

Interdependent Privacy: Let Me Share Your Data

Gergely Biczók¹ and Pern Hui Chia²

¹ Dept. of Telematics

Norwegian University of Science and Technology

`gbiczok@item.ntnu.no`

² Centre for Quantifiable Quality of Service (Q2S)

Norwegian University of Science and Technology

`chia@q2s.ntnu.no`

Abstract. Users share massive amounts of personal information and opinion with each other and different service providers every day. In such an interconnected setting, the privacy of individual users is bound to be affected by the decisions of others, giving rise to the phenomenon which we term as *interdependent privacy*. In this paper we define online privacy interdependence, show its existence through a study of Facebook application permissions, and model its impact through an Interdependent Privacy Game (IPG). We show that the arising negative externalities can steer the system into equilibria which are inefficient for both users and platform vendor. We also discuss how the underlying incentive misalignment, the absence of risk signals and low user awareness contribute to unfavorable outcomes.

Keywords: Interdependent Privacy, Application Permissions, Facebook Apps, Externality, Incentive Misalignment, Game Theory.

1 Introduction

In today's networked world, users of online services have become logically interconnected in many ways. Such relationships often involve sharing personal information or opinion via named or unnamed user accounts or pseudonyms. Privacy concerns arise along with data sharing. In such an intertwined setting, the privacy of individual users is bound to be affected by the decisions of others, and could be out of their own control. This gives rise to the phenomenon which we term as *interdependent privacy*. While there is a plethora of online services where privacy interdependence matters, including the blogosphere, forums, photo and video sharing portals, the low-hanging fruits in this context are Online Social Networks (OSNs). A particularly interesting example is Facebook and its platform for third-party apps. Through its app platform, Facebook provides an efficient way to create a lock-in effect for users. However, the structure of the permission system associated with the platform raises some questions [1].

In this paper we take a first step towards understanding interdependent privacy. Our contribution is threefold. First, we define the concept of interdependent

privacy in the context of today’s networked society. Second, through a study on third party Facebook apps, we point out the permissions causing privacy interdependence, and quantify their prevalence. Third, we present the Interdependent Privacy Game (IPG), and show how positive (network effect) and negative (privacy loss) externalities can shape the behavior of social network users with regard to app usage. Our main finding is that equilibrium outcomes may be inefficient and contrary to best interest of users and/or platform vendor. We discuss how the underlying incentive misalignment, the absence of risk signals and low user awareness contribute to the unfavorable outcome, and hint on designing a possibly better application installation mechanism by mitigating negative externalities. Our intention is to introduce privacy interdependence to the research community, and outline interesting future research directions of both theoretical and experimental nature.

The paper is organized as follows. Section 2 explains the notion of interdependent privacy, while Section 3 exemplifies privacy interdependence on the Facebook application platform through a measurement study. Motivated by the Facebook case study, Section 4 presents a game theoretical model of interdependent privacy and analyzes its potential equilibria. We discuss our findings in Section 5. Section 6 shortly describes related work. Finally, we provide a summary and potential topics for future research in Section 7.

2 Interdependent Privacy

An early definition given by Clarke [2] was that privacy is the interest that individuals have in sustaining a ‘personal space’, free from interference by other people and organizations. Clarke further outlined four dimensions of privacy: bodily, behavioral, communication and data privacy.*

Online Privacy. As the digital world evolves, and changes online user behavior and expectations, there is no single widely accepted definition of online privacy today. Yet, adapting from Clarke’s categorization, we can structure online privacy risks in three dimensions:

- **Personal:** Potential loss of information about a user and his behavioral data.
- **Relational:** Revelation of how a user relate to and communicate with others.
- **Spatial:** Invasion of the virtual space of an online user (e.g. uninvited postings on the user’s blog and social media spaces).

