

Using Conjoint Analysis to Investigate the Value of Interdependent Privacy in Social App Adoption Scenarios

Yu Pu*, Jens Grossklags†

College of Information Sciences and Technology

The Pennsylvania State University

Abstract

Applying a conjoint study approach, we conduct the first research to quantify the monetary value which users place on their friends' personal information. Utilizing the scenario of social app adoption, we further investigate the impact of the comprehensiveness of shared profile information on valuation, and vary the data collection context, i.e., friends' information is not relevant (T1), or is relevant (T2) to app functionality.

The monetary value (measured in US\$) which individuals associate with friends' full profile information in T1 (\$1.56) differs significantly from the valuation in T2 (\$0.98). However, the difference across data collection context is not significant for friends' less sensitive basic profile information (valued at \$0.23 in T1, and \$0.07 in T2). When considering the self-reported number of online friends, the average valuation for a single friend's profile information is not higher than three cents and as low as a mere fraction of a cent.

Keywords: Social Apps; Value of Interdependent Privacy; App Data Collection Context; Conjoint Analysis.

*E-mail: yxp134@ist.psu.edu

†E-mail: jensg@ist.psu.edu

1 Introduction

Within the past ten years, we have witnessed the increasing popularity of both social network sites (SNSs) and mobile platforms. In order to expand their functionalities, these platforms allow outside developers to interact with users through so-called third-party applications (or social apps). Those apps have met significant success in the marketplace ever since their emergence. As of July 2013, Google Play, one of the largest mobile app platforms, facilitated 50 billion downloads of 1 million apps [85]. Similarly, more than 75 billion times apps had been downloaded from another well-known mobile app platform, Apple Store, by June 2014 [79]. On Facebook, just to name one example, the Candy Crush Saga game app, has so far attracted more than 75 million fans worldwide [75].

Despite their worldwide popularity, social apps pose growing privacy risks to users since they collect and potentially misuse users' personal information. For example, since most users have little understanding of app permission management [8, 27], they tend to reveal more information to apps than they desire [8]. In addition, apps frequently collect more information than they need for their stated purposes, i.e., the apps are overprivileged [26, 15]. Perhaps even more troublesome, the interconnected settings of SNSs and mobile platforms have made it possible for apps to collect personal information about users' friends, who typically have limited opportunities to grant explicit consent or to prevent these practices. This problem has become known as *interdependent privacy*: the privacy of individual users does not only depend on their own decisions, but is also affected by actions of others [10].

The problem of interdependent privacy is common in social app marketplaces since social apps can request a broad range of information about users' friends. Such information includes but is not limited to friends' birthday information, friends' photos, friends' online presence, and friends' location data. Further, given the high frequency at which social apps are installed, the flow of information to outside developers can be substantial, even if only a limited number of apps request a specific type of information [83]. However, while the issue of interdependent privacy grows in practical importance, only a limited number of research studies have appeared on this subject. From an economic theory perspective, models have been developed to investigate how users make app adoption decisions in networks where privacy consequences are interdependent [69, 10]. However, to the best of our knowledge, no empirical research has investigated interdependent privacy in the scenario of social app adoption. Our work addresses this literature gap.

Our study is motivated by the theory of *other-regarding preferences* which is built on the observation that individuals' preferences consider their own material payoffs but also those of their peers [18]. In the social app context, the theory implies that individuals care about

their friends’ well-being when making app adoption decisions involving friends’ information. However, it is unknown to which degree individuals indeed exhibit such other-regarding preferences in interdependent privacy scenarios. Our first research objective is therefore to develop and execute a study to quantify the monetary value people place on their friends’ personal information in a social app adoption scenario.

The principle of *contextual integrity*, which is another key motivator for our study, suggests that privacy is context-dependent [65]. For example, when considering the installation of a birthday app that serves to remind individuals of their friends’ birthdays, users might feel less comfortable when that app also collects friends’ education information. In support of the theory, Wang et al. [84] conduct an online interaction study finding that the context-relevance of data requests influences users’ information disclosure behaviors and privacy perceptions. With our second research objective, we aim to complement these findings by quantifying to which degree *app data collection context*, i.e., the relevance of an information request, influences the valuation of friends’ information.

To address these research objectives, we follow a conjoint study approach, which allows for the measurement of users’ preferences in a trade-off scenario [35]. This methodology has been previously utilized to estimate the monetary value of individuals’ personal privacy concerns [41, 40, 58]. In this study, we adopt the conjoint study approach to quantify the monetary value of interdependent privacy. In addition, we manipulate app data collection context in two treatment conditions. An app’s request for personal information of a user’s friends is either presented as contextually relevant (i.e., it can improve the functionality of the app), or not.

The paper is organized as follows. We first discuss the research background for our work and develop hypotheses addressing the research objectives. In what follows, we provide details on our conjoint study approach and study design. We then present the analysis of the gathered data. Finally, we discuss implications and limitations of our findings, as well as present trajectories for future work.

2 Related Work

2.1 Interdependent Privacy

By investigating permission systems of apps on both SNSs and mobile platforms, a selected number of studies have highlighted the existence of interdependent privacy problems. Table 1, which is an abbreviated version from research by Wang et al. [83], summarizes the number of requests for the 16 types of permissions involving information of users’ friends in

Permission	Number of apps requesting permission	Percentage of apps requesting permission	Total times a permission is requested by apps
friends_birthday	206	2.19%	19,237,740
friends_photos	214	2.27%	13,051,340
friends_online_presence	121	1.29%	10,745,500
friends_location	104	1.11%	8,121,000
friends_hometown	21	0.22%	5,862,500
friends_work_history	86	0.91%	5,260,660
friends_education_history	14	0.15%	3,564,500
friends_activities	22	0.23%	3,448,300
friends_about_me	17	0.18%	3,328,000
friends_interests	13	0.14%	3,163,500
friends_relationships	3	0.03%	2,912,000
friends_photo_video_tags	32	0.34%	2,423,340
friends_likes	36	0.38%	2,385,960
friends_checkins	6	0.06%	1,350,000
friends_relationship_details	4	0.04%	741,000
friends_videos	2	0.02%	230,400

Table 1: Most Frequently Requested Facebook Permissions Explicitly Involving Information of Users’ Friends (Abbreviated Table from Wang et al. [83])

a sample of the 9411 most popular Facebook apps. These permissions not only cover friends’ basic information such as friends hometowns, but also involve more sensitive data such as friends birthdays, friends locations, or even friends photos and videos. Wang et al. report that while specific permissions are only collected by a small subset of apps, the impact of these data collection practices is nevertheless significant given the high frequency at which these apps are installed by users [83]. For example, friends’ birthday information is accessed by Facebook apps in almost 20 million cases. Taking into consideration the average number of friends of a typical user, the total amount of leaked information is considerable; i.e., a Pew Research Center survey found that users have on average over 300 friends [72].

Auditing over 800,000 apps in the Android play store, security firm BitDefender reports that apps frequently request information affecting users’ friends or other contacts. For instance, 10% of these apps request access to users’ contact lists, and a few of them even leak users’ call histories [52]. The same phenomena are observable in the iOS app market as demonstrated by a survey from Zscaler [70]. For example, they discover that among the 25 most popular apps 92% request access to users’ address books, and 32% include a permission for users’ calendars.

A small number of research papers have focused on the problem of interdependent pri-

vacy in the domains of online social networks and social app adoption; primarily from a theoretical perspective. Biczók and Chia [10] take a first step to study users' app adoption decisions by developing a game-theoretic 2-player model. By proposing and simulating a more comprehensive app adoption model for an entire networked system, Pu and Grossklags [69] find that even rational users who consider their friends' well-being may install apps that cause considerable interdependent privacy harm. Biasiola [9] visualize several concrete examples about the leakage of users' information when friends adopt social apps. In addition, Symeonidis et al. [81] propose a model and a privacy scoring formula to assess interdependent privacy consequences of app installations.

