An Economic Model and Simulation Results of App Adoption Decisions on Networks with Interdependent Privacy Consequences

Yu Pu and Jens Grossklags

College of Information Sciences and Technology
The Pennsylvania State University, University Park, PA, USA
{yxp134,jensg}@ist.psu.edu

Abstract. The popularity of third-party apps on social network sites and mobile networks emphasizes the problem of the interdependency of privacy. It is caused by users installing apps that often collect and potentially misuse the personal information of users' friends who are typically not involved in the decision-making process. In this paper, we provide an economic model and simulation results addressing this problem space. We study the adoption of social apps in a network where privacy consequences are interdependent. Motivated by research in behavioral economics, we extend the model to account for users' other-regarding preferences; that is, users care about privacy harms they inflict on their peers.

We present results from two simulations utilizing an underlying scale-free network topology to investigate users' app adoption behaviors in both the initial adoption period and the late adoption phase. The first simulation predictably shows that in the early adoption period, app adoption rates will increase when (1) the interdependent privacy harm caused by an app is lower, (2) installation cost decreases, or (3) network size increases. Surprisingly, we find from the second simulation that app rankings frequently will not accurately reflect the level of interdependent privacy harm when simultaneously considering the adoption results of multiple apps. Given that in the late adoption phase, users make their installation decisions mainly based on app rankings, the simulation results demonstrate that even rational actors who consider their peers' well-being might adopt apps with significant interdependent privacy harms. Our findings complement the usable privacy and security studies which show that users install privacy-invasive apps because they are unable to identify and understand apps' privacy consequences; however, we show that fully-informed and rational users will likely fall for privacy-invasive apps as well.

Keywords: Economic Model, Simulation, Interdependent Privacy, Other-Regarding Preferences, Scale-Free Networks, Social Network Sites, Mobile Networks, Third-Party Apps, App Adoption.

R. Poovendran and W. Saad (Eds.): GameSec 2014, LNCS 8840, pp. 246–265, 2014.

1 Introduction

Over the last ten years, we have witnessed the rapidly increasing popularity of social network sites, with Facebook being the most successful entity. In order to expand its service and functionality, Facebook opened its platform to allow outside developers to interact with users through so-called third-party Facebook applications (or social *apps*). Those applications gained worldwide popularity ever since their emergence. Similarly, the most important mobile platforms such as Android and iOS have enabled outside developers to create app content which met significant success in the marketplace.

Despite their high adoption rates, third-party apps pose privacy risks to users when they collect and potentially use user information. Some well-acknowledged issues are apps collecting more information than needed for their stated purposes [1,2]; and users demonstrating very little understanding of or ability with the management of app permissions [3,4].

A newly addressed problem associated with app permissions is the *interdependency of privacy*, which refers to the phenomenon that in an interconnected setting, the privacy of individual users not only depends on their own behaviors, but is also affected by the decisions of others [5].[1] The interdependent privacy issue is caused by users installing apps that often collect and potentially misuse the personal information of users' friends who are typically not involved in the decision-making process.

Research has not yet adequately investigated the problem of interdependent privacy, in particular, from an economic perspective. Most closely related to our work, Biczók and Chia aim to define interdependent privacy and to provide initial evidence from the Facebook permission system for social apps. They further develop a game-theoretic model to analyze users' app adoption decisions under the scenario of interdependent privacy. However, their study is limited to cases where two users are engaged in the decision-making over the adoption of one app, and therefore does not consider the complex dynamics of today's app adoption behaviors. To address this literature gap, we follow an economic approach to study how large groups of users, who are connected in a complex social network, act in an interdependent privacy scenario.

We develop an app adoption model of a rational consumer who considers cost of app adoption, benefits of an app, and the privacy consequences associated with an app adoption decision. Individuals in our model do not only consider personal costs and benefits of their decision. Instead, we consider that consumers have different levels of concern about the consequences of their adoption decisions for their peers. To accomplish this objective, we utilize the theory of

[1] In the security context, several studies have considered the interdependency of decision-making, but those models are less applicable to the app adoption scenario [6,7]. For a survey of the results in the area of interdependent security see [8].

other-regarding preferences which is well-established in psychology and economics, and has been demonstrated in various experimental studies [9,10]. In a nutshell, the theory of other-regarding preferences allows us to model users so that they consider their peers' utility when making adoption decisions.

In our research, we take a graph-theoretical approach and simulate app adoption decisions in scale-free networks to represent an approximate version of real social networks. More specifically, we conduct two simulations to investigate individuals' app adoption behaviors in two phases. One phase is the start-up period of new apps, the other phase is the later app adoption stage. More precisely, the first simulation, which considers the iterative/sequential adoption process of social apps, is used to study users' app adoption behaviors when an app is initially introduced. The second simulation, which is about comparing early adoption results of multiple apps, allows us to establish popularity rankings of the early adoption of those apps. We use those rankings to draw conclusions about the likely adoption processes of the considered apps in later adoption phases which are then heavily influenced by rankings [11].

As expected, we find that in the initial adoption phase, app adoption rates will increase when (1) the interdependent privacy harm caused by an app is lower, (2) installation cost decreases, or (3) network size increases. In the second simulation, interestingly, we find that app rankings frequently will not accurately reflect the level of interdependent privacy harm when considering the adoption results of multiple apps. Our analysis implies that in the later adoption period, even rational actors who consider their peers' well-being might adopt apps with invasive privacy practices. This helps us to explain why some apps that cause significant interdependent privacy issues are nevertheless highly popular on actual social network sites and mobile networks.

The paper is structured as follows. In Section 2, we discuss research on privacy consequences of installing third-party applications on social networking sites and mobile platforms. In Section 3, we develop our economic model of app adoption behavior. In Section 4, we describe our simulation setup. In Sections 5 and 6, we present our simulation results. Finally, we conclude in Section 7.

2 Related Work

2.1 Third-Party Applications on Social Network Sites

Primary motivators for our study are incidents that highlight the potential negative privacy and security consequences of third-party app adoption on social network sites. Several studies have documented how third-party apps are utilized to extract and to transfer user information not only to third-party app developers but also to advertising and data firms [12,13,14]. These studies are highly valuable because in most cases it is difficult to observe data practices once users have authorized third-parties to access their profiles (and their friends' profiles).