* Bodily privacy concerns the integrity of the individual’s body including issues such as blood transfusion without consent, and compulsory submission of body fluids or tissues. Meanwhile, behavioral privacy concerns all aspects of human behaviors including sensitive information such as sexual preferences, political activities and religious practices. On the other hand, communication privacy demands for the ability to communicate with intended targets without routine monitoring by others, while data privacy concerns the protection of personal data, and the ability to exercise control over data that is to be processed by others [2].

There exists a rich literature on the protection of online personal and relational privacy. Spatial privacy is another important subject as virtual spaces including blogs and social media spaces (e.g., Facebook’s user timeline, LinkedIn’s user profile) are being claimed by and associated with the users.

Privacy Interdependence and Externality. Rather than focusing on protecting each of the 3 online privacy dimensions from malicious actors, we present in this article an important aspect yet to be adequately addressed in the community – *privacy interdependence*. Indeed, the protection of personal, relational and spatial privacy of individuals is increasingly dependent on the actions of others, rather than the individuals themselves, in the interconnected digital world.

The interdependence in online privacy is perhaps not a new phenomenon. Alice could easily embarrass Bob by taking and sharing a “funny” photo of Bob through conventional mediums such as posters, emails or blogs. Yet, the advent of online social networking services has made data sharing much easier across networks of users and thus a higher concern for privacy interdependence. How appropriate it could be for an OSN service to allow a user to share an information concerning or on behalf of another user, based on their relationship, for an improved user experience?

Sharing a user’s information without his direct consent can lead to the emergence of externalities. We know from [3] that an externality arises when an entity engages in an activity that influences the well-being of a bystander and yet neither pays nor receives any compensation for that effect. If the impact on the bystander is beneficial, it is called a positive externality. On the contrary, a negative side-effect is termed as negative externality. While sharing someone else’s information may yield benefits for her (e.g., personalized experience), it is also almost certain to cause a decrease in her utility (i.e., loss of online privacy).

A straightforward example of privacy interdependence in OSN is with photo tagging. Consider the case where Alice tags Bob in a photo and shares it in the OSN without explicit consent from Bob. Both Alice’s and Bob’s friends gain access to the photo in the default setting. Another excellent example of privacy interdependence is exemplified by the Facebook application platform. How well a user can protect his privacy from third party developers depends not only on his decisions, but also the decisions of his friends. We leverage the case of Facebook applications as the primary example in this article.

3 Case Study: Facebook Application Platform

The Facebook Help Center [4] describes why apps need to access user information before she can use them. As expected, it says that apps look to maximize user experience by collecting personal information. Common ways of using this information are: helping you find friends using the same app, personalizing content, aiding content sharing and quick bootstrapping of required data. It is also stated that apps are not allowed to use information for advertisements or transfer user information without your consent. We will show in this section that the last statement is not entirely valid.

