacy by design. Once a technology is established, changing it to a privacy-friendly design is more effort and often requires regulations to level the field for all competitors. It is likely that AR will be used as a tool for an enhanced exploitation of personal information given the described possibilities to capture highly contextual information. Thus, it is required to think about ways to protect users' privacy when interacting with MAR apps. We discuss three possible solutions which address app developers, operators and regulators.

First, when manipulating the permissions for the hypothetical MAR app, it became apparent that the current set of available mobile permissions does not reflect all context-specific information that is potentially used by AR. In order to achieve a minimum level of agency, users have to know what permissions like the camera permission in the Android operation system exactly do and mean in the context of using an MAR app. For example, Android paraphrases each permission with other words. For camera, it says *Take pictures and videos*. But what does that mean for MAR apps? Is the app processing the raw output from the camera in the cloud or on the device itself when users measure their homes with the kind of app in our experiment? Is the app capturing any information before you start measuring (e.g., when you just open the app and “look at the real world through your device”)? What specific information from the camera output are saved on the app provider's servers? All these questions cannot be addressed at the

moment for AR, but they should be for the sake of transparency. By that, we could derive permissions which provide users with context-dependent information about *what* information is collected in *which context and at what time*. Developers have to pay attention to find a good balance between showing transparent context-dependent resource accesses to the user's personal information and permission fatigue of the user. Possible new permission can be implemented with the help of an "[...] intermediary protection layer between the applications and device resources" (de Guzman et al., 2018, p. 10). For example, such permissions could include the decision whether to allow or decline object or face recognition. Technical research on the issue of object and face sanitization exist for almost ten years (Jana et al. 2013). Such options are relevant since users might want to protect information on objects like medicine containers or want to disable the option for apps to capture every person's face while using the app. Against this backdrop, we argue that users should be able to decide whether they want object and face recognition in the first place in the form of a new permission.

Second, further implications can be derived from our results regarding the permission justifications. Our results show that the information about the purpose of accessing certain smartphone resources do not affect privacy concerns or trust in the app. Prior work even finds that such information can be misleading and damaging since it alleviates privacy concerns if developers just write any piece of information (Tan et al., 2014). Research on social networks shows that more privacy controls mislead users to feel a wrong sense of protection and cause more disclosure of private information (Brandimarte et al., 2013). All these examples show that it is apparently not clear to users that the prevailing economic imperative for most online services (including the smartphone ecosystem) is based on collecting, processing and selling personal information.

Third, even if users do have above-average knowledge on possible privacy issues on the internet, high levels of asymmetric information hinder users to make informed decisions. In our scenario, which reflects relatively complex information surrounding the app download phase,

trust cues serve as important drivers which alleviate *privacy concerns* and drive *download intentions* (see Table 4 for total effects of *perceived app popularity*). However, trust cues as download numbers are only a rough approximation about the popularity of an app and do by no means say anything about privacy issues of an app. The problem is that there is no direct solution or recommendation to decrease such effects. We rather suggest an approach taking into account two parts. On the one hand, developers, operators and app stores should be forced by regulation to decrease the immanent information asymmetry (this would be possible for MAR apps with new permissions as suggested before). On the other hand, users must be made aware about influencing cues in the app environment in order to recognize and judge them properly. This can be done with respective notes in app stores or with technical solutions which assess the privacy properties of apps (Bal et al., 2015; Wijesekera et al., 2015).

7. THEORETICAL CONTRIBUTIONS

This work is among the first ones investigating privacy concerns related to MAR apps based on a highly context-dependent model, thus, following calls and acknowledging the importance of context in privacy research (Acquisti et al., 2015; Nissenbaum, 2010; Xu et al., 2012).

We used the widely known APCO model from the privacy literature and augmented it with the framework of contextual integrity as a starting point to figure out contextually relevant factors which determine privacy concerns in our chosen context. Hereby, our work is among the few ones which systematically maps the theoretical aspects of CI with latent variables which can be manipulated and measured in an empirical model. Additionally, we showed the relevance of two other contextual factors for the context of MAR apps (*app price* and the *AR label*).

We contribute to literature by developing one of the first models which explains factors influencing privacy concerns related to MAR apps. Besides that, our empirical model is able to explain over 70% of the variance in *the privacy concerns* related to the app itself as well as over 60% of the variance in *download intentions*. Despite that our chosen research approach faces

the possible problem of underestimating behavioral factors influencing our endogenous variables (Adjerid et al., 2018), we could still show that analyzing the download stage based on a hypothetical scenario yields important results for practitioners and researchers alike. We presented a new structure of permissions for MAR apps which could be conceptually applied to other types of AR (e.g., smart glasses), too. By that, we contribute to the large stream of research on permissions and propose a new way of thinking about permissions according to the contextual information they represent.

8. FUTURE WORK

Our research provides rich opportunities for future work. Most importantly, future work could analyze how AR-specific permissions are perceived by users and whether they provide them more value in terms of better understanding what the MAR app does. Such research is important for future developments with respect to MAR apps as well as smart glasses since this is needed for enabling users to protect their privacy when interacting with this new type of technology.

A second avenue for future research is related to the interaction of permissions. The question arises how different permissions might interact with each other and create unknown harms to users. Such interactions are not transparent to the users and we should be thinking about analyzing them technically and making them transparent to the user.

Third, future work could consider altering specific treatment variables used in our experiment. For example, star-based ratings as additional source of information influencing trust of users in an app store environment could be investigated. Additionally, the AR labels and descriptions could be manipulated in different ways in order to investigate the effect of showing participants AR-related information.

Furthermore, conducting our study in another cultural setting could provide insights into the underlying social norms surrounding the research context. There are great challenges for future

work in developing normative recommendations for regulating AR systems in general due to the co-constitutive nature of AR and artificial intelligence (AI) (Benthall et al., 2017). AI is an important enabler for the proper functioning of AR in terms of providing real-time context-dependent information to the user. Thus, these two technologies are almost inseparably connected with each other (especially in the medium term when thinking about smart glasses in the B2C context). Hence, we should not only consider the capabilities and related risks of AR, but also the new issues which can come up due to the interconnection with AI.