To understand the problem space from a more user-centered perspective, several research papers focus on the disclosure and authorization procedures associated with third-party apps. User studies document the concerns users have

about app adoption, and their misunderstandings about the access of third-party developers to their profiles [3,15,16]. Similarly, the impact of interface improvements of the authorization dialogies for third-party apps on user behavior has been investigated in several user studies (see, for example, [17,18]).[2]

Table 1. Most frequently requested Facebook permissions explicitly involving information of users' friends (abbreviated table from Wang et al. [17])

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
user_work_history	73	0.78%	2,961,900
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

A selected number of studies have focused on measuring aspects of the permissions system for third-party apps on social network sites [1,21,22]. These studies identify the most requested permissions, and the average number of permissions for all apps and specific categories. In Table 1, we summarize data that relates to the sharing of other users' information from a study by Wang et al. [17]. They find that specific permissions (except for basic information and email) are only used by a subset of all apps. However, due to the popularity of the over 9000 surveyed apps, the impact of these data collection and usage practices is significant. As a result, even though less than 1% of the apps request the friends' employment history, this nevertheless means that the data is accessible to third-party developers (and potentially other parties) in over 5 Million cases.

In aggregate, these studies document many obstacles that users have to overcome to identify privacy consequences of social apps, and to implement their privacy preferences in practice during the app adoption process. We complement these studies by showing that even from a rational consumer perspective

[2] Already in the context of desktop computing, user studies have investigated how to inform users more effectively about third-party apps which collect personal information and potentially allow for privacy-invasive practices [19,20].

the severity of privacy intrusions does not always translate into a low ranking of an app in comparison to more privacy-friendly offerings.

2.2 Third-Party Applications on Mobile Networks

Security and privacy issues associated with third-party applications on mobile networks are increasingly gaining importance. Particularly troublesome, app developers have been trying to use unwitting users' devices for spam and unwanted costly premium services [23]. More broadly, measurements studies found that most apps include permission requests that enable potentially dangerous practices [24].

User studies of the utilized permission systems document comprehension and usability problems that are largely similar to the results in the Facebook context (see, for example, [4,25]). As a response, technological measures to help users to manage permissions on mobile systems have been proposed. For example, Beresford et al. introduced a system to disable information requests made by a mobile application and to disable unwanted permissions [26].

Similar to the context of apps on social network sites, mobile applications gain access in various ways to information of friends (or contacts more generally). Apps with multi-platform functionality that have access, for example, to a user's Facebook account will be able to share the same information also in the mobile context. However, apps will frequently enrich this data with additional information gathered in the mobile context. For example, the new Facebook mobile app has caused a stir due to the requirement to access a user's SMS and MMS (i.e., personal and professional communications with other users) [27].

Security firm BitDefender audited over 800000 apps in the Android Play Store and found that apps frequently require access to information that impacts friends and other contacts. For example, almost 10% of the surveyed apps can read your contact list, and a sizable minority leak a user's phone book and call history [28].

For iOS devices, security firm Zscaler discovered when it scanned the 25 most popular apps across five categories, that 92% require access to a user's address book, and 32% go through a user's calendar [29].

These examples highlight that the problem of sharing the information of friends or other contacts without their explicit consent goes well beyond the context of applications on social network sites. We aim to better understand the reasons for such sharing behaviors by developing an economic model that focuses on the adoption of apps with different interdependent privacy consequences.

3 Model Overview

The framework of our model builds on the local network effects research by Sundararajan [30]. His model studies the Bayes-Nash equilibria of a network game in which heterogeneous agents connected in a social network consider the purchase of a product with network effects. Individuals are rational and make their decisions based on a well-specified payoff function. A person who does not

purchase the product receives zero payoff, while the payoff of a purchaser is influenced by the actions of her peers, her own valuation type of the product, and the product cost.

Although, we use the basic structure of the payoff function of Sundararajan's model, the focus of our analysis is quite different. Sundararajan studied individuals' purchasing behavior in a scenario where all decisions are made simultaneously, while our goal is to discover users' behavior when they can make adoption decisions sequentially. In addition, we do not only consider positive network effects, but also integrate one specific type of negative network effects: *interdependent privacy harm*. We further consider individuals to have other-regarding preferences. That is, when making adoption decisions, individuals include in their evaluation the privacy harm they potentially inflict on their peers.

In the following, we present the model and break down its different constituent parts. For reference, we provide a complete list of symbols used in our paper.

a_i	User i's adoption choice (1 = adoption, 0 = no adoption)
c	Cost of app adoption
e	Individual's interdependent privacy harm resulting from her friend's adoption behavior
θ_i	User i's valuation of an app (also called her *type*)
k_i	User i's other-regarding preference
v_i	Number of user i's friends who have already adopted the app
N	Number of users in the network
n_i	Number of user i's friends
p_i	User i's payoff
M	Number of connections per additional node in the Barabási-Albert (BA) random graph model
M_0	Number of initial seeds in the BA random graph model
SI	Set of users that have already adopted the app
I	Set of users that choose to adopt the app in one step
F	Set of friends of users in I

In our work, we assume that individuals are rational in terms of their awareness of the privacy harm associated with app adoption. In addition, they make their adoption decisions based on the payoff of their actions. Extending Sundararajan's local network effects model, we propose that user i's payoff function is:

$$p_i = a_i[(v_i + 1)\theta_i - \frac{k_i}{v_i + 1} e \cdot n_i - c] \tag{1}$$

If the payoff from adopting the app is larger than zero, the individual will always install the app on her device; otherwise she would deny the installation offer.

There are three parts in the above payoff function: value gained from the app adoption; the perceived responsibility when inflicting privacy harm on peers (i.e., other-regarding preferences); and the cost of app adoption. We discuss each of these parts in the following subsections.

3.1 Value Gained from App Adoption

The value gained from app adoption is represented by $(v_i + 1)\theta_i$. This value can further be divided into two parts: the first part is the direct value gained from using the app; the second part refers to the positive network effects, for example, the extra enjoyment the individual will perceive when a game app can be played together with her friends rather than alone.