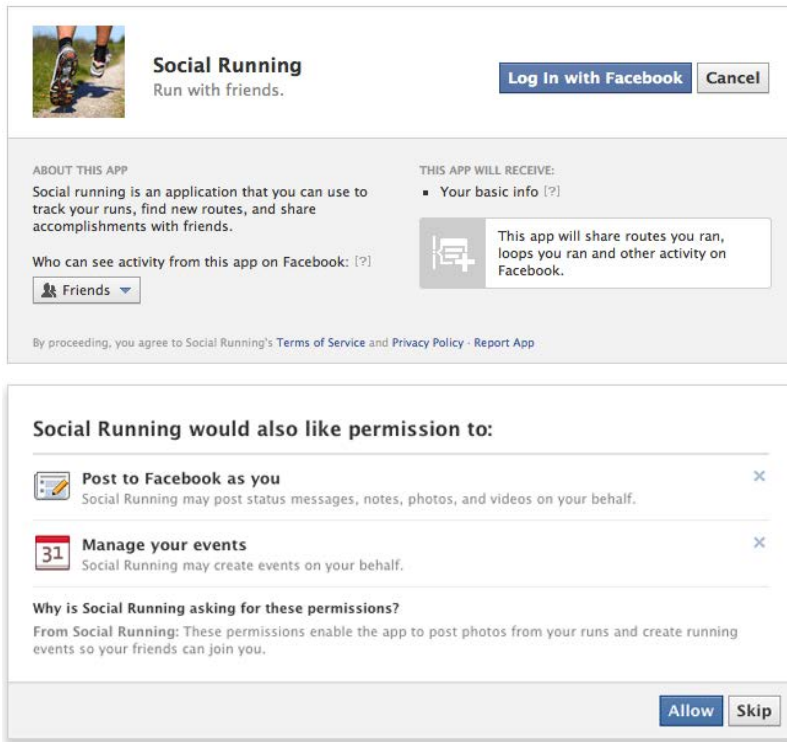


Fig. 1. The new permission request dialog of Facebook – Enhanced Auth Dialog. (Top) First screen displays the non-extended permissions requested by the app. (Bottom) Second screen shows the list of extended permissions requested, and an optional description by the developer to justify the need for extended permissions. Extended permissions can be individually rejected by the user.

Facebook relies on permission-based platform security to apply the least privilege principle to third party applications on the platform. Similarly to the Android mobile platform, third party Facebook apps wanting to access specific user information or account features are required first to ask for user consent to grant the relevant permissions. Chia et al. termed this as a user-consent permission system, as opposed to the centralized permission model of iOS where Apple decides which permissions can be requested by third party apps [5].

Application Permissions on Facebook. Facebook has a total of 65 permissions as of June 2012. They are categorized into 5 different types: basic, user or friend information, extended, open graph, and page permissions [6]. Towards the users, Facebook however distinguishes only between non-extended and extended permissions since early 2012. As shown in Fig. 1, the newly introduced Enhanced Auth Dialog of Facebook presents extended permissions to the users more prominently (on the second screen) as compared to the non-extended permissions (on the first screen).

We are most interested in investigating the interdependence aspect of different app permissions. Table 1 shows the three dimensions of online privacy risks affecting the user himself and his friends (hence the negative externality). Indeed, privacy control with app permissions on Facebook depends not only on the user’s own due diligence, but also the discipline of his friends. The table shows the number and percentage of apps exhibiting a particular privacy risk, derived from a data set of 27,029 apps constructed by Chia et al. [5]. The data set was constructed by first downloading the list of all Facebook applications on `socialbakers.com`, and then visiting each of the apps to save the list of permissions requested at install-time.

Personal privacy concerning the potential leak of user interests, birth date, education history and political views can depend on the user’s own, or any of his friends’ decisions to install a third party application. Facebook has 24 permissions as shown in Table 2 (right), which allow an app to obtain not the personal information of the user himself, but that of his friends. 1.92% of apps request for friends’ personal information. While it is a smaller sum compared to 17.15% of apps requesting for the user’s own personal information (excluding those that asks for only the `basic` permission), we believe that a majority of users are unaware of the privacy externality (or control dependency) with Facebook app permissions (see [7] for a different angle).

More than just friendship links, we consider relational privacy to include the conversations (chats, messages) and joint events between two users. Relational privacy thus affects both parties, and its protection depends on the actions of both. Installation of an app requesting a set of relational permissions, as listed in Table 3, can reveal the relation between the user and his friends. Note that the `basic` permission reveals the list of friends for a user, while `read_friendlists` reveals the custom lists of friends (e.g., close friends, band members, colleagues, relatives) that a user has made. The permission `manage_friendlists` further allows an app to edit the custom lists of friends. At the same time, `xmpp_login` exposes private chat messages, while `read_stream` allows an app to read the less private messages such as postings by friends onto user’s timeline. Excluding the `basic` permission, 1.75% of apps pose the risk of relational privacy breach.