9. CONCLUSION

We conducted a vignette-based online experiment with 1,100 smartphone users in Germany to elicit context-dependent drivers of *privacy concerns*, *trust* and *download intentions* for a hypothetical MAR app. Our results show that *privacy concerns* related to the MAR app are primarily driven by *permission sensitivity, trust in MAR the app* and *privacy concerns towards MAR apps in general*. *Download intentions* are driven by *trust in the MAR app* and the variables from the privacy calculus, i.e., *privacy concerns* and *perceived usefulness*. Based on these results, we conclude that transparency should be enhanced for future developments of AR, e.g., by adapting the current set of permission to the context-dependent nature of the technology. Augmented reality is a technology with great promises but also big risks for individual self-determination. Research like ours provides insights to inform the development and regulation of AR in order to fulfill the great promises of AR while protecting the privacy and autonomy of individuals.

REFERENCES

Ackerman, M. S., & Mainwaring, S. D. (2005). Privacy issues and human-computer interaction. In S. Garfinkel & L. Cranor (Eds.), *Security and Usability: Designing Secure Systems That People Can Use* (pp. 381–400). O'Reilly, Sebastopol, CA.

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, *347*(6221), 509–514. https://doi.org/10.2139/ssrn.2580411

Adjerid, I., Peer, E., & Acquisti, A. (2018). BEYOND THE PRIVACY PARADOX: OBJECTIVE VERSUS RELATIVE RISK IN PRIVACY DECISION MAKING. *MIS Quarterly*, *42*(2), 465–488. https://doi.org/10.25300/MISQ/2018/14316

Android Developers. (2019). *Android Permissions Overview*. https://developer.android.com/guide/topics/permissions/overview#permission-groups

Apple. (2019). *Apple Measure App*. https://support.apple.com/en-us/HT208924

Azuma, R. T., Baillot, Y., Feiner, S., Julier, S., Behringer, R., & Macintyre, B. (2001). Recent Advances in Augmented Reality. *IEEE Computer Graphics And Applications*, *21*(6), 34–47.

Bal, G. (2014). Explicitness of Consequence Information in Privacy Warnings: Experimentally Investigating the Effects on Perceived Risk, Trust, and Privacy Information Quality. *ICIS 2014 Proceedings*.

Bal, G., Rannenberg, K., & Hong, J. I. (2015). Styx: Privacy risk communication for the Android smartphone platform based on apps' data-access behavior patterns. *Computers & Security*, *53*(September), 187–202. https://doi.org/10.1016/j.cose.2015.04.004

Balebako, R., Schaub, F., Adjerid, I., Acquisti, A., & Cranor, L. (2015). The impact of timing on the salience of smartphone app privacy notices. *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, 63–74.

Bamberger, K. A., Egelman, S., Han, C., Elazari, A., On, B., & Reyes, I. (2020). CAN YOU PAY FOR PRIVACY ? CONSUMER EXPECTATIONS AND THE BEHAVIOR OF FREE AND PAID APPS. *BERKELEY TECHNOLOGY LAW JOURNAL (BTLJ)*, *35*.

Bansal, G., Zahedi, F. M., & Gefen, D. (2010). The impact of personal dispositions on information sensitivity, privacy concern and trust in disclosing health information online. *Decision Support Systems*, *49*(2), 138–150. https://doi.org/10.1016/j.dss.2010.01.010

Barth, A., Datta, A., Mitchell, J. C., & Nissenbaum, H. (2006). Privacy and contextual integrity:

Framework and applications. *IEEE Symposium on Security and Privacy (S&P'06)*, 184–198. https://doi.org/10.1109/SP.2006.32

Benthall, S., Gürses, S., & Nissenbaum, H. (2017). Contextual Integrity through the Lens of Computer Science. *Foundations and Trends® in Privacy and Security*, *2*(1), 1–69. https://doi.org/10.1561/3300000016

Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced Confidences: Privacy and the Control Paradox. *Social Psychological and Personality Science*, *4*(3), 340–347. https://doi.org/10.1177/1948550612455931

Chen, R., & Sharma, S. K. (2015). Learning and self-disclosure behavior on social networking sites: The case of Facebook users. *European Journal of Information Systems*, *24*(1), 93–106.

Chin, W. W. (1998). The Partial Least Squares Approach to Structural Equation Modeling. In G. A. Marcoulides (Ed.), *Modern Methods for Business Research* (pp. 295–336). Lawrence Erlbaum.

Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences*.

Crossler, R. E., Johnston, A. C., Lowry, P. B., Hu, Q., Warkentin, M., & Baskerville, R. (2013). Future directions for behavioral information security research. *Computers and Security*, *32*(June), 90–101. https://doi.org/10.1016/j.cose.2012.09.010

Dacko, S. G. (2017). Enabling smart retail settings via mobile augmented reality shopping apps. *Technological Forecasting and Social Change*, *124*, 243–256. https://doi.org/10.1016/j.techfore.2016.09.032

de Guzman, J. A., Thilakarathna, K., & Seneviratne, A. (2018). Security and Privacy Approaches in Mixed Reality: A Literature Survey. *Http://Arxiv.Org/Abs/1802.05797*. https://doi.org/arXiv:1802.05797v2

Degirmenci, K. (2020). Mobile users' information privacy concerns and the role of app permission requests. *International Journal of Information Management*, *50*(July 2019),

261–272. https://doi.org/10.1016/j.ijinfomgt.2019.05.010

Denning, T., Dehlawi, Z., & Kohno, T. (2014). In situ with bystanders of augmented reality glasses. *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems - CHI '14*, 2377–2386. https://doi.org/10.1145/2556288.2557352

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

Döring, N., & Bortz, J. (2016). *Forschungsmethoden und Evaluation in den Sozial- und Humanwissenschaften* (6th ed.). Springer-Verlag. https://doi.org/10.1007/978-3-642-41089-5