Since individuals have different assessments regarding an app's value, we use an individual's *type*, θ_i, to represent this heterogeneity. For example, those individuals with a higher valuation type, i.e., represented by a larger θ_i, will gain more direct value from an app compared with those who have a lower valuation type, which is represented by a smaller θ_i.

In addition, it is reasonable to consider v_i, the number of user i's friends who adopt the app, will affect the utility user i gains from installing and using the app. In particular, we assume that only the number of close friends, i.e., the neighbors in the network, will positively influence the individual's payoff. This is referred to as *positive local network effects* [30]. In practice, apps may also exhibit broader positive network effects; however, we assume that local network effects dominate the adoption decision.

3.2 Care for Privacy Harm Inflicted on Friends by Adoption Decisions

The central function of social network sites (from the user's perspective) is to find friends and interact with them. Typically, individuals will care about their close friends' well-being (however, the level of concern may differ) and try to avoid taking actions that negatively affect their friends. Experimental results provide substantial evidence of the existence of such other-regarding behaviors in group interactions [10]. Other-regarding preferences, which indicate whether and how much people tend to care about others' well-being, are described in detail in a recent review paper [9]. There are two primary types of other-regarding preferences: distributive and reciprocal. The distributive other-regarding preference is caused by people's aversion of outcome inequality [31,32]. The reciprocal aspect of the other-regarding preferences theory indicates that people tend to respond in kind to a peer's behavior [33], which means that people respond to kindness with kindness, and hostility with hostility.

Our paper focuses primarily on the reciprocal aspect of other-regarding preferences. That is, users consider the well-being of their close friends who presumably would act similarly. Under the scenario of interdependent privacy, if individual i chooses to adopt an app, she will inflict a certain amount of privacy harm, e, on her friends. More specifically, user i will incorporate partially the privacy harm she inflicts on all her n_i friends in her own payoff calculation. In our model, this other-regarding preference is represented by $e \cdot n_i$. We make the assumption that e is additive across users. In other words, if a user adopts a certain app which likely impacts her friends' privacy then her worry about this decision will increase with the number of close friends, n_i. We believe this assumption

is reasonable. For example, annoyances such as spam typically affect all close friends.[3]

Studies also found that group size likely reduces the impact of other-regarding preferences due to a diffusion of responsibility [34]. When individuals know that others have taken the same potentially harmful action, they do not experience the full burden of responsibility. In our case, the guilt of inflicting privacy harm on others will be diffused with each additional close friend who has already adopted the app. Likewise, reciprocity requires an agent to respond to previous installation decisions that also impose potential privacy harm on her.[4] We use $\frac{e \cdot n_i}{v_i+1}$ to represent the part of the remaining responsibility that user i shoulders when she calculates her payoff considering her diffused responsibility and reciprocal factors. In particular, we make the assumption that the guilt of causing privacy harm is split equally across the local peers who make the adoption decision.

In order to indicate to which degree an individual is generally concerned about privacy harm imposed on friends, we use k_i to represent agent i's other-regarding preference. A larger k indicates a higher other-regarding preference, a smaller k represents lower other-regarding preferences. Thus, $\frac{k_i}{v_i+1} e \cdot n_i$ reflects how agent i cares about her friends' privacy harm inflicted by herself.

Please note that users apply a heuristic evaluation when they calculate the privacy harm inflicted on others with the formula stated above. For example, an exact calculation would require an assessment of the overlap between her friends, and her friends of friends. Theoretically, a user should only experience partial emotional relief for the installation decisions of her friends when not all of her own friends were affected by her friends' app adoptions. However, while in practice it is relatively easy to determine how many friends have installed a particular app; it is extremely cumbersome (if not impossible for an average user) to determine this more specific figure on most social network sites and mobile networks. In addition, user i cannot easily reciprocate in the app installation context against a specific user since her adoption decision affects the whole groups of close friends.

3.3 Cost of App Adoption

All practical costs associated with an adoption decision, except the interdependent privacy harm experienced by her choice, are included in the installation cost. For example, the installation costs contain, but are not limited to, the cost of finding and installing an app, the cost of learning how to use the app, and user's personal privacy harm when she chooses to install the app.

[3] For example, if Bob installs Candy Crush, a very popular third-party Facebook app, then this installation will typically trigger invitations to both his friends Eve and Trudy.

[4] Note that interdependent privacy harm user i already is suffering from cannot be influenced by herself and is therefore not part of the payoff calculation. However, it finds consideration in her other-regarding preferences.

4 Simulation Setup

Given the model we proposed above, we conduct two simulations to investigate app adoption in both its early and late phases. Based on available empirical literature on the purchasing behavior for new products, we argue that the processes for early and late adoption differ significantly. In the early phase, a pool of potential first adopters is evaluating a newly introduced product (as described in our model) while considering social and privacy factors [35]. In the later stages of adoption, users are heavily influenced by available product rankings which are interpreted as a proxy for the quality of a product [36]. (Note that early adoption decisions can be also influenced by product rankings [36]. However, social networking sites and mobile networks typically only include apps in rankings once they have reached a certain popularity threshold.)

We proceed as follows. In the first simulation, we aim to understand the percentage of users who choose to adopt an app that collects information from users' friends from the first moment the app is introduced into an app marketplace. In addition, we will show how this percentage will be affected by network size, the level of an app's interdependent privacy harm, and installation cost. In the second simulation, we simultaneously derive early adoption results for multiple apps with different interdependent privacy harms. We then proceed to rank these apps according to their associated frequency of positive early adoption decisions. Based on these rankings, we then discuss the impact of these rankings on potential later adoption by a larger pool of users.

4.1 Scale-Free Network

Evidence from measurement studies suggests that social network sites and other human-formed networks exhibit properties of scale-free networks [37,38]. We therefore conduct our simulations within the framework of a scale-free network model. The model we use to generate the network is the Barabási-Albert (BA) model [39]. The central idea of the BA model is that in a network, the more nodes a particular node connects to, the more likely the node will attract new connections. In our model, this means that the more friends a user has, the more likely others are willing to be her friends (i.e., a notion of popularity).

When using the BA model to generate a scale-free network of N people, we first randomly connect M_0 initial nodes. Then, according to the principle that the probability of connecting to an existing node is proportional to the degree of that node, each new node is connected to M existing nodes. Following this procedure, the remaining $N - M_0$ nodes are then connected to the network one by one [39].