Interdependent privacy has also been addressed through the analysis of user data or surveys. Through analyzing users' comments posted on Facebook's official blog, Shi et al. [71] describe SNS users' interpersonal privacy concerns when their information is leaked by others' actions. Researchers have also shown how individuals' information disclosure behaviors are influenced by disclosure actions of other users and existing disclosure norms on marketplaces [11, 12]. In a survey of 199 Facebook users, Krasnova et al. [57] elicit levels of privacy concern regarding the release of 38 different information items including data about friends. When comparing the concern levels for data collection about the user and her friends, Krasnova et al. [57] find little support for the assumption of "privacy egoism." In their data, users frequently appear more concerned about data release involving their friends' data.

Although these studies highlight important aspects of the interdependent privacy problem, to the best of our knowledge, no empirical research has been proposed or conducted to quantify the monetary value of interdependent privacy.

2.2 Valuation of Privacy

To understand privacy decisions which are made in a deliberate fashion, researchers can take the perspective of the privacy calculus [87]. Viewing privacy as an economic good [55], the privacy calculus perspective posits that individuals are consciously conducting a cost-benefit analysis when they are confronted with an information disclosure scenario [22, 73, 44, 19], which requires a balancing of perceived costs against perceived benefits. Ideally, this deliberation should lead to the selection of the most favorable outcomes [80]. Analyses of survey responses suggests that individuals take multiple factors into consideration (such as trust in a website) when they decide to partially sacrifice privacy for benefits arising from data disclosure [22].

The notion of privacy calculus is relevant for diverse empirical approaches in the domain

of privacy valuation research. In particular, we describe research works that put individuals in implicit or explicit trade-off scenarios, such as surveys, field experiments, discrete choice experiments, and conjoint analyses, and shed light at the value individuals place on their own personal privacy.

Several studies directly elicit users' valuation of their privacy in survey settings. Acquisti and Grossklags [2] elicit individuals' valuations for information items with varying degrees of sensitivity (e.g., name, and personal health history), and show how this valuation can be manipulated through reframing of the scenario. Chellappa and Sin [14] investigate the trade-off between the value of personalization and privacy concerns in an online context. Wathieu and Friedman [86] use responses from a survey including questions on the dissemination of personal information in a commercial context to better understand the value of personal information. Similarly, Bauer et al. [5] and Spiekermann et al. [77] aim to understand how Facebook users value privacy by asking them how much they would be willing to pay to protect their SNS profile information.

Further, a variety of context-rich experiments tackle the measurement challenge to understand the value of personal information. Conducting experiments where participants have to make purchasing decisions that affect their privacy, some studies find consumers are willing to pay a premium in order to purchase more privacy-friendly products [47, 82, 6]. In contrast, Spiekermann et al. [76] demonstrate how actual behaviors differ from reported privacy concerns, and how even sensitive information is given away in exchange for potentially better recommendations during online shopping. By comparing users' choices for selling and protecting personal information, Grossklags and Acquisti [36] demonstrate that on average, the amount of money users are willing to accept to reveal their information is higher than the amount they are willing to pay for protecting their privacy. Similar results were provided by Acquisti [3]. Applying a second-price auction mechanism, Huberman et al. [43] measure the amount of money individuals would demand to reveal their weight or height information. Using a related method, Danezis et al. [20] evaluate how the monetary valuation for location data differs for individuals from various European Union countries. Christin et al. [16] show in an online experiment how even modest monetary enticements lead many users to sacrifice the security of their computing systems which typically contain substantial amounts of personal information.

A different set of studies investigate the valuation of privacy by applying discrete choice experiments. For example, Potoglou et al. [68] use this method to estimate the value of personal information in three real-life contexts and situations. Their findings reveal that respondents have a low willingness to pay for control over their personal data. However, they also suggest that the extent of data collection of personal information by third parties is the

most important factor affecting users’ online retailer choice. Applying the same method, Krasnova et al. [56] and Egelman [24] study users’ concerns about the release of personal information when presented with single sign-on mechanisms such as Facebook Connect.

Closely related to our work, conjoint analysis has been used to investigate individuals’ privacy valuations of their own information. Hann et al. [41] and Hann et al. [40] explore the trade-off between the benefits and costs of revealing personal information online. They also derive the monetary value of individuals’ personal information. Similarly, conjoint analysis allows Krasnova et al. [58] to infer the Euro value respondents associate with their own information on SNSs.

In aggregate, these studies utilize different methods to quantify the value individuals place on their own personal information. To the best of our knowledge, no published research has investigated the monetary value of interdependent privacy. Our work further applies the conjoint analysis method to the scenario of social app adoption. We extend previous studies by considering interdependent privacy as a key attribute, as well as by introducing and comparing different app data collection contexts.

3 Hypotheses

Experimental studies provide substantial evidence that individuals deviate from pure self-interest in many decision-making situations [38, 50, 29, 25, 7]. Individuals exhibit other-regarding preferences [18, 46]; they care about the well-being and welfare of others, behave altruistically, and reciprocate kindness of others [46]. It has been shown that such preferences are particularly strong when choices directly determine others’ welfare [78]. Applied to our context of interest, we anticipate that social app users demonstrate other-regarding preferences, and hence care about their friends’ privacy (as an important component of overall well-being).

In addition, previous research reported that computer-mediated interactions help individuals to develop and maintain online social capital [39, 53]. Social capital, which refers to immaterial resources accumulated through the relationships among people [17], is correlated with positive outcomes, such as emotional support [4, 42], increased chances of exposure to diverse ideas [67], and opportunities for access to non-redundant information [31]. As part of online communities, social app users have likely developed some online social capital. In order to maintain relationships and continue to enjoy the benefits of social capital, app users would likely carefully evaluate actions that might harm other community members. In this manner, we expect that social app users demonstrate concerns over their friends’ privacy when they are presented with interdependent privacy decision-making situations.

Hypothesis 1: *Social app users place a positive (non-zero) value on their friends' personal information suggesting the existence of other-regarding preferences in the scenario of interdependent privacy decision-making.*

Individuals, including social app users, express privacy concerns when they are faced with decisions about disclosing their own personal information [74, 74, 61]. In addition, as argued above, app users are expected to value their friends' personal information. Further, while in many decision-making scenarios individuals do not act purely based on self-interest, the strength of other-regarding preferences may vary for individuals and across contexts [38, 50, 29, 25, 7]. It is, therefore, an empirical question whether personal privacy concerns trump the concerns for friends' privacy.

In particular, while individuals build online social capital through their interactions [39, 53], research has shown that the social ties (i.e., intensity of relationships) between SNS users, e.g., on Facebook, are comparatively weak [46]. As opposed to strong ties which occur between people who trust each other and whose social circles are tightly overlapping [31], weak ties are often connections between individuals with different national, linguistic, or social backgrounds [31, 46]. We expect that such weak ties translate in comparatively lower valuations for friends' information in comparison to the valuation for a user's own information.

Hypothesis 2: *Social app users express a higher valuation for their own personal information than their friends' personal information when making app adoption choices.*

Contextual integrity, which ties adequate protection for privacy to norms which have developed in specific situations, demands that information gathering and dissemination should be appropriate to a particular context [64]. Further, prior research has found that privacy perceptions are context-dependent [65], that is when individuals feel that contextual integrity is violated (e.g., irrelevant or inappropriate information is collected or disseminated), their concerns for privacy tend to increase. For example, Wang et al. [84] find that people are typically unconcerned about giving away their friends' birthday information to a birthday app, but quickly become uncomfortable when that app also attempts to access information unrelated to its stated purpose. In our study, we manipulate app data collection context in two treatment conditions. An app's request for personal information of a user's friends is either presented as contextually relevant (i.e., it can improve the functionality of the app), or not. We expect that social app users express a higher valuation for friends' personal information when irrelevant information is requested by an app.