Duan, W., Gu, B., & Whinston, A. B. (2009). Informational Cascades and Software Adoption on the Internet: An Empirical Investigation. *MIS Quarterly*, *33*(1), 23–48. https://doi.org/10.2307/20650277

Frik, A., Malkin, N., Harbach, M., Peer, E., & Egelman, S. (2019). A Promise Is A Promise The. *CHI '19*, 1–12. https://doi.org/10.1177/001452467808901006

Google Play Developer Policy Center. (2020). *Privacy, Security, and Deception*. https://play.google.com/about/privacy-security-deception/permissions/

Gu, J., Xu, Y. (Calvin), Xu, H., Zhang, C., & Ling, H. (2017). Privacy concerns for mobile app download: An elaboration likelihood model perspective. *Decision Support Systems*, *94*, 19–28. https://doi.org/10.1016/j.dss.2016.10.002

Hair, J., Hult, G. T. M., Ringle, C. M., & Sarstedt, M. (2017). *A Primer on Partial Least Squares Structural Equation Modeling (PLS-SEM)*. SAGE Publications.

Hair, J., Ringle, C. M., & Sarstedt, M. (2011). PLS-SEM: Indeed a Silver Bullet. *The Journal of Marketing Theory and Practice*, *19*(2), 139–152. https://doi.org/10.2753/MTP1069-6679190202

Han, C., Reyes, I., Elazari, A., On, B., Reardon, J., Feal, Á., Bamberger, K. A., Egelman, S., & Vallina-rodriguez, N. (2019). Do You Get What You Pay For? Comparing The Privacy

Behaviors of Free vs. Paid Apps. *Workshop on Technology and Consumer Protection (ConPro '19)*, 1–7.

Harborth, D. (2017). Augmented Reality in Information Systems Research: A Systematic Literature Review. *Twenty-Third Americas Conference on Information Systems (AMCIS)*, 1–10.

Harborth, D. (2019). Unfolding Concerns about Augmented Reality Technologies: A Qualitative Analysis of User Perceptions. *Wirtschaftsinformatik (WI19)*, 1262–1276.

Harborth, D., Hatamian, M., Tesfay, W. B., & Rannenberg, K. (2019). A Two-Pillar Approach to Analyze the Privacy Policies and Resource Access Behaviors of Mobile Augmented Reality Applications. *Hawaii International Conference on System Sciences (HICSS) Proceedings*, 5029–5038.

Harborth, D., & Kreuz, H. (2020). Exploring the Attitude Formation Process of Individuals Towards New Technologies: The Case of Augmented Reality. *International Journal of Technology Marketing*, *14*(2), 125–153. https://doi.org/10.1504/IJTMKT.2020.10031990

Harborth, D., & Pape, S. (2018). Privacy Concerns and Behavior of Pokémon Go Players in Germany. In M. Hansen, E. Kosta, I. Nai-Fovino, & S. Fischer-Hübner (Eds.), *Privacy and Identity Management. The Smart Revolution. Privacy and Identity 2017. IFIP Advances in Information and Communication Technology, vol 526* (pp. 314–329). Springer, Cham. https://doi.org/https://doi.org/10.1007/978-3-319-92925-5_21

Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the Academy of Marketing Science*, *43*(1), 115–135. https://doi.org/10.1007/s11747-014-0403-8

Jana, S., Narayanan, A., & Shmatikov, V. (2013). A Scanner Darkly: Protecting User Privacy From Perceptual Applications. *IEEE Symposium on Security and Privacy*, 349–363. https://doi.org/10.1109/SP.2013.31

Kang, J. Y. M., Mun, J. M., & Johnson, K. K. P. (2015). In-store mobile usage: Downloading and

usage intention toward mobile location-based retail apps. *Computers in Human Behavior*, *46*, 210–217. https://doi.org/10.1016/j.chb.2015.01.012

Keith, M. J., Thompson, S. C., Hale, J., Lowry, P. B., & Greer, C. (2013). Information disclosure on mobile devices: Re-examining privacy calculus with actual user behavior. *International Journal of Human Computer Studies*, *71*(12), 1163–1173. https://doi.org/10.1016/j.ijhcs.2013.08.016

Kelley, P. G., Cranor, L. F., & Sadeh, N. (2013). Privacy as part of the app decision-making process. *CHI '13 Proceedings*, 3393–3402. https://doi.org/10.1145/2470654.2466466

Kim, S. C., Yoon, D., & Han, E. K. (2016). Antecedents of mobile app usage among smartphone users. *Journal of Marketing Communications*, *22*(6), 653–670. https://doi.org/10.1080/13527266.2014.951065

Kohn, V., & Harborth, D. (2018). AUGMENTED REALITY – A GAME CHANGING TECHNOLOGY FOR MANUFACTURING PROCESSES? *Twenty-Sixth European Conference on Information Systems (ECIS2018)*, 1–19.

Kummer, M., & Schulte, P. (2019). When private information settles the bill: Money and privacy in Google's market for smartphone applications. *Management Science*, *65*(8), 3470–3494.

Lebeck, K., Ruth, K., Kohno, T., & Roesner, F. (2018). Towards Security and Privacy for Multi-user Augmented Reality: Foundations with End Users. *IEEE Symposium on Security and Privacy*, 392–408. https://doi.org/10.1109/SP.2018.00051

Lin, J., Amini, S., Hong, J. I., Sadeh, N., Lindqvist, J., & Zhang, J. (2012). Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 501–510.

Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, *15*(4), 336–355. https://doi.org/10.1287/isre.1040.0032

Martin, K. E., & Nissenbaum, H. (2016). MEASURING PRIVACY: AN EMPIRICAL TEST USING CONTEXT TO EXPOSE CONFOUNDING VARIABLES. *The Columbia Science & Technology Law Review*, *18*(Fall), 176–218.