The third dimension is spatial privacy. Again, the protection of the user’s digital space depends on his own and his friends’ decisions. Third party apps with `publish_actions` and `publish_checkins` permissions can post to the user’s own

Table 1. Online privacy dimensions, dependency of privacy control (equivalently, the affected victim), and the number of apps posing the respective risks. Figures in [brackets] exclude apps that request only the single `basic` permission.

Dimension	Dependency (Affecting)	# app	% app
Personal	Self	18204 [4634]	67.35 [17.15]
	Friends	518	1.92
Relational	Both self and friends	18204 [480]	67.35 [1.78]
Spatial	Self	494	1.83
	Friends	6249	23.12

Facebook timeline. Meanwhile, `publish_streams` has been designed to be the superset permission of `publish_actions`, allowing an app to post also onto the friends' spaces. This single permission of `publish_streams`, which has been requested by 23.12% of apps, is the main culprit for uninvited and often disgraceful postings on a user's timeline or feeds, including invitations from spammy apps or obscene postings. We regard this as a violation of spatial privacy.

Table 2. Facebook permissions with implications to personal privacy; control lies with the user himself (left), or depends on the decisions of his friends (right)

Permission	# app	% app
basic	18204	67.35
email	3766	13.93
user_about_me	284	1.05
user_activities	67	0.25
user_birthday	914	3.38
user_checkins	24	0.09
user_education_history	67	0.25
user_events	27	0.10
user_games_activity	5	0.02
user_groups	35	0.13
user_hometown	204	0.75
user_interests	94	0.35
user_likes	314	1.16
user_location	412	1.52
user_notes	12	0.04
user_online_presence	67	0.25
user_photos	574	2.12
user_questions	-	-
user_relationships	77	0.28
user_relationship_details	21	0.08
user_religion_politics	50	0.18
user_status	131	0.48
user_subscriptions	-	-
user_videos	187	0.69
user_website	12	0.04
user_work_history	107	0.40

Permission	# app	% app
friends_about_me	25	0.09
friends_activities	23	0.09
friends_birthday	162	0.60
friends_checkins	15	0.06
friends_education_history	30	0.11
friends_events	7	0.03
friends_games_activity	5	0.02
friends_groups	8	0.03
friends_hometown	44	0.16
friends_interests	33	0.12
friends_likes	51	0.19
friends_location	62	0.23
friends_notes	3	0.01
friends_online_presence	89	0.33
friends_photos	256	0.95
friends_questions	-	-
friends_relationships	19	0.07
friends_relationship_details	8	0.03
friends_religion_politics	20	0.07
friends_status	16	0.06
friends_subscriptions	-	-
friends_videos	75	0.28
friends_website	2	0.01
friends_work_history	29	0.11

4 The Interdependent Privacy Game

Motivated by our measurement findings, we propose a game-theoretic model called the Interdependent Privacy Game (IPG). While IPG is general enough to model most decision scenarios involving privacy interdependence, here we choose to focus on the inter-user effects of app installations on Facebook. We concentrate on the 2-player-1-app case, while also providing some insight for multiple players and apps. The main feature of IPG is the possibly simultaneous emergence of both positive and negative externalities, and the effect of this phenomenon on the stable outcome.

Table 3. Facebook permissions with implications to relational privacy. Control depends on both the user himself and his friends. The **basic** permission reveals the list of friends, while **read_friendlists** exposes the custom lists of friends of the user.

Permission	# app	% app
basic	18204	67.35
read_friendlists	114	0.42
read_mailbox	1	0.00
read_requests	5	0.02
read_stream	356	1.32
rsvp_event	12	0.04
xmpp_login	14	0.05
manage_friendlists	1	0.00
manage_notifications	7	0.03

4.1 Game Structure

Assumptions. We assume that the players are non-cooperative. Players have a connection in the social network, hence they are “friends”. We only consider apps that ask for permissions affecting the privacy of the friends of the user. We assume that these permissions are independent of the set of permissions requested by the app directly from the user. For tractability reasons, we focus on the scenario of two players and a single app.