4.2 Simulation Process

Users make their decisions according to the payoff function stated in Equation (1). Thus, before we can simulate users' behaviors, we have to decide on the parameters that appear in Equation (1). These unknowns include the number

of users in the network, topology of the network, valuation type and the other-regarding preference type of each user, installation cost, and the level of an app's interdependent privacy harm. In order to decide on these unknowns, we make the assumption that given the overall number of users in the network, nature will determine how those people are connected, and what valuation type and other-regarding preference type each user has. In addition, installation cost, c, the level of an app's interdependent privacy harm, e, and the network size, N, are predefined by us (i.e., we will indicate the specific values in the following tables and figures).

Hence, before we simulate adoption rates of a newly introduced app that causes interdependent privacy harm, we need to set values for unknown parameters, i.e., e, c and N. We then use the BA model to generate a scale-free network of N users. Next, we attribute a valuation type θ_i and an other-regarding preference type k_i to each user. Although, we assume both types to follow the uniform distribution over the interval $(0, 1)$, we do not randomly attribute types to users. Instead, we assign types according to the assumption that friends tend to have similar preferences (which is motivated by social science research, e.g., [40]). This means, users are assigned to types in such a way that people who are friends tend to have similar θ_i and k_i.

After setting values for unknown parameters, we follow a fixed simulation methodology and average the percentage of individuals who install the app across 10000 simulation rounds. In our simulation procedure each individual has the opportunity to make a positive adoption decision more than once. In other words, even if a user declines to adopt an app at first inspection, she can reconsider her decision when more friends chose to adopt that app. The simulation is set up as follows:

1. For each individual in the network, set her v_i to be 0. This is reasonable since none of the N individuals has yet installed the app. In addition, we use SI to denote the set of people that have already installed the app. Here, SI is $\emptyset$.
2. Check adoption decisions of all N individuals according to the payoff function. Use set I to record the individuals that choose to adopt the app in this step. Add each person in I to SI.
3. For each individual in set I, find the friends of them and record these friends in set F. For each person in F, find her current v_i.
4. Check the adoption decision of each person in F. Change set I so that it records all the new individuals who adopted the app in this step. Add each element in I to the set SI.
5. Repeat step 3 and step 4 until there are no individuals left in set I.
6. Divide the number of individuals in SI by the total number of users in the network. Output this result, which denotes the percentage of users who have eventually decided to adopt the app.
7. Terminate this round.

The above simulation determines the adoption result for a particular app with a given combination of values e, c, and N. To help us understand how adoption

Table 2. Distribution of app adoption outcomes for various values of e and constant $c = 0$, $N = 100$

e	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.9
$\leq 10\%$	0	0.87	4.58	11.95	22.7	32.39	43.0	53.03	61.17	69.43
$10\% \sim 90\%$	0.01	0.01	0.03	0.18	0.59	1.75	2.61	3.85	4.92	4.99
$\geq 90\%$	99.99	99.12	95.39	87.87	76.71	65.86	54.39	43.11	33.91	25.58

results change with respect to each of these app dimensions, we systemically vary these parameters.

5 Simulation Results 1: Individuals' App Installation Behaviors in Early Adoption Stage

In this section, we provide simulation results that describe early stage adoption outcomes of apps with interdependent privacy consequences. The results help us understand future app adoption outcomes which we will discuss in Section 6.

In the following subsections, we focus our analysis on one particular parameter (i.e., e, c or N) to analyze its impact on app adoption.

How Is Adoption Impacted by Interdependent Privacy Harm? For this analysis, we consider changes of the level of interdependent privacy harm, e, and keep constant the values for c and N. We consider 4 different sets of (c, N) and plot graphs to show the distribution of app adoption rates for each of these 4 sets (Figure 1). The horizontal axis represents adoption rates, while the vertical axis indicates the percentage of 10,000 simulation rounds that fall into a particular range of adoption rates. Here, we consider three ranges of app adoption results: less than 10% adoption rate; adoption rates between 10% and 90%; and adoption rates above 90%.

Figure 1 shows that with increasing privacy harm the percentage of positive adoption decisions decreases. E.g., adoption rates between 90% and 100% occur much less frequently, while there is an increased possibility of falling into the lower range of adoption rates, i.e., 0% to 10%. Numeric figures are provided in Table 2.

How Is App Adoption Impacted by Installation Cost? In this subsection, we vary the installation cost, c, from 0 to 1, while keeping the parameters for privacy harm, e, and network size, N, constant. Similar to the previous analysis, we consider 4 fixed sets of (e, N). Figure 2 demonstrates that when installation costs increase, there is a higher probability that the app will suffer from a lower adoption rate. Numeric values for fixed (e, N) equaling (0.5, 100) are provided in Table 3.

How is adoption impacted by network size? We consider different app network sizes from 100 to 2,000 nodes, and keep constant installation cost and

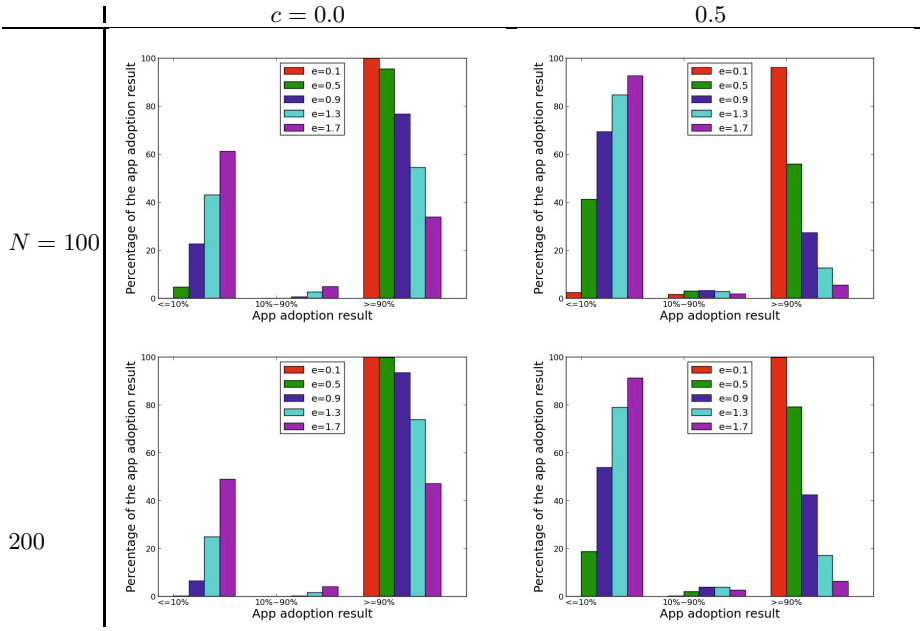

Fig. 1. Distribution of app adoption outcomes for different e with fixed c and N

Table 3. Distribution of app adoption outcomes for various values of c and constant $e = 0.5$, $N = 100$