Martin, K. E., & Nissenbaum, H. (2019). What Is It About Location? *Available at: Https://Ssrn.Com/Abstract=3360409*, 101–174. https://doi.org/https://doi.org/10.15779/Z382F7JR6F

McKnight, D. H., Carter, M., Thatcher, J. B., & Clay, P. F. (2011). Trust in a Specific Technology: An Investigation of Its Components and Measures. *ACM Transactions on Management Information Systems (TMIS)*, *2*(2), 1–25. https://doi.org/10.1145/1985347.1985353

McNaney, R., Vines, J., Roggen, D., Balaam, M., Zhang, P., Poliakov, I., & Olivier, P. (2014). Exploring the Acceptability of Google Glass as an Everyday Assistive Device for People with Parkinson's. *32nd Annual ACM Conference on Human Factors in Computing Systems*, 2551–2554.

Microsoft. (2017). *Microsoft HoloLens*. https://www.microsoft.com/microsoft-hololens/en-us/buy

Nellis, S. (2017). *Google, Apple face off over augmented reality technology*. Reuters. https://www.reuters.com/article/us-google-apple/google-apple-face-off-over-augmented-reality-technology-idUSKCN1BA001

Nicas, J., & Zakrzewski, C. (2016). *Augmented Reality Gets Boost From Success of 'Pokémon Go.'* Wall Street Journal. https://www.wsj.com/articles/augmented-reality-gets-boost-from-success-of-pokemon-go-1468402203

Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy and the Integrity of Social Life*. Stanford University Press.

Pavlou, P. A. (2003). Consumer Acceptance of Electronic Commerce: Integrating Trust and Risk with the Technology Acceptance Model. *International Journal of Electronic Commerce*, *7*(3), 101–134. https://doi.org/10.1080/10864415.2003.11044275

Peterson, A. (2016). *Pokémon Go had "full access" to the Google accounts of some iPhone players*. Washington Post. https://www.washingtonpost.com/news/the-switch/wp/2016/07/12/pokemon-go-had-full-access-to-the-google-accounts-of-some-iphone-players/

Petrock, V. (2020). *US Virtual and Augmented Reality Users 2020*. EMarketer. https://www.emarketer.com/content/us-virtual-and-augmented-reality-users-2020

Podsakoff, P. M., MacKenzie, S. B., Lee, J. Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research: a critical review of the literature and recommended remedies. *Journal of Applied Psychology*, *88*(5), 879–903. https://doi.org/10.1037/0021-9010.88.5.879

Rauschnabel, P. A., He, J., & Ro, Y. K. (2018). Antecedents to the adoption of augmented reality smart glasses: A closer look at privacy risks. *Journal of Business Research*, *92*, 374–384. https://doi.org/10.1016/J.JBUSRES.2018.08.008

Ringle, C. M., Wende, S., & Becker, J. M. (2015). *SmartPLS 3*. Boenningstedt: SmartPLS GmbH, http://www.smartpls.com.

Rossolillo, N. (2020). *3 Ways to Invest in Virtual and Augmented Reality*. The Motley Fool. https://www.fool.com/investing/2020/07/16/3-ways-to-invest-in-virtual-and-augmented-reality.aspx

Smith, H. J., Dinev, T., & Xu, H. (2011). Theory and Review Information Privacy Research: An Interdisciplinary Review. *MIS Quarterly*, *35*(4), 989–1015.

Tan, J., Nguyen, K., Theodorides, M., Negrón-Arroyo, H., Thompson, C., Egelman, S., & Wagne, D. (2014). The effect of developer-specified explanations for permission requests on smartphone user behavior. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, 91–100. https://doi.org/10.1145/2556288.2557400

van der Meulen, E., Cidotã, M.-A., Lukosch, S. G., Bank, P. J. M., van der Helm, A. J. C., & Visch, V. T. (2016). A Haptic Serious Augmented Reality Game for Motor Assessment of

Parkinson's Disease Patients. In E. E. Veas, T. Langlotz, J. Martinez-Carranza, R. Grasset, M. Sugimoto, & A. Martin (Eds.), *International Symposium on Mixed and Augmented Reality, ISMAR 2016 Adjunct* (pp. 102–104). IEEE. https://doi.org/10.1109/ISMAR-Adjunct.2016.0050

Venkatesh, V., Thong, J., & Xu, X. (2012). Consumer Acceptance and User of Information Technology: Extending the Unified Theory of Acceptance and Use of Technology. *MIS Quarterly*, *36*(1), 157–178.

Wijesekera, P., Baokar, A., Hosseini, A., Egelman, S., Wagner, D., & Beznosov, K. (2015). Android Permissions Remystified: A Field Study on Contextual Integrity. *Proceedings of the 24th USENIX Security Symposium*, 499–514.

Wu, J. J., Lien, C. H., Mohiuddin, M., Chien, S. H., & Yang, X. J. (2016). The effects of smartphone users' core self-evaluations and stickiness on intentions to download free social media apps. *Journal of Decision Systems*, *25*(3), 263–272. https://doi.org/10.1080/12460125.2016.1187549

Xu, H., Teo, H.-H., Tan, B. C. Y., & Agarwal, R. (2012). Research Note-Effects of Individual Self-Protection, Industry Self-Regulation, and Government Regulation on Privacy Concerns: A Study of Location-Based Services. *Information Systems Research*, *23*(4), 1342–1363. https://doi.org/10.1287/isre.1120.0416

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books.

APPENDIX

A. Survey Items

All items are measured on a 7-point Likert scale ranging from "strongly disagree" to "strongly agree", if not otherwise indicated.

Perceived Permission Sensitivity (Gu et al., 2017)

PPS1. Measure it! (AR) requests many permissions.

PPS2. Measure it! (AR) requests sensitive permissions.

PPS3. The potential risk related to the permission requests of Measure it! (AR) is high.

Perceived App Popularity (Gu et al., 2017)

PAP1. I think Measure it! (AR) is popular.

PAP2. Measure it! (AR) is downloaded numerous times.

PAP3. I think Measure it! (AR) is hot among users.

Privacy Concerns related to Measure it! (AR) (Gu et al., 2017)

PC1. I think Measure it! (AR) will over-collect my personal information.

PC2. I will worry that Measure it! (AR) leaks my personal information to irrelevant third-parties.

PC3. If I were to download and use this app, I would be concerned that Measure it! (AR) would violate my privacy.