Players. IPG is played by two players. Players embody users of Facebook, who also have an established friend connection.

Strategies. In the simplest form of the game, the decision is whether to install or not to install the given single app. Formally, the strategy space is $S = \{i, n\}$ for both players. Mixed strategies are probability distributions over these (pure) strategies.

Payoffs. The main characteristics of IPG is that both *positive* and *negative externalities* could emerge from the decisions of the two players. The positive externality is the so-called *network effect*: having more users install the same app could actually improve user experience [4]. On the other hand, negative externality can emerge as the other user’s decision to install an app imposes privacy risks to this user. As detailed in Section 3, Facebook’s app permission model exhibits characteristics of interdependent privacy. Many of the app permissions poses personal, relational or spatial privacy risks to the friends of an app user. Adding to positive and negative externalities, users also have their own valuations of each app. The valuation v can be positive (e.g., fun, useful) or negative (e.g., useless, waste of time, buggy); note that v represent the app’s value without network effect, and therefore is independent from other parameters. Formally:

$$\pi_{(s_1, s_2)} = f(v, e^+(s_1, s_2), e^-(s_1, s_2)), \quad (1)$$

Table 4. Facebook permissions with implications to spatial privacy. Control lies with the user himself (top), or his friends (bottom). The **publish_actions** permission allows an app to post to the user’s spaces (wall, timeline) while **publish_stream** allows an app to post to friends’ spaces.

Permission	# app	% app
publish_actions	485	1.79
publish_checkins	9	0.03

Permission	# app	% app
publish_stream	6249	23.12

Table 5. Payoff matrix for IPG

	(I)nstall	(N)ot
(I)nstall	$(v + e^+ + e^-, v + e^+ + e^-)$	(v, e^-)
(N)ot	(e^-, v)	$(0, 0)$

where s_1 (s_2) is the strategy that player 1 (player 2) plays, $v \in R$ is the user's own valuation for the app, while $e^+ > 0$ ($e^- < 0$) represents the positive (negative) externality. Note that we assume identical players with respect to valuation of both app and externalities, therefore the payoff matrix is symmetric. Payoff values are shown in Table 5.

It is important to emphasize that the positive externality emerges only when both players choose to install the app, while the negative externality appears if either one of them does. Also, the signum of the payoff has a more important message than the exact value; especially, since user valuations are hard to estimate/measure.

4.2 Analysis

In such a 2-player matrix game with parametric payoffs, the equilibrium depends on the relation of payoffs for the possible 4 outcomes. In fact, 4 inequalities (1 for each neighboring outcome-pair) are satisfactory to describe these relations. These are:

$$\pi_{ii} > \pi_{in} \quad \text{iff} \quad e^+ + e^- > 0 \quad (2)$$

$$\pi_{ii} > \pi_{ni} \quad \text{iff} \quad v + e^+ > 0 \quad (3)$$

$$\pi_{in} > \pi_{nn} \quad \text{iff} \quad v > 0 \quad (4)$$

$$\pi_{nn} > \pi_{ni} \quad \text{always.} \quad (5)$$

It can be seen that there are 8 possible cases with regard to the three parametric inequalities. Since we aim to explore the solution space, we assign a single bit to each inequality in the order of Eq. (2)-(3)-(4), and set it to 1 if the inequality holds and to 0 if it does not. After checking for conflicts, we find that Case 101 and 001 are not feasible. This leaves us with 6 potential cases.

Possible Nash Equilibria. Let us identify Nash equilibria (NE) on a case-by-case basis. The first three cases have intuitive outcomes, and a simple NE. We also characterize Pareto-optimality (PO) and social optimality (SO).

- *Case 111.* Unique NE: (i, i) . This is an intuitively beneficial outcome: direct valuation is positive, and positive network effects are stronger than the privacy loss; the result is mutual installation. The NE is both PO and SO.