Hypothesis 3: *Compared with the case where information collected about users' friends is app-relevant, social app users value their friends' personal information higher when contextual integrity of apps is violated.*

4 Methodology

4.1 Conjoint Study Design

Conjoint analysis is an experimental approach that aims to uncover the hidden rules individuals use to make trade-off decisions between products or services. It assumes that consumers view a product as a bundle of certain features (*attributes*), which have different values (*levels*) [34]. In a conjoint study, researchers would then test individuals' preferences for multiple versions of a product (*profiles*), which are formed by combining different attributes and levels. Through applying statistical analysis on these preferences, researchers are able to quantify the value individuals place on each attribute level, as well as to understand the role each attribute plays on individuals' decision-making [48].

In our case, users view a social app as a combination of multiple app features. For example, the fact that apps collect information about users' friends can be understood as an app feature. Correspondingly, its levels will be the different types of friends' information an app collects. Then, through the analysis of how people evaluate versions of a social app, we are able to infer how certain factors, especially revealing friends' personal information, affect users' choice of an app.

In the following subsections, we discuss how we select app attributes, attribute levels, as well as the different app versions for the conjoint study. We also provide details on the mathematical method to estimate conjoint parameters.

4.1.1 Determination of Attributes

As suggested by Green and Krieger [32], we conducted in-depth semi-structured interviews with social app users to determine app attributes for the conjoint study. To reduce the bias of our study, we paid particular attention to the recruitment process. Since our study focuses on privacy issues, we targeted individuals with unrelated employment backgrounds. In addition, we aimed for a balanced sample by interviewing both individuals with and without technical backgrounds. Individuals with expertise in fields such as computer science and information science were considered to have a technical background, while those without such experiences were regarded as non-technical. Conducting the interviews helped us to confirm appropriate attributes, thereby reducing the likelihood of model misspecification.

Note here, we define social apps as any apps that can be installed or used from mobile app stores and/or social network app platforms. In other words, social apps in our study not only refer to apps used on smartphones and tablets, but also web-based apps, e.g., on predominantly web-based SNSs. As part of the interview, we asked participants to identify

factors that affect their decisions to install an app. Since our study mainly focuses on privacy issues, and previous research indicated that online users exhibit privacy concerns over both their own information [58, 61] and their friends’ data [57], we directly probed whether these different dimensions of privacy would be a concern if participants did not mention anything about information collected by apps. Prior survey research has demonstrated that individuals often cannot initially recall important privacy-relevant issues that appear obvious to them after being prompted [1].

In total, we interviewed 18 individuals, among which 10 had technical expertise and 8 had non-technical backgrounds. Most interview participants (17 out of 18) singled out app price (*price*) as one of the key factors that influences their choice of a social app. In addition, those 17 participants identified the level of an app’s popularity among friends (*network popularity*) as an important decision-making factor. This is in line with existing research documenting that positive network effects are an important motivator for people to use technologies [37]. In addition, some interviewees (13 out of 18) reported they care a lot about whether their own information was collected by apps (i.e., *own privacy*). Without being directly probed, 4 among them immediately self-reported that their *own privacy* is a concern. “I will not consider the app if it tries to collect my cellphone number”, said one interviewee. “I always post my fake information on the Facebook so that when I am installing some Facebook apps, they won’t get access to my real information”, revealed another respondent. This is consistent with findings that people are typically reporting a moderate to high concern for personal privacy [61]. Realizing that some apps are collecting the information of users’ friends, our interviewees (13 out of 18), indicated that the type and procedure for the collection of friends’ information by an app (i.e., *friends’ privacy*) matters to them. Among them, 2 participants immediately self-reported this as a factor affecting their app choices. One participant said, “As they are my friends, I should also care about their privacy.” This finding is supported by the theory on other-regarding preferences which explains how people care about other individuals’ well-being [18]. In addition, other factors, such as the level of interest in an app, and whether a third-party app developer is trustworthy were also identified by a minority of our interviewees (7 out of 18, and 1 out of 18, respectively). Overall, interviewees’ responses indicate that price, popularity among friends (i.e., network popularity), users’ own information collected by apps (i.e., own privacy), and the information apps collect about users’ friends (i.e., friends’ privacy) constitute four important dimensions of an app-selection decision and are suitable for the usage in a conjoint study.

4.1.2 Determination of Attribute Levels

In this section, we explain the levels chosen for the four attributes. The interview results helped us to determine levels for *price*. Although, most apps which our interviewees had previously installed were free of charge, our participants were willing to pay a small amount for installing apps they strongly desired. Most of them thought a reasonable price for an app would be about \$2.00. Based on this finding, two different levels for *price* were selected: “\$0.00” and “\$1.99”. We also assigned two levels for *network popularity*. These two levels were: “5%” and “25%”. Note, the number here represents the percentage of a user’s friends who have already installed the app, not the popularity across all network users. By investigating app permission systems, previous researchers found that Facebook apps typically collect a user’s basic information such as user name and ID [83]. Besides accessing basic information, some apps request additional information such as the user’s birthday and location information [83]. However, we did not rule out the possibility that some apps may collect no information about users. Therefore, we chose three levels for *own privacy* to indicate which information is collected by an app about a user: “none”, “basic profile” and “full profile”. “None” indicates that the app does not collect any information about users. “Basic profile” includes a user’s name, profile picture, gender, user ID, number of user’s friends, and any other information the user made public. “Full profile” includes user’s email-address, birthday, all photos, and location information, plus all the information included in “basic profile”.

Besides collecting users’ own information, apps frequently collect information about users’ friends; however, not all apps do aim for friends’ information. Similarly, the range of friends’ information an app collects is subject to substantial variation [83]. Some apps only collect friends’ basic information, while others gather more sensitive information about friends. Hence, in our study, *friends’ privacy* differed in three levels: “none”, “basic profile”, and “full profile”. “None” means that the app does not collect any information about friends. “Basic profile” indicates that the app solicits friends’ names, profile pictures, gender, user IDs, number of friends’ friends, and any other information friends have made public on their profiles. Similarly, “full profile” includes friends’ “basic profile”, as well as friends’ email-addresses, birthdays, all photos, and location information. Table 2 summarizes the attributes and levels used in our conjoint study.

4.1.3 Selection of Conjoint Analysis Method and Number of Stimuli

Conjoint analyses can be conducted in several ways. Among them, the three most popular methods are full-profile conjoint analysis, choice-based conjoint analysis, and adaptive

Attributes	Attribute Descriptions	Attribute Levels
Price	Price of the app	\$0 \$1.99
Network Popularity	Percentage of a user’s friends who installed the app	5% 25%
Own Privacy	Information the app collects about a user	None Basic profile Full profile
Friends’ Privacy	Information the app collects about a user’s friends	None Basic profile Full profile

Table 2: Summary of Attributes and Levels

conjoint analysis. For our study, we selected the full-profile method. First, full-profile conjoint analysis is considered preferable in the case of less than six attributes, while adaptive and choice-based method are more suitable for studies that have more than six attributes [62]. Second, when it comes to the choice-based conjoint method, although it can be easily implemented, the results have traditionally been analyzed at the aggregate level. However, in the context of privacy, measuring the preferences of individuals is critical [1]. Therefore, aggregate-level models, which assume respondent homogeneity, cannot be as accurate as individual-level models [66]. Further, while the Bayesian estimation for choice-based conjoint analysis permits estimating individual-level utilities [49], such estimated utilities are not as representative as the individual results that are directly obtained by applying the full-profile conjoint method.

When applying the full-profile approach, respondents are asked to rank a set of product profiles (*stimuli*) [33]. In our case, the stimuli are different app versions that are formed by combining different levels of the four attributes. The attributes and levels in Table 2 give rise to a total of 36 ($2 \times 2 \times 3 \times 3$) stimuli. Clearly, this ranking task would become too complex for individuals to complete in a meaningful way. In order to solve this problem, we utilized the SPSS Conjoint 22 package to reduce the number of profiles used in our study. More specifically, SPSS Conjoint 22 uses a fractional factorial design to generate an orthogonal array. The resulting array, which is a suitable fraction of all possible combinations of the factor levels, is designed to capture the main effects of each factor level. By applying this approach, we reduced the design from 36 possible profiles to 9 profiles.