PC4. If I were to download and use this app, I would be concerned that Measure it! (AR) would misuse my personal information.

Trust in Measure it! (AR) (Pavlou, 2003)

TRUST1. Measure it! (AR) is trustworthy.

TRUST2. Measure it! (AR) keeps promises and commitments.

TRUST3. I trust Measure it! (AR) because they keep my best interests in mind.

Perceived Usefulness of Measurement Apps (Venkatesh et al., 2012)

PU1. I find Measurement Apps like Measure it! (AR) useful in my daily life.

PU2. Using Measurement Apps like Measure it! (AR) increases my chances of achieving things that are important to me.

PU3. Using Measurement Apps like Measure it! (AR) helps me accomplish things more quickly.

PU4. Using Measurement Apps like Measure it! (AR) increases my productivity.

Intention to Download Measure it! (AR) (Gu et al., 2017)

DI1. I am willing to download Measure it! (AR).

DI2. After reading the related information of Measure it! (AR), I am willing to try Measure it! (AR).

DI3. Based on the given information, I would prefer Measure it! (AR) over comparable apps.

(Mobile) Privacy Victim Experience (Gu et al., 2017; Malhotra et al., 2004)

VIC. How frequently have you personally been the victim of what you felt was an improper privacy invasion from your installed mobile apps?

Note: measured on a 7-point frequency scale ranging from “never” to “very frequently”.

Privacy Concerns about Mobile Augmented Reality Apps in General (self-made)

Measure it! is a so-called mobile augmented reality (MAR) application aligning digital objects with the real environment.

PCMAR1. I perceive MAR applications as more privacy-invasive compared to non-MAR applications.

PCMAR2. I am concerned being surveilled by MAR applications.

PCMAR3. I am concerned being filmed by MAR applications.

PCMAR4. I am concerned that MAR applications distribute the gathered data without my knowledge to third-parties.

Trust in the App Store (Pavlou, 2003)

TRUSTAS1. App Stores are trustworthy in protecting my privacy against malicious apps.

TRUSTAS2. App Stores keep promises and commitments

TRUSTAS3. I trust App Stores because they keep my best interests in mind.

Overall Preexisting Attitude Towards Augmented Reality (Chen & Sharma, 2015)

Your overall attitude toward using Augmented Reality in general is:

ATT1. Good

ATT2. Beneficial

ATT3. Positive

ATT4. Favorable

B. Permission Distribution and Exemplary Mockups

Figure 3

Comparison of the Requested Permissions of MAR Apps Compared to Non-MAR Apps

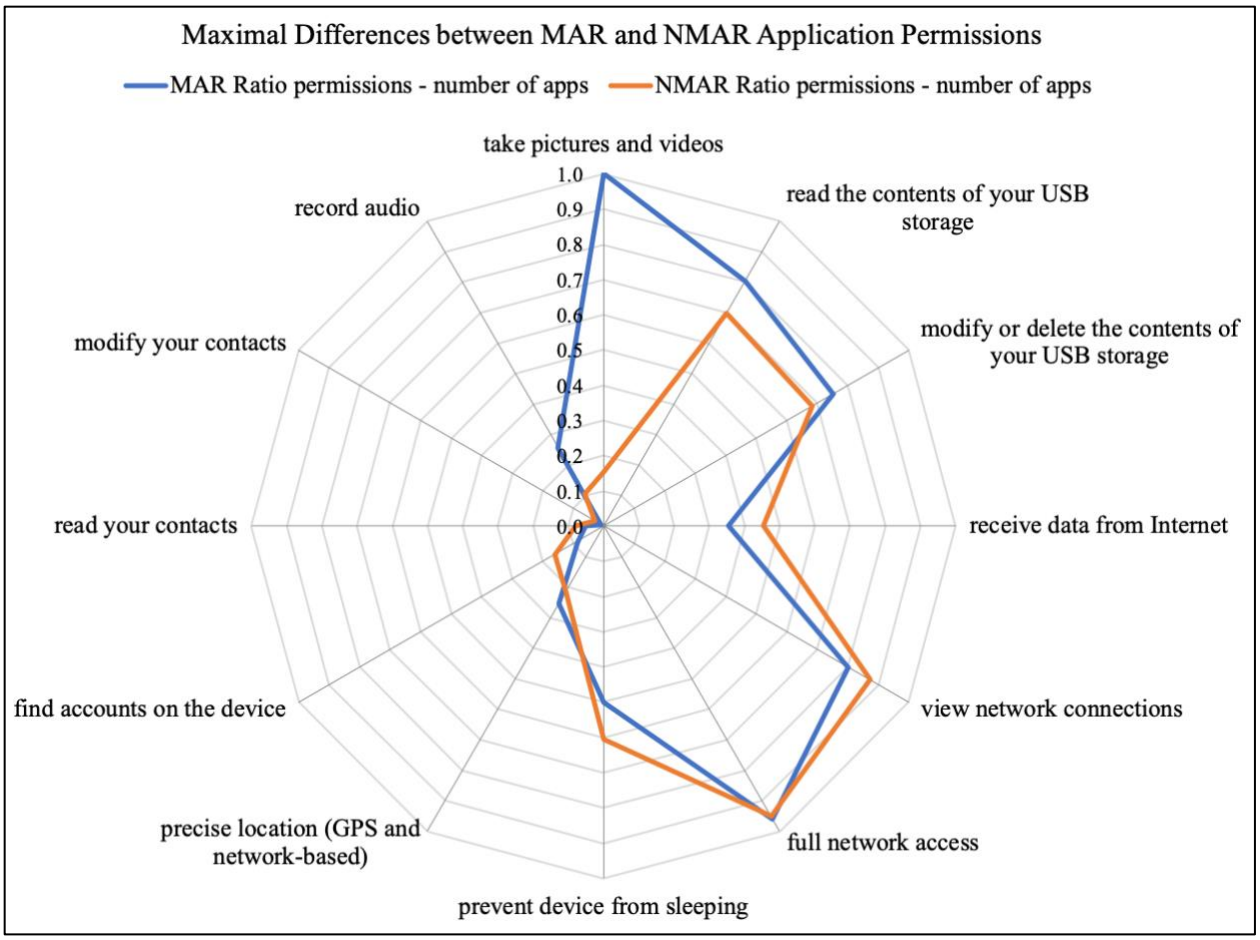

Figure 4

Exemplary Mockup with Minimum Set of Treatments (as Defined in Our Design)