c	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
$\leq 10\%$	4.82	8.02	12.4	18.78	28.68	41.6	55.72	72.93	87.69	96.77	99.99
$10\% \sim 90\%$	0.01	0.16	0.64	1.11	2.04	2.85	3.01	2.86	1.63	0.53	0.01
$\geq 90\%$	95.17	91.82	86.96	80.11	69.28	55.55	41.27	24.21	10.68	2.7	0

privacy harm. We consider 4 different fixed sets of (e, c), and plot the results in Figure 3. In Table 4, we provide numeric results for (e, c) equaling $(1.0, 0.5)$. From both Figure 3 and Table 4, we can observe that as the network size increases the probability of an app being adopted increases as well.

6 Simulation Results 2: App Installation Behaviors in the Late Adoption Stage

In this section, we aim to understand app installation results once rational early adopters have evaluated new apps and a ranking has become available that increases the prominence of the new apps (proportional to its ranking) to a larger user group. Rankings based on early adoption results play an important role in

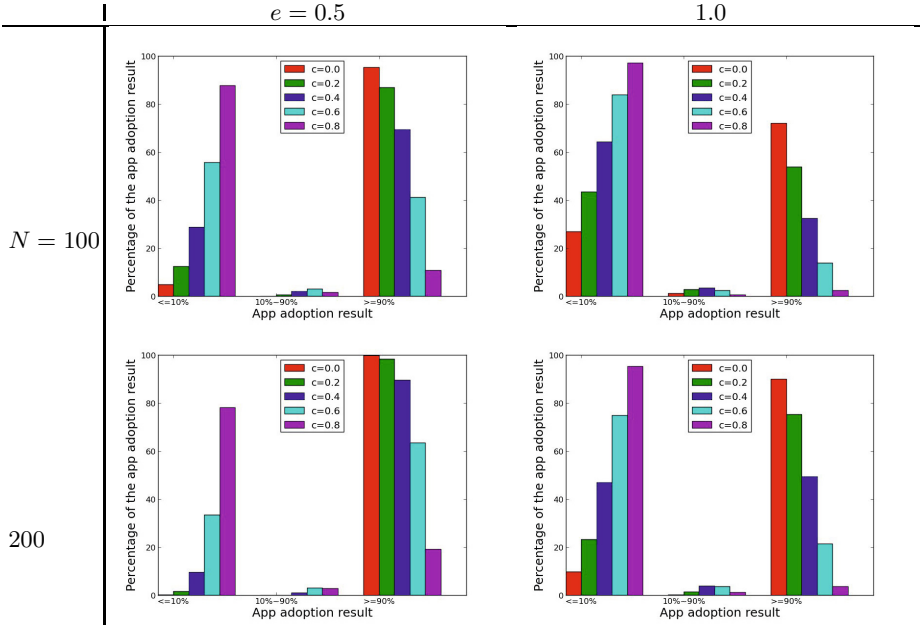

Fig. 2. Distribution of app adoption outcomes for different c with fixed e and N

Table 4. Distribution of app adoption outcomes for various values of N and constant $e = 1.0$, $c = 0.5$

N	100	200	500	1000	2000
$\leq 10\%$	74.21	61.07	37.54	19.07	5.23
$10\% \sim 90\%$	3.12	4.0	3.94	2.11	0.55
$\geq 90\%$	22.67	34.93	58.52	78.82	94.22

shaping users' adoption behavior in the later adoption phase since consumers will frequently rely on these rankings during their own adoption decisions [41,42].

For example, based on evidence from an iOS app market, Garg and Telang found that top-ranked for-pay apps generated about 150 times more downloads compared to apps ranked at about 200 [42]. Similarly, Carare showed that consumers' willingness to pay for a top-ranked app is about $4.50 greater than for the same unranked app [41]. Further direct evidence of the impact of app rankings is provided by a more recent study. Applying a data-driven approach, Ifrach and Johari studied the effect of the top-rank position on demand in the context of mobile app markets [11]. They found that the demand for an app almost doubles when its rank shifts from position 20 to position 1. Taken together, rankings can serve as an important indicator for future adoption outcomes in app markets.

In addition, previous research showed that users adopt apps with high privacy harms because they did not understand the fact that apps maliciously harvest their profile information [3,4]. Some researchers expect (e.g., [4]), it would be sufficient to protect users from undesirable apps if at least some users demonstrated