4.1.4 Estimation of Conjoint Model

Main effects analysis of variance (ANOVA) is the basic estimation procedure underlying the full-profile conjoint analysis [40]. This procedure computes the utility of each attribute level such that the rank ordering of utility sums of all levels in a profile equals the actual rank ordering of that profile. Thus, the predicted rank ordering for profile j can be specified as follows:

$$R_j = \beta_0 + \sum_{i=1}^T \beta_i X_{ij} + \varepsilon_j \quad (1)$$

where R_j is the ranking of profile j , β_0 is a utility constant, T is the total number of attribute levels, and β_i is the coefficient (utility value) to be estimated for attribute level i . X_{ij} is a $\{0, 1\}$ variable that equals 1 if profile j has attribute level i , otherwise it equals 0. ε_j is a stochastic error term.

To estimate the utility value of each attribute level, we use the SPSS Conjoint 22 package. More specifically, based on each participant’s rankings for the 9 app versions, the utility of each attribute level is estimated on an individual basis.

4.2 Design of the Conjoint Study Experimental Survey

We conducted a web-based, between-subject online experiment (with treatments T1 and T2; see below) by using a combination of Qualtrics, an online survey software, and Amazon Mechanical Turk (MTurk), a recruitment source that is very popular for conducting online user experiments [30]. We constructed the survey questions and implemented the experimental conditions on Qualtrics. We then recruited participants from MTurk and asked them to access the Qualtrics link for our experiment.

4.2.1 Treatments

The theory of contextual integrity demands that information gathering and dissemination should be appropriate to that context [64], and prior research discovered that individuals’ privacy concerns are affected by whether or not information requests are context-relevant [65, 84]. For example, Wang et al. [84] find users are typically unconcerned about giving away their friends’ birthday information to a birthday app, but become uncomfortable when that app also tries to get access to information unrelated to its stated purpose. Motivated by the theory of contextual integrity and empirical findings, we aim to understand the effect of app data collection context on how app users value their friends’ personal information.

Individual privacy perceptions may also depend on a variety of other contextual factors [65]. In particular, while our study considers on an abstract level whether an information request is app-relevant or not, we do not test whether concrete types of app purpose and presentation have a measurable impact on users' privacy valuation. Our objective is to conduct a conjoint study with a generic description which applies to a variety of practical scenarios. Future work may target whether specific usage contexts or app purposes influence the valuation of personal information.

We varied the data collection context in our experiment by randomly placing participants in one of the following two treatment scenarios which were included and highlighted in the instructions immediately preceding the ranking task:

T1: The information the app collects about user's friends does not improve usability or functionality of the app.

T2: The information the app collects about user's friends improves usability or functionality of the app.

4.2.2 Procedures

Next, we discuss the procedures of the experiment.

At the beginning of the study, each participant was presented with a consent form, which described the aim of the study (i.e., to rank versions of a social app) without highlighting the topic of privacy, the high-level procedure they were going to follow, and additional related information. After consenting to the terms of the task, participants proceeded to the main part of the study.

The main part of the online study can be further divided into three parts. A short survey in the first part collected participants' demographic information such as gender, age, education level and income level. In the second part, participants proceeded to the core task. They were asked to use a drag-and-drop interface to rank the 9 app versions. Figure 1 is a screenshot of our conjoint study interface. The ranking task was preceded by a detailed explanation of the attributes and levels, and the interface. Further, the treatments were administered as part of the instructions for the ranking task. In the third part, each participant was asked to complete a post-experimental survey, which was used to further understand people's privacy preferences and app-adoption behaviors. After completing the survey, participants exited our study, and entered a completion code on MTurk to receive payment.

Below is a list of 9 different app versions, which differ in the 4 product dimensions: price (**Price**), percentage of your friends who have installed the app (**Popularity**), information the app collects about you (**Own privacy**), and information the app collects about your friends (**Friends' privacy**). Please rank them in order of preference from 1 to 9 (1 = most preferred, 9 = least preferred).

You can return to the previous page to study the instructions in more detail.

Price: \$0	Popularity: 5%	Own privacy: None	Friends' privacy: Basic Profile	1
Price: \$0	Popularity: 5%	Own privacy: Basic Profile	Friends' privacy: Full Profile	2
Price: \$0	Popularity: 25%	Own privacy: Full Profile	Friends' privacy: None	3
Price: \$1.99	Popularity: 5%	Own privacy: Full Profile	Friends' privacy: Basic Profile	4
Price: \$0	Popularity: 25%	Own privacy: Basic Profile	Friends' privacy: Basic Profile	5
Price: \$0	Popularity: 5%	Own privacy: None	Friends' privacy: None	6
Price: \$1.99	Popularity: 5%	Own privacy: Basic Profile	Friends' privacy: None	7
Price: \$1.99	Popularity: 25%	Own privacy: None	Friends' privacy: Full Profile	8
Price: \$0	Popularity: 5%	Own privacy: Full Profile	Friends' privacy: Full Profile	9

Figure 1: Screenshot of Experiment Interface (Drag and Drop Interaction)

4.2.3 Participant Pool and Payment

We recruited participants through MTurk. Eligible participants included social network site account holders and/or smart phone users. In addition, eligible participants should have previously installed at least one app on a social network site and/or smart phone. Furthermore, we restricted participants to Turkers who had completed over 50 Human Intelligence Tasks (HITs) with a HIT approval rating of 95% or better, as well as those who had United States IP addresses. We put these restrictions in place to ensure that the participants were familiar with social apps, and to aim for higher response quality. We paid \$1.00 to each participant after survey completion. We further prevented MTurk users from participating multiple times by controlling for their MTurk IDs.

5 Data Analysis

5.1 Participant Data

Data collection was conducted in April 2015 when we collected data for a total of 400 Mechanical Turk users. Among them, 198 users were assigned to T1, and 202 users were assigned to T2. A key advantage of Mechanical Turk is that the demographic mix of participants is typically more diverse than university student convenience samples [45, 51]. Further, the ratio of payment to average quality of the responses is considerably lower compared to traditional laboratory studies [63]. However, studies provide evidence that a substantial

fraction of Mechanical Turk users tend to shirk on assignments or even use automated bots to complete surveys [23]. In particular, tasks of an increased complexity (such as full-profile conjoint analyses) may result in lower attention of some participants, and therefore require careful filtering [23, 30].

The orthogonal design which we utilize in the full-profile conjoint analysis allows for a straightforward identification of low effort submissions that violate the most basic consumer behaviors. For example, consider two app versions, say app A and app B, that have the same level of popularity and collect the same information about friends. However, they differ in price and in what information they collect about users. Specifically, app A accesses users' full profiles and costs \$1.99, while app B does not collect any information about users and is free of charge. We then expect that reasonable consumers would prefer app B to app A. If app ranking results indicate otherwise, we believe such results demonstrate irregular consumer behaviors and thus exclude them from analysis. This exclusion criterion segmented the subject population effectively, since our further investigation showed that these excluded individuals did not exercise enough care with the completion of the ranking task and spent only an insufficient amount of time on the task which essentially led often to random responses.

Out of 198 Turkers in T1, responses from 100 Turkers were used for analysis. For T2, data from 101 Turkers of the 202 total responses showed sensible consumer behaviors. The exclusion criteria were applied consistently across the two treatments. Chi-square tests revealed that participants whose responses were excluded from the study and participants whose responses were kept did not differ significantly in terms of gender (T1: $p = 0.24$; T2: $p = 0.57$), age (T1: $p = 0.27$; T2: $p = 0.30$), education level (T1: $p = 0.73$; T2: $p = 0.93$) or income level (T1: $p = 0.92$; T2: $p = 0.53$). In addition, Chi-square tests demonstrated that answers in the post-experimental survey did not systematically differ across the two groups, indicating that participants in these two groups have similar privacy preferences. Among the 100 participants in T1 whose responses are used in our analysis, 55% were male and 45% were female; they belonged to a wide range of age categories (from 18 to over 50), and covered a wide range of education levels (from less than high school to PhD) and yearly income levels (from less than \$25,000 to more than \$100,000). The final T2 sample consisted of 62.4% male and 37.6% female subjects. Similar to the final sample in T1, these subjects belonged to a wide range of age categories and covered a wide range of education levels as well as income levels.