- *Case 110*. Unique NE: (n, n) . Strong positive externality cannot overcome the negative direct valuation resulting in not installing the app. The NE is both PO and SO.
- *Case 000*. Unique NE: (n, n) . This is the mirror image of Case 111: negative direct valuation and strong negative externality result in not installing the app. The NE is both PO and SO.

In the following case, IPG turns into the classic *prisoner's dilemma* (PD), when the payoff is negative in equilibrium. Note that cooperating in the original PD is analogous to not installing in IPG.

- *Case 011*. Unique NE: (i, i) . When $v < |e^- + e^+|$, the payoff is negative in the NE. Under this condition, the relation of payoffs becomes $\pi_{in} > \pi_{nn} > \pi_{ii} > \pi_{ni}$, which in turn leads to a PD-type game. Hence, the NE is neither Pareto- nor socially optimal. The strategy profile (n, n) is both PO and SO. Putting it differently, users will install the app because they fear that the other player would possibly inflict a negative externality on them.

In the last two cases IPG turns into a *coordination game*.

- *Case 110*. Three possible NE: (i, i) , (n, n) and a symmetric mixed NE; in the mixed NE players play i with a positive probability $p = \frac{-v}{e^+}$, while they play n with a probability of $1 - p$. The strategy profile (i, i) is PO and SO if $|e^+| > |v| + |e^-|$; (n, n) is PO and SO otherwise.
- *Case 010*. Three possible NE: (i, i) , (n, n) and a symmetric mixed NE; in the mixed NE players play i with a positive probability $p = \frac{-v}{e^+}$, while they play n with a probability of $1 - p$. The strategy profile (n, n) is both PO and SO. All three other cases result in negative payoffs. In fact, (i, i) could yield the worst possible aggregate social payoff (sum of the 2 players), if $|e^-| - e^+ > |e^+| - |v|$. This could be interpreted as the following: users can punish each other with installing the app, even if their own valuation v and their total payoff is negative.

Evolutionary Stability. In order to view the IPG from a different angle, here we apply the evolutionary perspective. We think this is particularly fitting in the case of social networks, since organisms (persons) of the same population (users) do interact (establish connections, comment, chat) with each other, and user behavior could be considered inherent regarding certain actions (“genetic” strategies). While we can consider the same strategies as above, the solution concept here is Evolutionarily Stable Strategy (ESS). ESS is analogous to NE in an evolutionary setting: a genetically determined strategy that tends to persist once it is prevalent in a population [8]. On the other hand, ESS is a refinement of NE, and relies on a stricter definition of equilibrium. Due to this, an ESS is always a NE, but a NE is not always an ESS.

Consider the general, symmetric game shown in Table 6. A well-known result for ESS is that in a two-player, two-strategy, symmetric game, S is an ESS iff (i) $a > c$ or (ii) $a = c$ and $b > d$ [8]. Now, putting it into the context of IPG,

Table 6. General symmetric game

	S	T
S	(a, a)	(b, c)
T	(c, b)	(d, d)

let $S = I$ (install) and $T = N$ (not install). Further along, this makes $a = \pi_{ii}$, $b = \pi_{in}$, $c = \pi_{ni}$ and $d = \pi_{nn}$.

The relations between payoffs of the two strategies are different for the 6 cases described above. We now determine in which cases the ESS conditions are satisfied with regard to strategy I (install); these are 111, 110, 011 and 010, respectively. Interestingly, with the exception of 111, these are exactly the controversial or uncertain cases going back to the prisoner's dilemma and coordination games. Put plainly, both players installing the app is an equilibrium in a stricter sense (ESS is stricter than NE) in the exact scenarios, where this is contrary to the individual and aggregate interest of players (inefficient equilibrium).