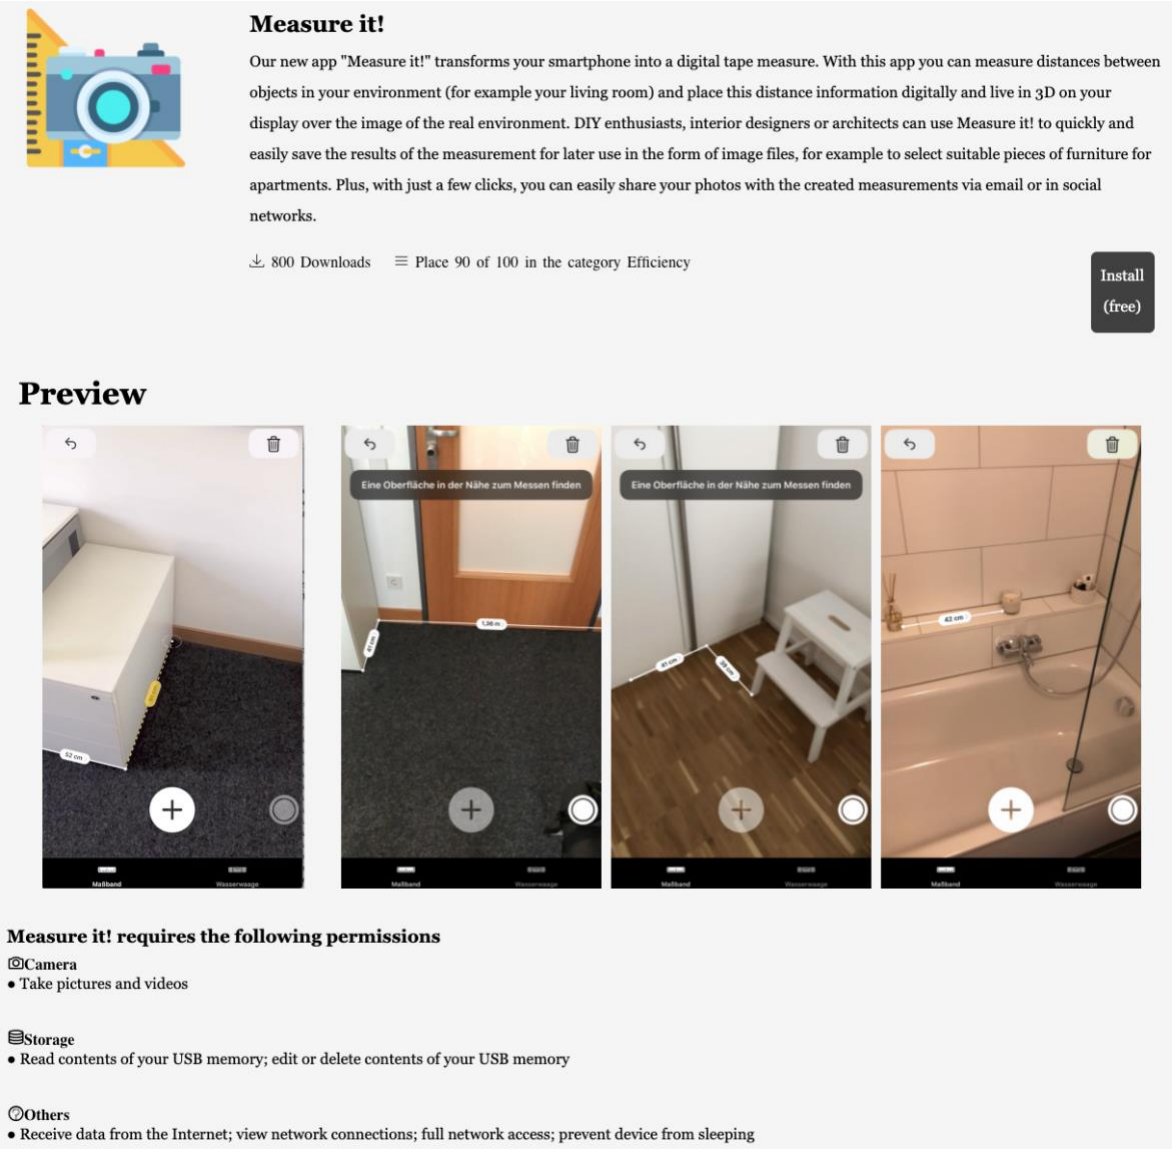

Figure 5

Exemplary Mockup with Maximum Set of Treatments (as Defined in Our Design)