	Value	
	T1	T2
Pearson’s r	0.99	0.99
Kendall’s tau	0.94	1.00

Table 3: Correlation Between the Observed and Estimated Preferences

Attributes	Importance Values	
	T1 (%)	T2 (%)
Price	37.84	43.47
Network Popularity	17.00	13.80
Own Privacy	26.62	26.96
Friends’ Privacy	18.55	15.77

Table 4: Averaged Importance Values

5.2 Analysis of Part-worths and Importance Values

In this section, we describe final utilities (i.e., part-worths) of attribute levels, which are represented by β_i in Equation 1, and the relative importance of each attribute. Part-worths reflect the attractiveness of an attribute level, while relative importance of an attribute allows us to draw conclusions about the role that the attribute plays in users’ app selection decisions.

The goodness-of-fit data for the estimated model is reported in Table 3. Models in both treatments have a Pearson’s r value of 0.99 (see Table 3), which indicates a high level of correlation between observed preferences and estimated preferences. Similarly, high values of Kendall’s tau in both treatments (T1: 0.94; T2: 1.00) also demonstrate the goodness-of-fit of these two models.

Part-worths and relative importance values are shown in Figure 2 and Table 4, respectively. Rankings of the four attributes, which are based on their importance to app adoption decisions, are consistent across the two treatments. With an importance score of 37.84% in T1 and 43.47% in T2, *price* is the most important factor in the choice of a social app. As expected, utility decreases when app price increases. *Own Privacy*, which has an importance value of 26.62% in T1 and 26.96% in T2, is the second most important factor in both treatments. We can observe that utility is higher when less information is collected about the user. This is consistent with the finding in [59], which confirms the link between privacy concerns and adoption behaviors.

Friends’ privacy is the third most important factor with 18.55% in T1 and 15.77% in T2. We find that utility decreases when the app collects more information about friends. This is an important finding, as it shows that users indeed take friends’ privacy into consideration in the choice of a social app, and provides evidence for the relevance of other-regarding preferences theory in the app selection context and privacy decision-making [18]. At the same time, the observation that own privacy importance scores are higher than those for friends’ privacy indicates that individuals care more about their own privacy than others’ privacy. With an importance score that is slightly lower than the importance value of friends

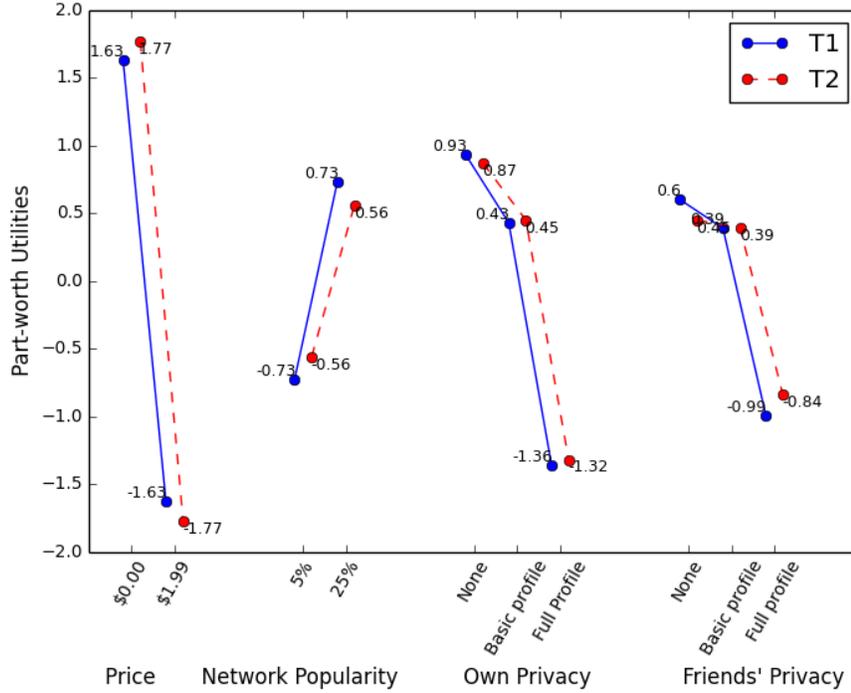


Figure 2: Averaged Part-worth Utilities

privacy, *network popularity* is the least important factor. As expected, part-worth increases with an increasing level of network popularity.

5.3 Analysis of Monetary Values of Level Changes

In this section, we first offer an overview of users’ utility changes when an app switches from one level of an attribute to another level. Note, the utility change here is the averaged utility change across the entire subject sample *within* each treatment. We then compute monetary equivalents for each change in the attribute levels. The latter step is important to provide insights regarding the economic value individuals place on attribute level changes; including level changes of *own privacy* and *friends’ privacy*.

We present our analysis results in Table 5. The “Utility Change” column presents the exact values for utility changes of corresponding level changes. For example, we observe that in both treatments, the biggest utility change occurs when an app’s *price* switches from “\$0.00” to “\$1.99”. This utility change highlights users’ aversion of shifting from a free app to a costly version. The second biggest utility change is associated with the level change related to *own privacy*, i.e., when own privacy changes from “none” to “full profile”. Note that the utility change is negative indicating that users prefer less privacy-invasive apps.

Attributes	Level Change	Utility Change		Dollar Value		P-value	
		T1	T2	T1	T2	T1	T2
Price	\$0.00 \Rightarrow \$1.99	-3.25	-3.54	-1.99	-1.99	-	-
Network Popularity	5% \Rightarrow 25%	1.46	1.12	1.12	0.77	-	-
Own Privacy	None \Rightarrow Basic profile	-0.49	-0.42	-0.55	-0.48	0.00	0.00
	Basic profile \Rightarrow Full profile	-1.79	-1.78	-1.76	-1.56	0.00	0.00
	None \Rightarrow Full profile	-2.28	-2.20	-2.31	-2.04	0.00	0.00
Friends' Privacy	None \Rightarrow Basic profile	-0.21	-0.06	-0.23	-0.07	0.03	0.15
	Basic profile \Rightarrow Full profile	-1.39	-1.22	-1.33	-0.91	0.00	0.00
	None \Rightarrow Full profile	-1.60	-1.28	-1.56	-0.98	0.00	0.00

Table 5: Utility Change and Monetary Value of Change

Next, we discuss how to translate the utility changes between levels to monetary values. We observe that the total change in final utility from “\$1.99” to “\$0.00” implies a change of $3.25/1.99 = 1.63$ units of final utility per dollar in T1, and $3.54/1.99 = 1.78$ in T2. We then use these values to calculate the dollar equivalents for level changes in other attributes. We present these results in the “Dollar Value” column in Table 5.

For example, the change from collecting basic profile information to accessing no information from friends is worth \$0.23 for participants in T1, and \$0.07 for participants in T2. This means, participants in T1 are ready to pay 23 cents in order to prevent an app from collecting their friends’ basic profile, and respondents in T2 are likely to pay 7 cents for the same reason. Similarly, in order to prevent apps from accessing full profile information of friends rather than merely their basic profiles, participants in T1 are willing to pay \$1.33. For the same reason, respondents in T2 would pay \$0.91. Hence, in total, the monetary value individuals place on full profile information collected from friends is \$1.56 in T1, and \$0.98 in T2. When it comes to users’ *own privacy*, the amount participants in T1 and T2 are willing to pay for preventing access to their basic profile information is \$0.55 and \$0.48, respectively. In addition, social app users in T1 and T2 are willing to pay \$1.76 and \$1.56, respectively, in order to protect their sensitive full profile. Therefore, the dollar value that participants place on their own full profile information is \$2.31 in T1, and \$2.04 in T2.