Multiple Players and Apps. When considering more than two players, an interesting phenomena arises. First, the positive externality is getting stronger with the increasing number of users installing the same app (network effect). On the other hand, the maximum impact of the negative externality is already present with a single other user deciding for installing the app: the privacy loss is already there. Also factoring in multiple different applications j , the payoff function of a user i will be composed of the aggregate valuation, network effects and negative externalities:

$$\pi_i = \sum_j s_{ij} \left(v_{ij} + \left(\sum_k s_{kj:k \neq i} \right) e_j^+ + I_{\{\sum_k s_{kj:k \neq i} > 0\}} e_j^- \right), \quad (6)$$

where $s_{ij} \in \{0, 1\}$ is the decision of user i whether to install app j , v_{ij} is the valuation of app j by user i , $I \in \{0, 1\}$ is a variable indicating whether at least one friend of user i installed app j , while $e_j^+ > 0$ ($e_j^- < 0$) denotes the unit positive (negative) externality inflicted by app j . Note that apps differ in nature and permissions requested, so their inflicted externalities can be very different (e.g., game app vs. news feed app). Solving such an extended game is not straightforward. We plan to utilize the measurement results to give a numerical solution in future work.

5 Discussion

Table 7 summarizes the analysis of the different scenarios of our IPG model. We discuss several insights and implications in the following.

Sub-optimal Equilibrium. Notice that the Nash equilibrium is not always socially and/or Pareto optimal. For example, in case 011, when negative externality e^- outweighs the sum of network effect e^+ and positive user evaluation

Table 7. Nash Equilibrium (NE) as well as the Social Optimality (SO), Pareto Optimality (PO) and Vendor Optimality (VO) of different app scenarios. v denotes if initial user valuation on app is positive (+) or negative (-), while $e^+ + v$ indicates if the network effect (e^+) offsets a negative initial valuation

Case	v	$e^+ + v$	NE	SO	PO	VO
111	+	+	(i, i)	Y	Y	Y
100	-	-	(n, n)	Y	Y	N
000	-	-	(n, n)	Y	Y	N
011	+	+	(i, i)	Y/N	Y/N	Y
110 010	-	+	$(n, n), (i, i),$ mixed	Y/N	Y/N	Y/N

v , the game becomes the classic Prisoner's Dilemma. In this situation, while it is socially optimal not to install the app, both users do otherwise. Similarly, inefficiency can arise in the coordination game scenarios (i.e., case 110 and 010). The question is who will be incentivized to remedy the situation. Inefficiency can cause users to suffer from installing potentially risky or useless apps, as well as to miss out on potentially good or useful apps.

Incentive Misalignment. It is not counter-intuitive to assume that a platform vendor such as Facebook has in its best interest to stimulate app installation and data sharing. We thus define vendor optimality (VO) based on the users' decision whether to install an app in equilibrium. Comparing the SO and VO columns in Table 7, one can quickly notice the mismatched interests between platform vendor and user. This has serious implications. For example, in the Prisoner's Dilemma version of case 011, the vendor may not have direct incentives to warn against the potentially privacy-invasive apps. A similar situation can be observed on mobile application platforms. As platform owners compete in boosting the number of apps to increase platform attractiveness, problems with inappropriate apps, coupled with a lax app review process, have not been adequately addressed [9].

Absence of Risk Signaling. Attributable to the incentive mismatch, we see that the prominent cue for users to avoid bad apps today has remained with community app ratings. Unfortunately, most of the rating systems have neither factored in the risk aspects of an app nor the negative externality e^- . In their current form, ratings are thus not helpful for privacy control, particularly in the Prisoner's Dilemma situation of case 011. Specifically to Facebook, the platform has stopped displaying the average community rating of an app in the permission request dialog with the launch of the Enhanced Auth Dialog (see Fig. 1). Instead, Facebook displays a list of friends, if any, who have installed the app. This is an interesting move. Knowing that friends have installed a particular app can indeed help users better estimate positive network effects and negative privacy externality; this helps in the coordination cases of 110 and 010. However, the platform does not inform users when a friend uninstalls an app. Omitting the app ratings completely may also be unhelpful. While current user ratings do not

How people bring your info to apps they use

People on Facebook who can see your info can bring it with them when they use apps. This makes their experience better and more social. Use the settings below to control the categories of information that people can bring with them when they use apps, games and websites.