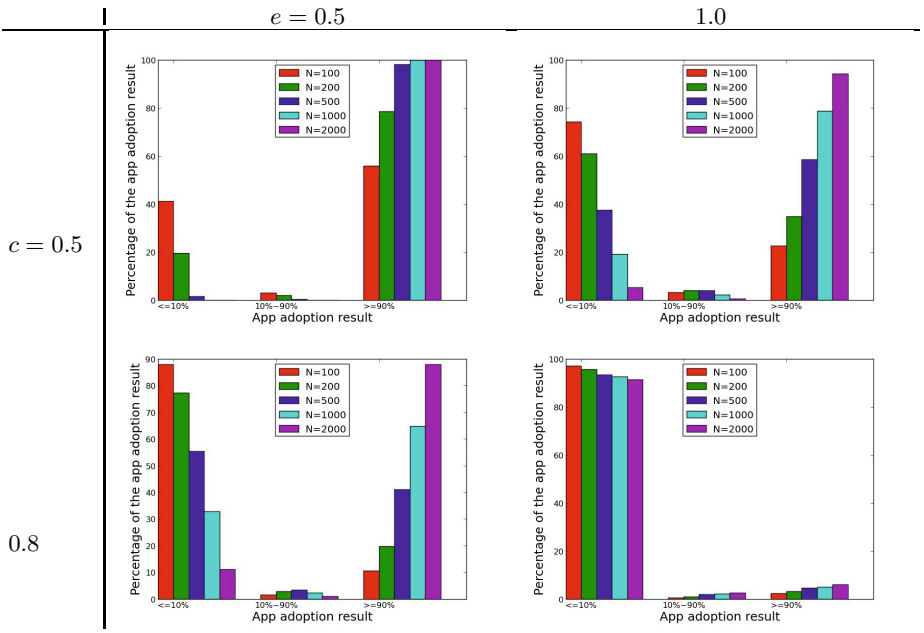

Fig. 3. Distribution of app adoption outcomes for different N with fixed e and c

awareness and understanding of permissions. However, in this section we show that even if an early adopter group rationally evaluates the different aspects of new apps, then the resulting app ranking can provide misleading signals to less savvy consumers in the later adoption phase.

We take the following approach. We first investigate early adoption results for multiple apps (by following the general methodology outlined in the previous section). More precisely, we first simulate early adoption results of 100 apps with different levels of e for 50 times. Note here, the level of privacy harm ranges from 0.1 to 10, the network size is fixed at $N = 100$, and installation cost is constant at $c = 0$. We then rank those apps based on their early adoption rates for each of the 50 simulations, i.e., we collect 50 rankings. By analyzing the variation in the early adopter rankings, we can then gain some insights about apps' likely future adoption results. Or to put it differently, we can discuss the informativeness of the app ranking as a signal to the consumer.

We show the simulation results in Figures 4 and 5. Note that in both figures, each number on the x-axis indicates a particular app with a value of e equal to that number. For example, 4 represents the app with $e = 4$. In Figure 4, the y-axis represents the number of individuals who adopt the app. The y-axis in Figure 5 represents the cardinal number of each app's rank. In both figures, the dots represent the mean value, while the ranges that include the dot in the vertical direction denote the standard deviation of the relevant value.

As we can see in Figure 4, the more interdependent privacy harm an app causes, the lower its adoption rate will become. In addition, adoption outcomes of apps with either a particularly high e or low e do not vary too much. In most

cases, those apps either have close to 0% adoption rate or a very high adoption rate, respectively. However, adoption results of apps with a medium e change a lot and rarely result in extremely high or low initial adoption rates. This indicates that by comparing adoption results, we can differentiate among apps with particularly low levels, medium levels and high levels of interdependent privacy harm.

By inspecting the mean value in Figure 5, we can observe that the lower the interdependent privacy harm, e, an app is associated with, the higher the ranking it will receive. However, from a practical perspective this basic observation can be challenged once we examine the standard deviation of rankings.[5] Most apps (except the very privacy-friendly apps) have a very high standard deviation concerning their ranks; the result of which is the phenomenon that wide ranges of apps' potential ranking outcomes are overlapping. As is indicated in the figure, this is particularly relevant for apps with $e > 2$. In other words, it is highly possible that an app with a quite high privacy harm e ends up with a favorable ranking, while an app with a comparatively low privacy harm e receives a very low ranking. For example, observing Figure 5, simulation outcomes are quite feasible in which a privacy-unfriendly app ($e = 10$) ranks in the 30th place while the app with a low privacy harm ($e = 2$) receives rank 50. That is, it may be very misleading to rely on rankings for apps that do not fall into the category of the lowest privacy harm ($e < 2$) even if the initial ranking was determined by a set of rational early adopters.

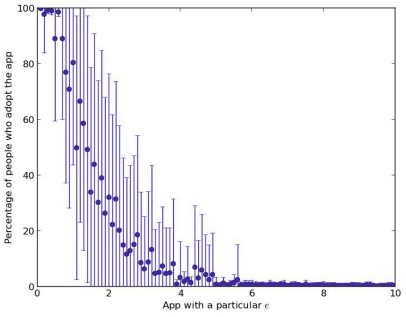

Fig. 4. Adoption results of 100 apps with e changing from 0.1 to 10

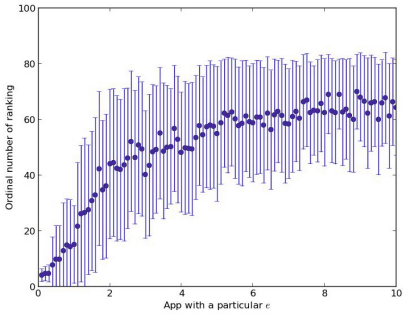

Fig. 5. Rankings of 100 apps with e changing from 0.1 to 10

To better illustrate that rankings cannot accurately reflect apps' interdependent privacy harm, we compare rankings between pairs of apps with one app having a medium level of privacy harm and one app with a relatively high level of privacy harm. Here, we compare four groups of apps: (1) app with $e = 2.6$

[5] Essentially, we investigate the variation between a low number of different ranking outcomes which in our case are 50 alternate universes of app rankings.

and app with $e = 9.5$, (2) app with $e = 3.4$ and app with $e = 9.5$, (3) app with $e = 3.8$ and app with $e = 9.8$ and (4) app with $e = 4.5$ and app with $e = 8.8$. For each group, we plot the 50 rankings of a particular pair of apps in Figure 6. As we can see in each figure, blue dots and red dots fall into the same range, and reveal no discernible pattern. We use Welch's $t-test$ to examine the relationship between the 50 rankings for each pair. The statistical results, shown in Table 5, indicate that for all of these four pairs, the potential rankings of the app with a medium level of privacy harm are not significantly different from the rankings of the app with a high level of privacy harm.