In the next step, we conduct one-tailed one-sample *t*-tests to better understand whether social app users place positive values on *own privacy* and *friends’ privacy*, and present the results in the “P-value” column in Table 5. Importantly, we observe that the valuations for own and friends’ privacy are significantly positive in both treatments. An exception is the valuation for friends’ basic information in T2. This result suggests that basic profile information of friends is considered to have a very low value when the information request is contextually relevant. Therefore, **Hypothesis 1** is partially supported.

5.4 Comparison of Dollar Values of Own Privacy and Dollar Values of Friends’ Privacy

In this section, we investigate the differences between the dollar value that individuals place on their *own privacy* with the value of their *friends’ privacy* (see Table 6).

Level Change	Privacy Dollar Value in T1			Privacy Dollar Value in T2		
	Own	Friends	P-value	Own	Friends’	P-value
None \Rightarrow Basic profile	-0.55	-0.23	0.04	-0.48	-0.07	0.00
Basic profile \Rightarrow Full profile	-1.76	-1.33	0.01	-1.56	-0.91	0.00
None \Rightarrow Full profile	-2.31	-1.56	0.01	-2.04	-0.98	0.00

Table 6: Comparison of Monetary Values of Own Privacy and Friends’ Privacy (*t*-test; one-tailed)

From Table 6 we can conclude that (in both T1 and T2) the dollar values participants place on their *own privacy* are significantly higher than the corresponding values for their *friends’ privacy*. These observations support the common-held belief that individuals care more about themselves than others.

Importantly, when we refer to the monetary values of *friends’ privacy*, we refer to the dollar values an individual places on the profile information of *all* her friends. Next, we investigate how an individual values the profile information of a single (average) friend. We determine the average value of the profile of a single friend by dividing the dollar values of all friends’ profile information by the self-reported number of friends an individual has (which is data we collected in the post-experimental survey). We show the results in Figure 3, where the dollar values for a single friend’s private information are represented by (small) black bars.

We observe from Figure 3 that the value that participants associate with the information about an average friend is very small. For example, in both treatments, the monetary value an individual places on her average friend’s full profile information is less than \$0.03, and the value for basic profile information is considerably less (i.e., as low as \$0.001). A comparison with the high value that users associate with their *own privacy* (shown with white bars) supports the notion that users should be considered “privacy egoists” with respect to the privacy of individual friends. Considering recent research which shows that most friendship ties are weak on social networks [21], the observation that social app users only care a little (on average) about the privacy of each of their friends is perhaps not surprising. However, our study quantifies the strength of this effect. Therefore, **Hypothesis 2** is supported.

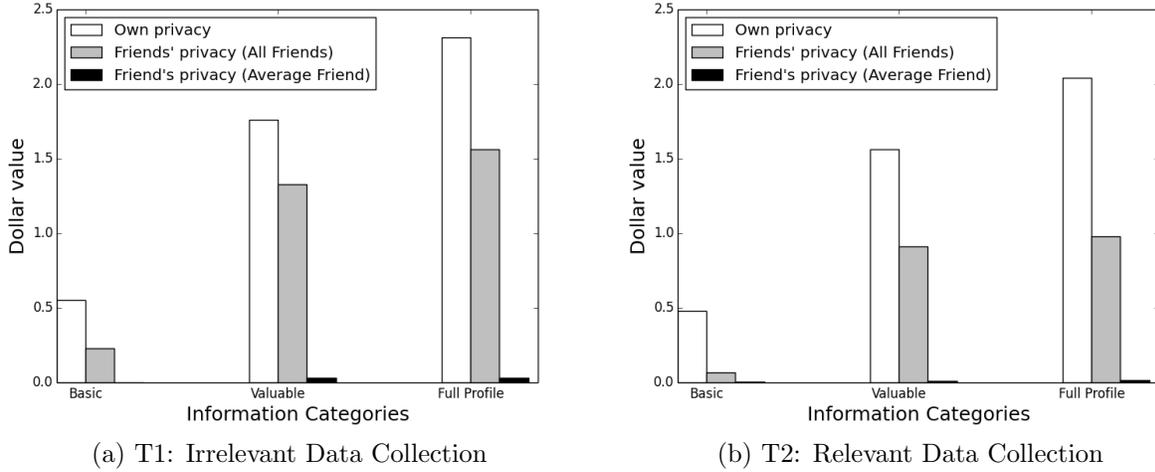


Figure 3: Privacy Dollar Value

5.5 Comparison of Dollar Values of Privacy-Related Level Changes between Treatments

In the next step, we examine whether dollar values of level changes related to *friends' privacy* and *own privacy* differ *between* the two treatments. In other words, we want to test **Hypothesis 3**, that is whether app data collection context affects how participants value the privacy of their friends.

Attributes	Level Change	Dollar Value		P-value
		T1	T2	
Own Privacy	None \Rightarrow Basic profile	-0.55	-0.48	0.36
	Basic profile \Rightarrow Full profile	-1.76	-1.56	0.23
	None \Rightarrow Full profile	-2.31	-2.04	0.24
Friends' Privacy	None \Rightarrow Basic profile	-0.23	-0.07	0.12
	Basic profile \Rightarrow Full profile	-1.33	-0.91	0.04
	None \Rightarrow Full profile	-1.56	-0.98	0.02

Table 7: Comparison of Monetary Values between Treatments (*t*-test; one-tailed)

We show dollar values and significance values in Table 7. We observe that the absolute monetary value of each level change in T1 is higher than in T2. This suggests that individuals in T1 place more value on privacy, both own privacy and friends' privacy, than individuals in T2. These observations are anticipated, since T1 represents the less desirable case where friends' personal information collected by an app is not related to its functionality, while T2 represents the situation where such information can improve the functionality of the app. We further think it is reasonable that participants in T1 place a slightly higher valuation

on their *own privacy* than their counterparts in T2 since they may become more vigilant to privacy issues after observing the unjustified request for friends’ personal information by the app.

However, the statistical analysis indicates that *only* level changes for friends’ information differ significantly between T1 and T2. First, the monetary value participants place on *friends’ privacy* level change from “basic profile” to “full profile” is significantly higher in T1 compared to T2. Importantly, this level change indicates that an app is requesting friends’ particularly valuable data, such as their photos and location data. The significant difference about this value indicates that individuals are more sensitive to friends’ valuable data than basic data. The absence of a significant effect for basic profile information is perhaps unsurprising since such less sensitive data can be more easily obtained from other sources. As such, although participants positively value friends’ basic profile information, this valuation is not significantly affected by app data collection context. In addition, the value for full profile information (compared to “none”) is significantly different between T1 and T2, which is expected since this level change includes the particularly sensitive data of friends. These findings demonstrate that app data collection context significantly affects the value that individuals place on their friends’ sensitive full profile information, but not basic profile information, or their own personal information. Thus, **Hypothesis 3** is partially supported.

6 Conclusions

To the best of our knowledge, this paper presents the first approach to quantify the monetary value of interdependent privacy with a particular focus on app adoption scenarios. In addition, motivated by the principle of *contextual integrity* [65], we examine the effect of *app data collection context* on the value individuals place on their friends’ personal information by introducing two treatments into our study: (T1) the case where friends’ personal information does not improve an app’s functionality, and (T2) the case where friends’ personal information improves an app’s functionality.

Regarding the *relative importance* of the different app features we studied, we show that *friends’ privacy* is the third most important factor in both treatments (after price and own privacy). We further find that the *monetary value* individuals associate with their own full profile information is \$2.31 in T1 and \$2.04 in T2. When studying the value individuals place on friends’ information, we observe that full profile information of friends is valued at \$1.56 in T1, and \$0.98 in T2. These findings initially suggest that users pay considerable attention to *friends’ privacy* when they reason about app adoption decisions in a survey

setting. However, when considering the self-reported number of friends, then the value associated with the profile of an average single friend is a very small fraction of the value associated with a user’s own personal information. Our results suggest that users behave like “privacy egoists” when making social app adoption decisions.

When examining the impact of the different data collection context treatments more closely, we find that the monetary values participants place on both their *own privacy* and *friends’ privacy* are higher in T1 than in T2. However, only the treatment difference for the monetary value of friends’ more valuable full profile information such as friends’ photos and locations is statistically significant. This suggests that the protection of basic profile information is considered of low importance to participants, but that unwarranted information requests for sensitive information are detrimental to the positive evaluation of an app.