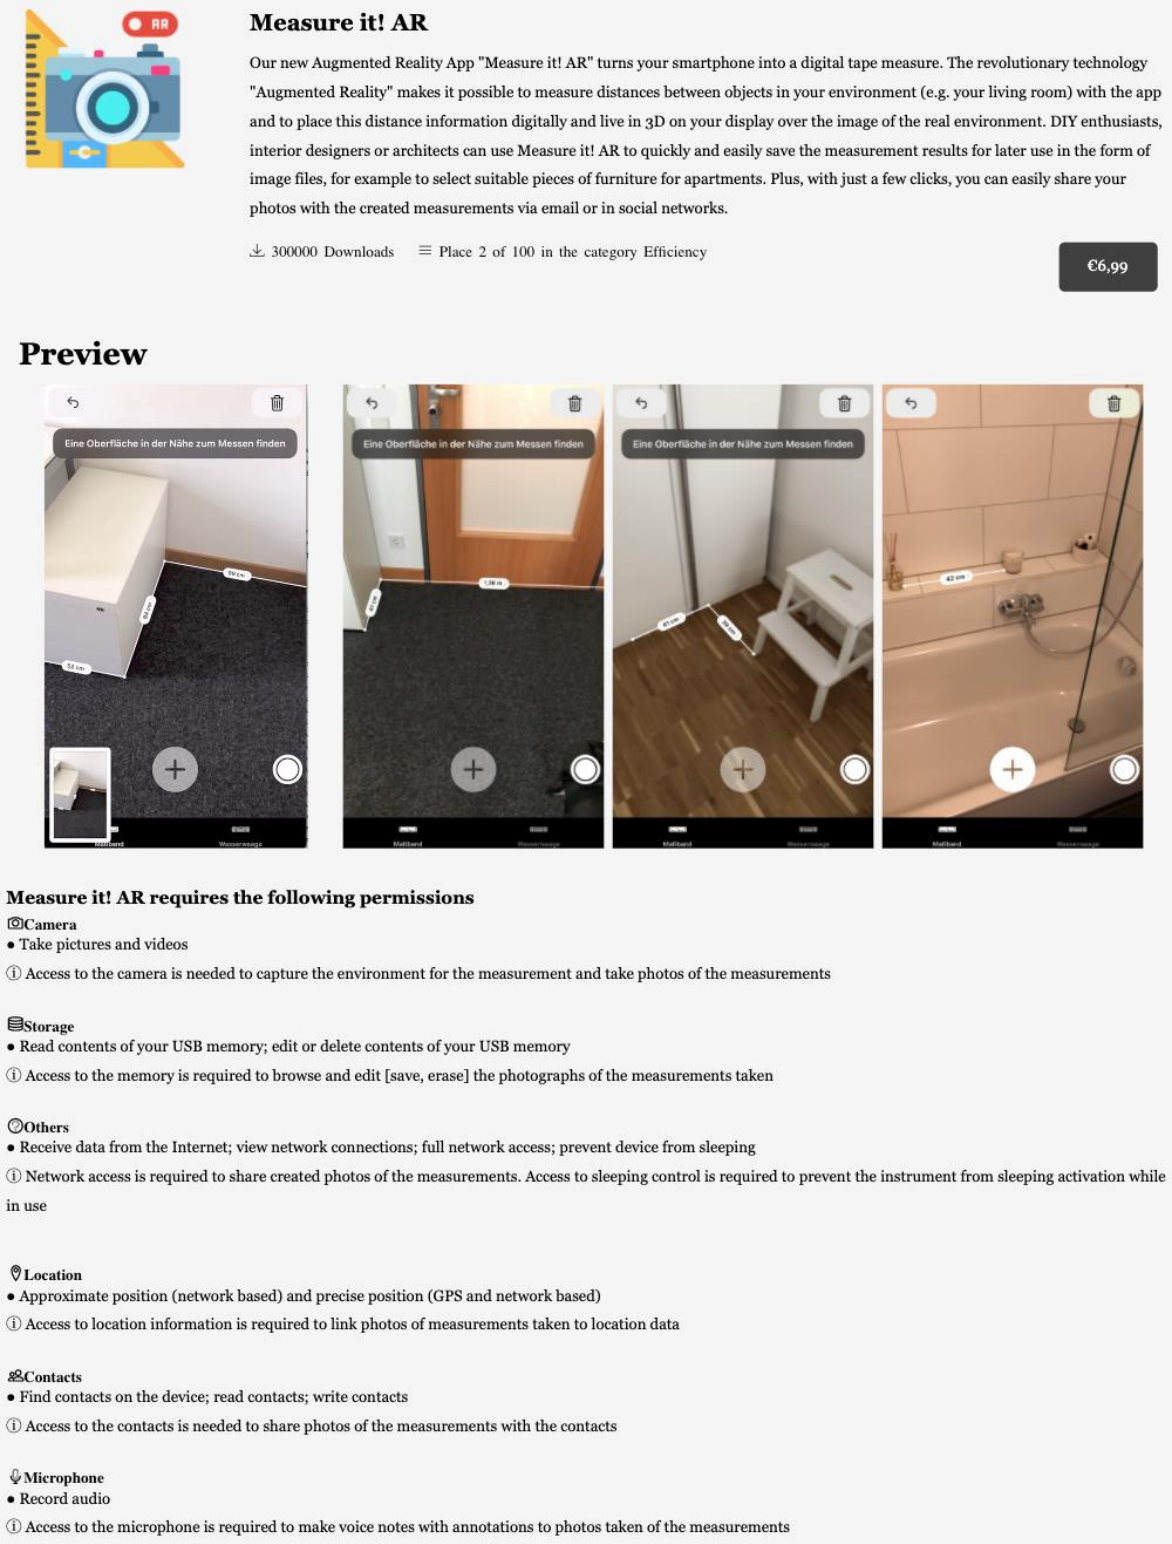

C. Vignettes

Table 6

Group Distribution

Group	Number of participants	App Popularity	Permission Sensitivity	Permission Justification	AR label	Free app
1	39	1	1	1	0	0
2	32	1	1	1	1	0
3	35	1	1	1	0	1
4	34	1	1	1	1	1
5	32	1	1	0	0	0
6	36	1	1	0	1	0
7	33	1	1	0	0	1
8	34	1	1	0	1	1
9	33	1	0	1	0	0
10	33	1	0	1	1	0
11	35	1	0	1	0	1
12	35	1	0	1	1	1
13	35	1	0	0	0	0
14	35	1	0	0	1	0
15	34	1	0	0	0	1
16	33	1	0	0	1	1
17	34	0	1	1	0	0
18	37	0	1	1	1	0
19	32	0	1	1	0	1
20	32	0	1	1	1	1
21	33	0	1	0	0	0
22	33	0	1	0	1	0
23	33	0	1	0	0	1
24	39	0	1	0	1	1
25	34	0	0	1	0	0
26	33	0	0	1	1	0
27	33	0	0	1	0	1
28	41	0	0	1	1	1
29	35	0	0	0	0	0
30	34	0	0	0	1	0
31	35	0	0	0	0	1
32	34	0	0	0	1	1
Sum	1100	1: 548	1: 548	1: 545	1: 548	1: 548
		0: 552	0: 552	0: 555	0: 552	0: 552