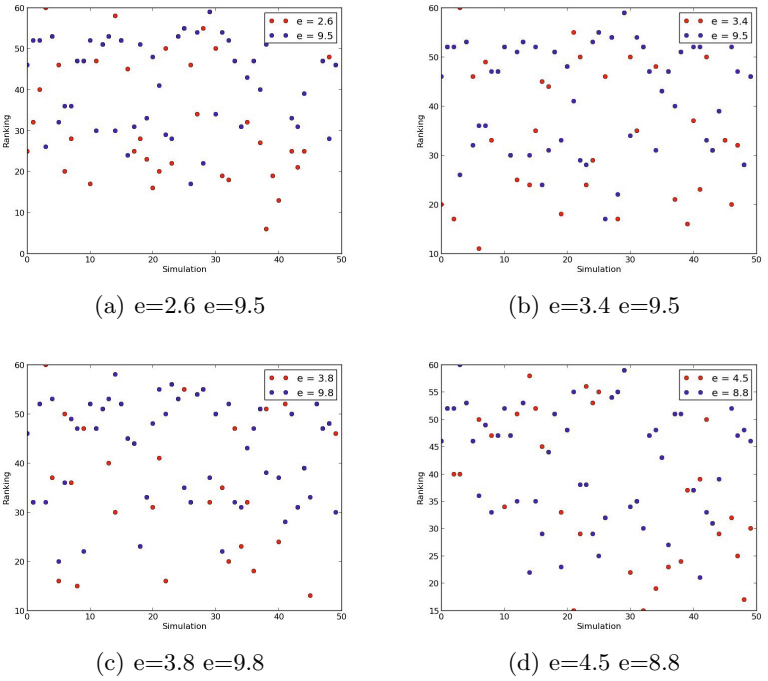

(a) e=2.6 e=9.5 (b) e=3.4 e=9.5

(c) e=3.8 e=9.8 (d) e=4.5 e=8.8

Fig. 6. Comparison of app rankings for four sample groups

Table 5. Statistical results of Welch's t-test for four sample groups

Group	F-Value	p-Value
$e = 2.6$ $e = 9.5$	-1.84	0.07
$e = 3.4$ $e = 9.5$	-1.42	0.16
$e = 3.8$ $e = 9.8$	-1.02	0.31
$e = 4.5$ $e = 8.8$	-0.63	0.53

In summary, we assume that in the early adoption period, rational users are able to identify privacy-intrusive apps; for example, partly because of the lack

of market signals (i.e., a ranking) they have a higher incentive to inspect applications. However, after a sufficiently large group of early adopters has inspected the app, the platform provider will typically include the app in its rankings. Rational as well as less savvy adopters will now likely rely on the app rankings to guide their adoption behaviors. However, since the resulting ranking is not informative enough to reflect app's interdependent privacy harm level, users are likely to also fall for apps with significant privacy harm. This observation complements the findings in the behavioral literature that users adopt apps with high privacy harms mostly due to their unawareness of apps' malicious and intrusive privacy practices. However, our take-away is somewhat disillusioning. Even if we can motivate a group of early adopters to rationally evaluate apps and we assume that they understand the privacy consequences of the installation, then the long-term outcomes might still disappoint privacy and consumer advocates.

7 Conclusion

In the interconnected setting of social network sites and mobile networks, apps' practices to collect personal information of users' friends and to allow for potential misuse amplifies the importance of interdependent privacy. That is, the privacy of an individual user does not only depend on her own behavior, but it is also the result of the decisions of her friends.

Taking an economic perspective, we propose a model of the adoption behavior for social apps in a networked system where privacy consequences are interdependent. Motivated by behavioral economics research, we model users to exhibit other-regarding preferences about the privacy well-being of their peers. We present two simulation approaches to investigate individuals' app adoption behaviors: early adoption of individual apps, and later adoption of a pool of apps with different privacy harms. The simulation results indicate that in the early adoption period, either lowering the level of interdependent privacy harm or decreasing the installation cost will increase the app adoption rates. The results also show app adoption rates will increase with a growing network size. Based on the second simulation approach, we conclude that rankings based on early adoption results frequently will not accurately reflect the level of apps' interdependent privacy harm. This is especially relevant for rankings of apps that have medium and high level of privacy harm.

While further study is needed, for example, we are investigating the robustness of our results to different specifications of the model, we believe that our study can contribute to the policy discussion on app privacy [43]. Privacy advocates should cautiously reconsider the expected impact of added scrutiny by early adopters in a marketplace; that is, encouraging individuals to pay more attention to the potential privacy harm of apps may not create the anticipated ripple-effects in the marketplace. We believe that in many cases it is likely misleading to rely on such market signals when we are considering products with strong network effects and interdependent privacy harm. In addition, our work highlights the important role of the platform provider. In particular, the design

and scope of rankings should be carefully tested to increase the likelihood that market signals are meaningful. For example, rankings could be limited to data with low variability of ratings.

To better understand the impact of different rankings we intend to work on actual app adoption data. Unfortunately, publicly available data usually does not provide details of app adoption dynamics. Instead, we favor an experimental approach (similar to [36]) to further calibrate our model.

References

1. Chia, P., Yamamoto, Y., Asokan, N.: 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), pp. 311–320 (April 2012)
2. Felt, A., Evans, D.: Privacy protection for social networking APIs. In: Proceedings of the 2008 Workshop on Web 2.0 Security and Privacy (W2SP) (May 2008)
3. Besmer, A., Lipford, H.: Users' (mis)conceptions of social applications. In: Proceedings of Graphics Interface (GI), pp. 63–70 (May 2010)
4. Felt, A., Ha, E., Egelman, S., Haney, A., Chin, E., Wagner, D.: Android permissions: User attention, comprehension, and behavior. In: Proceedings of the 7th Symposium on Usable Privacy and Security (SOUPS), pp. 3:1–3:14 (July 2012)
5. Biczók, G., Chia, P.H.: Interdependent privacy: Let me share your data. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 338–353. Springer, Heidelberg (2013)
6. Grossklags, J., Christin, N., Chuang, J.: Secure or insure?: A game-theoretic analysis of information security games. In: Proceedings of the 17th International World Wide Web Conference (WWW), pp. 209–218 (April 2008)
7. Kunreuther, H., Heal, G.: Interdependent security. Journal of Risk and Uncertainty 26(2), 231–249 (2003)
8. Laszka, A., Felegyhazi, M., Buttyán, L.: A survey of interdependent information security games. ACM Computing Surveys (forthcoming)
9. Cooper, D., Kagel, J.: Other regarding preferences: A selective survey of experimental results (forthcoming), `http://myweb.fsu.edu/djcooper/research/otherregard.pdf`
10. Stahl, D., Haruvy, E.: Other-regarding preferences: Egalitarian warm glow, empathy, and group size. Journal of Economic Behavior & Organization 61(1), 20–41 (2006)
11. Ifrach, B., Johari, R.: The impact of visibility on demand in the market for mobile apps. Technical report, SSRN Working Paper (February 2014)
12. Book, T., Wallach, D.: A case of collusion: A study of the interface between ad libraries and their apps. In: Proceedings of the 3rd Annual ACM CCS Workshop on Security and Privacy in Smartphones & Mobile Devices (SPSM), pp. 79–86 (November 2013)
13. Krishnamurthy, B., Wills, C.: On the leakage of personally identifiable information via online social networks. In: Proceedings of ACM SIGCOMM Workshop on Online Social Networks (WOSN), pp. 7–12 (August 2009)
14. Steel, E., Fowler, G.: Facebook in privacy breach. The Wall Street Journal (October 2010)
15. King, J., Lampinen, A., Smolen, A.: Privacy: Is there an app for that? In: Proceedings of the 7th Symposium on Usable Privacy and Security (SOUPS), pp. 12:1–12:20 (July 2011)