Our study contributes to the policy discussion on app privacy. The relatively low value that users associate with the information of their individual friends (even if notified that the data collection serves no legitimate purpose) suggests that such data sharing should be more rigorously limited, or involve the affected friends of a user more directly in the decision-making process. As an initial step, privacy advocates should encourage a revision of apps’ privacy notice dialogues so that they can *more effectively* notify users of apps’ data collection practices (i.e., whether or not apps can gain access to their friends’ personal information, what types of friends’ information they collect, and how many individuals are affected by the data sharing). In addition, our findings highlight the importance of embedding technologies in the app adoption process which help individuals understand whether apps’ practices of collecting data, especially the data about friends’ full profile, are necessary [28]. We believe that, collectively, these interventions will help app users to make more informed decisions about their own and their friends’ data.

There are several limitations of this study, some of which provide insights for further research. First, privacy concerns were given more prominence in our experiment than they would likely receive in practical decision-making scenarios. It is possible that this pronounced focus elevated the measured monetary valuations. Taking this into account, the very low valuations for an individual friend’s personal information stand out even more. Recent research focused on the idea to measure privacy in survey settings *without asking about it* [13]. We consider it as a fruitful research direction to implement similar ideas in experiments and trade-off scenarios such as conjoint studies. Second, our investigation about the influence of context on privacy valuations is limited to the measurable impact of app data collection context (i.e., whether the requested data can improve the functionality of the app). However, other factors that might have considerable impact on users’ privacy concerns and behaviors [65], such as app category, are not part of the study. To better understand the relationship

between other contextual factors and privacy valuations, more context dimensions should be evaluated in follow-up research. Further, in this study, we restricted our participants to Mechanical Turk users with United States IP addresses. However, previous research shows that individuals in different regions exhibit different attitudes towards privacy issues [60, 82]. Therefore, in the next step, we plan a cross-cultural study with participants from different countries to further evaluate the robustness of our study results. Finally, although our study confirms the existence and evaluates the strength of interdependent privacy concerns, we are unaware of the determinants of such concerns. Prior research demonstrates that factors, such as past experience, predispositions to trust and other personality factors affect individuals' own privacy concerns [54, 60]. In future work, we plan to apply the method of structural equation modeling to understand how such factors impact the valuation of interdependent privacy.

References

- A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005.
- A. Acquisti and J. Grossklags. An online survey experiment on ambiguity and privacy. *Communications & Strategies*, 88(4):19–39, 2012.
- A. Acquisti, L. John, and G. Loewenstein. What is privacy worth? *Journal of Legal Studies*, 42(2):249–274, 2013.
- J. Bargh and K. McKenna. The internet and social life. *Annual Review of Psychology*, 55:573–590, 2004.
- C. Bauer, J. Korunovska, and S. Spiekermann. On the value of information - What facebook users are willing to pay. In *Proceedings of the European Conference on Information Systems (ECIS)*, 2012.
- A. Beresford, D. Kübler, and S. Preibusch. Unwillingness to pay for privacy: A field experiment. *Economics Letters*, 117(1):25–27, 2012.
- J. Berg, J. Dickhaut, and K. McCabe. Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1):122–142, 1995.
- A. Besmer and H. Lipford. Users' (mis)conceptions of social applications. In *Proceedings of Graphics Interface (GI)*, pages 63–70, 2010.

- S. Biasiola. What friends are for: How network ties enable invasive third party applications on Facebook. In *Proceedings of Measuring Networked Privacy Workshop at CSCW*, 2013.
- G. Biczók and P. Chia. Interdependent privacy: Let me share your data. In A.-R. Sadeghi, editor, *Financial Cryptography and Data Security*, volume 7859 of *Lecture Notes in Computer Science*, pages 338–353. Springer, 2013.
- R. Böhme and J. Grossklags. Vanishing signals: Trading agent kills market information. In *Proceedings of the 6th Workshop on the Economics of Networks, Systems and Computation (NetEcon)*, 2011.
- R. Böhme and J. Grossklags. Trading agent kills market information: Evidence from online social lending. In *Proceedings of the 9th Conference on Web and Internet Economics (WINE)*, pages 68–81, 2013.
- A. Braunstein, L. Granka, and J. Staddon. Indirect content privacy surveys: Measuring privacy without asking about it. In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS)*, 2011.
- R. Chellappa and R. Sin. Personalization versus privacy: An empirical examination of the online consumer’s dilemma. *Information Technology and Management*, 6(2-3):181–202, 2005.
- P. Chia, Y. Yamamoto, and N. Asokan. Is this app safe?: A large scale study on application permissions and risk signals. In *Proceedings of the 21st International World Wide Web Conference (WWW)*, pages 311–320, 2012.
- N. Christin, S. Egelman, T. Vidas, and J. Grossklags. Its all about the benjamins: An empirical study on incentivizing users to ignore security advice. In G. Danezis, editor, *Financial Cryptography and Data Security*, volume 7035 of *Lecture Notes in Computer Science*, pages 16–30. Springer Berlin Heidelberg, 2012.
- J. Coleman. Social capital in the creation of human capital. *American Journal of Sociology*, 94:S95–S120, 1988.
- D. Cooper and J. Kagel. Other regarding preferences: A selective survey of experimental results. *Handbook of Experimental Economics*, 2, 2009.
- M. Culnan. ’how did they get my name?’: An exploratory investigation of consumer attitudes toward secondary information use. *MIS Quarterly*, 17(3):341–363, Sept. 1993.
- G. Danezis, S. Lewis, and R. Anderson. How much is location privacy worth? In *Proceedings of the Workshop on the Economic of Privacy (WEIS)*, 2005.

- P. De Meo, E. Ferrara, G. Fiumara, and A. Provetti. On facebook, most ties are weak. *Communications of the ACM*, 57(11):78–84, 2014.
- T. Dinev and P. Hart. An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1):61–80, 2006.
- J. Downs, M. Holbrook, S. Sheng, and L. Cranor. Are your participants gaming the system?: Screening mechanical turk workers. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, pages 2399–2402, 2010.
- S. Egelman. My profile is my password, verify me!: The privacy/convenience tradeoff of Facebook Connect. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, pages 2369–2378, 2013.
- E. Fehr, G. Kirchsteiger, and A. Riedl. Does fairness prevent market clearing? an experimental investigation. *The Quarterly Journal of Economics*, 108(2):437–459, 1993.
- A. Felt and D. Evans. Privacy protection for social networking APIs. In *Proceedings of the 2008 Workshop on Web 2.0 Security and Privacy (W2SP)*, 2008.
- A. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the 8th Symposium On Usable Privacy and Security (SOUPS)*, 2012.
- A. P. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS)*, pages 627–638, 2011.
- R. Forsythe, J. Horowitz, N. Savin, and M. Sefton. Fairness in simple bargaining experiments. *Games and Economic behavior*, 6(3):347–369, 1994.
- J. Goodman, C. Cryder, and A. Cheema. Data collection in a flat world: The strengths and weaknesses of mechanical turk samples. *Journal of Behavioral Decision Making*, 26(3):213–224, 2013.
- M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.
- P. Green and A. Krieger. Segmenting markets with conjoint analysis. *The Journal of Marketing*, 55(4):20–31, 1991.