D. Measurement Model Assessment

Table 7

Loadings and Cross-Loadings of the Reflective Items and Internal Consistency Reliability

Construct	DI	TRUST	PC	PU	PAP	PPS	$TRUST_{AS}$	PC_{MAR}	ATT
DI1	**0.951**	0.608	-0.387	0.679	0.309	-0.313	0.342	-0.326	0.491
DI2	**0.962**	0.629	-0.432	0.650	0.299	-0.350	0.357	-0.362	0.486
DI3	**0.927**	0.634	-0.408	0.609	0.362	-0.298	0.346	-0.338	0.453
TRUST1	0.629	**0.947**	-0.555	0.477	0.320	-0.429	0.472	-0.411	0.457
TRUST2	0.556	**0.923**	-0.468	0.471	0.296	-0.327	0.465	-0.359	0.432
TRUST3	0.655	**0.934**	-0.576	0.549	0.344	-0.416	0.477	-0.435	0.450
PC1	-0.359	-0.482	**0.892**	-0.183	-0.106	0.711	-0.237	0.612	-0.192
PC2	-0.395	-0.532	**0.954**	-0.226	-0.117	0.658	-0.300	0.696	-0.233
PC3	-0.430	-0.565	**0.967**	-0.263	-0.146	0.679	-0.321	0.702	-0.250
PC4	-0.444	-0.581	**0.961**	-0.291	-0.156	0.663	-0.351	0.695	-0.274
PU1	0.653	0.469	-0.237	**0.908**	0.280	-0.150	0.320	-0.212	0.505
PU2	0.634	0.521	-0.251	**0.924**	0.337	-0.180	0.392	-0.226	0.460
PU3	0.620	0.501	-0.232	**0.951**	0.324	-0.146	0.377	-0.204	0.501
PU4	0.636	0.504	-0.235	**0.947**	0.322	-0.155	0.373	-0.216	0.495
PAP1	0.334	0.334	-0.133	0.341	**0.950**	-0.109	0.236	-0.138	0.232
PAP2	0.244	0.250	-0.075	0.225	**0.913**	-0.074	0.205	-0.088	0.159
PAP3	0.364	0.366	-0.170	0.366	**0.964**	-0.149	0.272	-0.157	0.258
PPS1	-0.246	-0.308	0.572	-0.104	-0.062	**0.891**	-0.102	0.434	-0.075
PPS2	-0.268	-0.338	0.619	-0.123	-0.116	**0.937**	-0.126	0.491	-0.074
PPS3	-0.397	-0.485	0.758	-0.221	-0.147	**0.929**	-0.257	0.629	-0.203
$TRUST_{AS}1$	0.323	0.455	-0.285	0.355	0.200	-0.152	**0.945**	-0.275	0.417
$TRUST_{AS}2$	0.344	0.473	-0.297	0.366	0.257	-0.157	**0.960**	-0.298	0.422
$TRUST_{AS}3$	0.374	0.498	-0.328	0.388	0.264	-0.212	**0.930**	-0.316	0.401
$PC_{MAR}1$	-0.304	-0.307	0.530	-0.186	-0.081	0.461	-0.173	**0.781**	-0.259
$PC_{MAR}2$	-0.329	-0.398	0.656	-0.229	-0.154	0.513	-0.302	**0.928**	-0.233
$PC_{MAR}3$	-0.343	-0.451	0.724	-0.217	-0.147	0.584	-0.335	**0.935**	-0.247
$PC_{MAR}4$	-0.297	-0.349	0.601	-0.177	-0.101	0.456	-0.278	**0.878**	-0.246
ATT1	0.491	0.469	-0.254	0.498	0.222	-0.145	0.429	-0.271	**0.955**
ATT2	0.467	0.429	-0.215	0.501	0.234	-0.102	0.393	-0.228	**0.949**
ATT3	0.478	0.468	-0.262	0.488	0.219	-0.144	0.432	-0.295	**0.962**
ATT4	0.490	0.460	-0.234	0.526	0.229	-0.124	0.418	-0.264	**0.960**
Cronbach's α	0.942	0.928	0.959	0.950	0.938	0.909	0.940	0.904	0.969
Composite Reliability	0.963	0.954	0.970	0.964	0.960	0.942	0.962	0.933	0.977

Table 8

Discriminant Validity with AVEs and Construct Correlations (AVEs in Parentheses)

Constructs	ATT	DI	PAP	PC	PC_{MAR}	PPS	PU	TRUST	$TRUST_{AS}$
ATT (0.915)	0.957								
DI (0.897)	0.504	0.947							
PAP (0.889)	0.236	0.341	0.943						
PC (0.892)	-0.253	-0.432	-0.140	0.944					
PC_{MAR} (0.779)	-0.277	-0.361	-0.140	0.717	0.882				
PPS (0.844)	-0.135	-0.339	-0.122	0.717	0.574	0.919			
PU (0.870)	0.526	0.682	0.339	-0.256	-0.230	-0.169	0.933		
TRUST (0.874)	0.478	0.658	0.343	-0.573	-0.431	-0.420	0.535	0.935	
$TRUST_{AS}$ (0.893)	0.437	0.368	0.255	-0.322	-0.314	-0.185	0.392	0.504	0.945

Table 9

HTMT-Values for Assessing Discriminant Validity

Constructs	ATT	DI	PAP	PC	PC_{MAR}	PPS	PU	TRUST	$TRUST_{AS}$
ATT									
DI	0.527								
PAP	0.241	0.355							
PC	0.261	0.454	0.140						
PC_{MAR}	0.298	0.392	0.144	0.765					
PPS	0.135	0.357	0.123	0.758	0.619				
PU	0.548	0.720	0.349	0.267	0.248	0.175			
TRUST	0.503	0.702	0.359	0.603	0.464	0.444	0.568		
$TRUST_{AS}$	0.458	0.390	0.267	0.336	0.334	0.189	0.414	0.539	