16. Tam, J., Reeder, R., Schechter, S.: I'm allowing what? Disclosing the authority applications demand of users as a condition of installation. Technical Report MSR-TR-2010-54, Microsoft Research (2010)
17. Wang, N., Grossklags, J., Xu, H.: An online experiment of privacy authorization dialogues for social applications. In: Proceedings of the Conference on Computer Supported Cooperative Work (CSCW), pp. 261–272 (February 2013)
18. Wang, N., Xu, H., Grossklags, J.: Third-party apps on Facebook: Privacy and the illusion of control. In: Proceedings of the ACM Symposium on Computer Human Interaction for Management of Information Technology (CHIMIT), pp. 4:1–4:10 (December 2011)
19. Good, N., Dhamija, R., Grossklags, J., Aronovitz, S., Thaw, D., Mulligan, D., Konstan, J.: Stopping spyware at the gate: A user study of privacy, notice and spyware. In: Proceedings of the Symposium on Usable Privacy and Security (SOUPS), pp. 43–52 (July 2005)
20. Good, N., Grossklags, J., Mulligan, D., Konstan, J.: Noticing notice: A large-scale experiment on the timing of software license agreements. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), pp. 607–616 (April-May 2007)
21. Shehab, M., Marouf, S., Hudel, C.S.: ROAuth: Recommendation based open authorization. In: Proceedings of the 7th Symposium on Usable Privacy and Security (SOUPS), pp. 11:1–11:12 (July 2011)
22. Wang, N.: Third-party applications' data practices on Facebook. In: Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems, Extended Abstracts (CHI EA), pp. 1399–1404 (May 2012)
23. Felt, A., Finifter, M., Chin, E., Hanna, S., Wagner, D.: A survey of mobile malware in the wild. In: Proceedings of the ACM Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM), pp. 3–14 (October 2011)
24. Felt, A., Greenwood, K., Wagner, D.: The effectiveness of application permissions. In: Proceedings of the 2nd USENIX Conference on Web Application Development (WebApps), p. 7 (June 2011)
25. Kelley, P., Cranor, L., Sadeh, N.: Privacy as part of the app decision-making process. In: Proceedings of the ACM Annual Conference on Human Factors in Computing Systems (CHI), pp. 3393–3402 (April 2013)
26. Beresford, A., Rice, A., Skehin, N., Sohan, R.: Mockdroid: Trading privacy for application functionality on smartphones. In: Proceedings of the 12th Workshop on Mobile Computing Systems and Applications (HotMobile), pp. 49–54 (March 2011)
27. Woollaston, V.: Is Facebook reading your TEXTS? Android update lets app access your written and picture messages. Daily Mail Online (January 2014)
28. Karambelkar, D.: Spyware: A bird's-eye view. Gulf News (February 2014)
29. Robertson, J.: Google+, 'Candy Crush' show risk of leakiest apps. Bloomberg Technology (January 2014)
30. Sundararajan, A.: Local network effects and complex network structure. The BE Journal of Theoretical Economics 7(1) (January 2007)
31. Fehr, E., Schmidt, K.: A theory of fairness, competition, and cooperation. The Quarterly Journal of Economics 114(3), 817–868 (1999)
32. Bolton, G., Ockenfels, A.: ERC: A theory of equity, reciprocity, and competition. American Economic Review 90(1), 166–193 (2000)
33. Berg, J., Dickhaut, J., McCabe, K.: Trust, reciprocity, and social history. Games and Economic Behavior 10(1), 122–142 (1995)

34. Darley, J., Latane, B.: When will people help in a crisis? In: Hochman, S. (ed.) Readings in Psychology, pp. 101–110. MSS Information Corporation (1972)
35. Fisher, R., Price, L.: An investigation into the social context of early adoption behavior. Journal of Consumer Research 19(3), 477–486 (1992)
36. Salganik, M., Dodds, P., Watts, D.: Experimental study of inequality and unpredictability in an artificial cultural market. Science 311(5762), 854–856 (2006)
37. Ahn, Y., Han, S., Kwak, H., Moon, S., Jeong, H.: Analysis of topological characteristics of huge online social networking services. In: Proceedings of the 16th International World Wide Web Conference (WWW), pp. 835–844 (May 2007)
38. Mislove, A., Marcon, M., Gummadi, K., Druschel, P., Bhattacharjee, B.: Measurement and analysis of online social networks. In: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC), pp. 29–42 (October 2007)
39. Barabási, A., Albert, R.: Emergence of scaling in random networks. Science 286(5439), 509–512 (1999)
40. Verbrugge, L.: The structure of adult friendship choices. Social Forces 56(2), 576–597 (1977)
41. Carare, O.: The impact of bestseller rank on demand: Evidence from the app market. International Economic Review 53(3), 717–742 (2012)
42. Garg, R., Telang, R.: Inferring app demand from publicly available data. MIS Quarterly 37(4), 1253–1264 (2013)
43. Good, N., Grossklags, J., Thaw, D., Perzanowski, A., Mulligan, D., Konstan, J.: User choices and regret: Understanding users' decision process about consensually acquired spyware. I/S: A Journal of Law and Policy for the Information Society 2(2), 283–344 (2006)