- P. Green and V. Rao. Conjoint measurement for quantifying judgmental data. *Journal of Marketing Research*, 8(3):355–363, 1971.
- P. Green and V. Srinivasan. Conjoint analysis in consumer research: Issues and outlook. *Journal of Consumer Research*, 5(2):103–123, 1978.
- P. Green and V. Srinivasan. Conjoint analysis in marketing: New developments with implications for research and practice. *The Journal of Marketing*, 54(4):3–19, 1990.
- J. Grossklags and A. Acquisti. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, 2007.
- S. Gupta and C. Mela. What is a free customer worth? Armchair calculations of nonpaying customers’ value can lead to flawed strategies. *Harvard Business Review*, 86(11):102–109, 2008.
- W. Güth, R. Schmittberger, and B. Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, 3(4):367–388, 1982.
- K. Hampton and B. Wellman. Neighboring in netville: How the internet supports community and social capital in a wired suburb. *City and Community*, 2(4):277–311, 2003.
- I.-H. Hann, K.-L. Hui, S.-Y. T. Lee, and I. P. Png. Overcoming online information privacy concerns: An information-processing theory approach. *Journal of Management Information Systems*, 24(2):13–42, 2007.
- I.-H. Hann, K.-L. Hui, T. Lee, and I. Png. Online information privacy: Measuring the cost-benefit trade-off. *Proceedings of the International Conference on Information Systems (ICIS)*, 2002.
- J. Helliwell and R. Putnam. The social context of well-being. *Philosophical Transactions of the Royal Society B - Biological Sciences*, 359(1449):1435–1446, Sept. 2004.
- B. Huberman, E. Adar, and L. Fine. Valuating privacy. *IEEE Security & Privacy*, 3(5):22–25, 2005.
- K.-L. Hui, B. Tan, and C.-Y. Goh. Online information disclosure: Motivators and measurements. *ACM Transactions on Internet Technology*, 6(4):415–441, 2006.
- P. Ipeirotis. Demographics of Mechanical Turk. Technical report, Social Science Research Network, Technical Report No. 1585030, 2010.

- H. Itoh. Moral hazard and other-regarding preferences. *Japanese Economic Review*, 55(1):18–45, 2004.
- N. Jentzsch, S. Preibusch, and A. Harasser. Study on monetising privacy: An economic model for pricing personal information. *ENISA*, 2012.
- R. Johnson. Trade-off analysis of consumer values. *Journal of Marketing Research*, pages 121–127, 1974.
- R. Johnson. The cbc system for choice-based conjoint analysis. *Sawtooth Software*, 1994.
- D. Kahneman, J. Knetsch, and R. Thaler. Fairness and the assumptions of economics. *Journal of Business*, 59(4):S285–S300, 1986.
- C. Kam, J. Wilking, and E. Zechmeister. Beyond the “narrow data base”: Another convenience sample for experimental research. *Political Behavior*, 29(4):415–440, Dec. 2007.
- D. Karambelkar. Spyware: A bird’s-eye view. *Gulf News*, Feb. 2014.
- A. Kavanaugh, J. Carroll, M. Rosson, T. Zin, and D. D. Reese. Community networks: Where offline communities meet online. *Journal of Computer-Mediated Communication*, 10(4), 2005.
- J. King, A. Lampinen, and A. Smolen. Privacy: Is there an app for that? In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS)*, 2011.
- P. Klopfer and D. Rubenstein. The concept privacy and its biological basis. *Journal of Social Issues*, 33(3):52–65, 1977.
- H. Krasnova, N. Eling, O. Abramova, and P. Buxmann. Dangers of facebook loginfor mobile apps: Is there a price tag for social information? In *Proceedings of the International Conference on Information Systems (ICIS)*, 2014.
- H. Krasnova, N. Eling, O. Schneider, H. Wenninger, and T. Widjaja. Does this app ask for too much data? the role of privacy perceptions in user behavior towards facebook applications and permission dialogs. In *Proceedings of the European Conference on Information Systems (ECIS)*, 2013.
- H. Krasnova, T. Hildebrand, and O. Günther. Investigating the value of privacy in on-line social networks: Conjoint analysis. In *Proceedings of the International Conference on Information Systems (ICIS)*, 2009.

- H. Krasnova, E. Kolesnikova, and O. Günther. ‘it won’t happen to me!’: Self-disclosure in online social networks. *Proceedings of the Americas Conference on Information Systems (AMCIS)*, 2009.
- H. Krasnova and N. Veltri. Privacy calculus on social networking sites: Explorative evidence from germany and usa. In *Proceedings of the Hawaii International Conference on System Sciences (HICSS)*, 2010.
- P. Kumaraguru and L. F. Cranor. Privacy indexes: A survey of westin’s studies. *Institute for Software Research International*, 2005.
- M. Majláth et al. Evaluation of environmentally friendly product attribute—results of an empirical research. In *Proceedings of the MEB 7th International Conference on Management, Enterprise and Benchmarking*, pages 201–212, 2009.
- W. Mason and S. Suri. Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior Research Methods*, 44(1):1–23, Mar. 2012.
- H. Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1), 2004.
- H. Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
- B. Orme. Which conjoint method should i use. *Sawtooth Software Research Paper Series*, 2003.
- P. Paxton. Is social capital declining in the united states? a multiple indicator assessment. *American Journal of Sociology*, 105(1):88–127, 1999.
- D. Potoglou, S. Patil, C. Gijón, J. F. Palacios, and C. Feijóo. The value of personal information online: Results from three stated preference discrete choice experiments in the uk. In *Proceedings of the European Conference on Information Systems (ECIS)*, 2013.
- Y. Pu and J. Grossklags. An economic model and simulation results of app adoption decisions on networks with interdependent privacy consequences. In R. Poovendran and W. Saad, editors, *Decision and Game Theory for Security*, Lecture Notes in Computer Science, pages 246–265. Springer, 2014.
- J. Robertson. Google+, ‘Candy Crush’ Show Risk of Leakiest Apps. *Bloomberg Technology*, Jan. 2014.

- P. Shi, H. Xu, and Y. Chen. Using contextual integrity to examine interpersonal information boundary on social network sites. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, pages 35–38, 2013.
- A. Smith. 6 new facts about facebook. <http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/>, 2014. Accessed: 2015-09-09.
- J. Smith, T. Dinev, and H. Xu. Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4):989–1016, 2011.
- J. Smith, S. Milberg, and S. Burke. Information privacy: Measuring individuals’ concerns about organizational practices. *MIS Quarterly*, 20(2):167–196, 1996.
- Socialbakers. Candy crush saga facebook statistics. <http://www.socialbakers.com/statistics/facebook/pages/detail/244944385603396-candy-crush-saga>, 2015. Accessed: 2015-09-09.
- S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC)*, pages 38–47, 2001.
- S. Spiekermann, J. Korunovska, and C. Bauer. Psychology of ownership and asset defense: Why people value their personal information beyond privacy. In *Proceedings of the International Conference on Information Systems (ICIS)*, 2012.
- D. Stahl and E. Haruvy. Other-regarding preferences: Egalitarian warm glow, empathy, and group size. *Journal of Economic Behavior & Organization*, 61(1):20–41, Sept. 2006.
- Statista. Statistics and facts about mobile app usage. <http://www.statista.com/topics/1002/mobile-app-usage/>, 2015. Accessed: 2015-09-09.
- E. Stone and D. Stone. Privacy in organizations: Theoretical issues, research findings, and protection mechanisms. *Research in Personnel and Human Resources Management*, 8(3):349–411, 1990.
- I. Symeonidis, F. Beato, P. Tsormpatzoudi, and B. Preneel. Collateral damage of facebook apps: An enhanced privacy scoring model. 2015. <https://eprint.iacr.org/2015/456.pdf>.
- J. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2):254–268, 2011.

- N. Wang, J. Grossklags, and H. Xu. An online experiment of privacy authorization dialogues for social applications. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW)*, pages 261–272, 2013.
- N. Wang, P. Wisniewski, H. Xu, and J. Grossklags. Designing the default privacy settings for facebook applications. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 249–252, 2014.
- C. Warren. Google play hits 1 million apps. <http://mashable.com/2013/07/24/google-play-1-million/>, 2013. Accessed: 2015-09-09.
- L. Wathieu and A. Friedman. An empirical approach to understanding privacy valuation. *HBS Marketing Research Paper*, (07-075), 2007.
- H. Xu, H.-H. Teo, B. Tan, and R. Agarwal. The role of push-pull technology in privacy calculus: The case of location-based services. *Journal of Management Information Systems*, 26(3):135–174, 2009.