<input checked="" type="checkbox"/> Bio	<input checked="" type="checkbox"/> My videos
<input checked="" type="checkbox"/> Birthday	<input checked="" type="checkbox"/> My links
<input checked="" type="checkbox"/> Family and relationships	<input checked="" type="checkbox"/> My notes
<input type="checkbox"/> Interested in	<input checked="" type="checkbox"/> Hometown
<input type="checkbox"/> Religious and political views	<input checked="" type="checkbox"/> Current city
<input checked="" type="checkbox"/> My website	<input checked="" type="checkbox"/> Education and work
<input checked="" type="checkbox"/> If I'm online	<input checked="" type="checkbox"/> Activities, interests, things I like
<input checked="" type="checkbox"/> My status updates	<input checked="" type="checkbox"/> My app activity
<input checked="" type="checkbox"/> My photos	

If you don't want apps and websites to access other categories of information (like your friend list, gender or info you've made public), you can turn off all Platform apps. But remember, you will not be able to use any games or apps yourself.

Save Changes **Cancel**

Fig. 2. Default interdependent privacy settings on Facebook

warn against privacy risks, they can be at least useful to filter off apps with low user valuation (i.e., a low v in our model).

User Awareness and Default Settings. Rather than blaming the users for not paying enough attention, we argue that a greater effort should be made to raise the user awareness on privacy implications. Privacy interdependence with regard to apps on Facebook is certainly an area urgently requiring more attention. We expect very few users to actually realize the interdependent privacy control with Facebook apps. Permissions requesting friends' personal information are not extended permissions; they are also not prominently shown on the (first screen of the) Enhanced Auth Dialog. Furthermore, the default settings in Facebook do not protect users against the negative privacy externality. As shown in Fig. 2, other than user interests, religious views and political activities, all other personal information may be “brought to others as friends use their apps”. There is certainly more that Facebook can do to protect the users. By default, Facebook users also cannot review photo tagging by friends. However, when configured, users are prompted to approve or disprove individual photo tags. Such fine grain control is missing for apps; one cannot specify which apps that friends are using could gain access to their personal data currently.

6 Related Work

Since privacy is both an inherent human need and a complex technological challenge in OSNs, there has been a flurry of research in this area. Almost all authors agree on the fact that the privacy settings of OSNs are both complicated [10]

and non-uniform [11]. This hinders the users' ability to protect their online privacy [12], and gives rise to strange privacy patterns [13].

Chia et al. [5] studied the effectiveness of user-consent permission systems across three different platforms – Facebook, Chrome and Android. They found an absence of effective risk signals in the current app markets in addition to evidence of attempts to entice or trick users into compromising their privacy through free and mature apps [5]. Related to Facebook apps, King et al. [7] conducted a survey on the privacy knowledge, behavior and concerns of Facebook app users. They found that while almost all survey participants had heard of Facebook apps, only 77% of them were aware that apps are both created by Facebook and third parties. In addition, half of them were uncertain if Facebook reviews the apps [7]. A number of other works (e.g., [14–17]) have commented on the weaknesses of app permission systems in safeguarding user privacy and presented ways for improvement. In particular, Wang et al. [17] presented some insights regarding app permission dialogues, and gave an example where installing a given calendar app would violate the user's global privacy setting.

Observations that online privacy may be out of the control of the user himself have been made earlier. Researchers [18, 19] demonstrated the ability to infer private user information using only friendship links, group memberships and information shared by others publicly. Albeit similar, there is a distinctive difference between our work and theirs. Rather than focusing on unintended disclosure of private information inferred by combining pieces of public information, our work has looked into the case of explicit sharing of friends' data (through Facebook permissions).

Dealing with explicit collaborative information sharing, Hu et al. [20] proposed a method to detect and resolve privacy conflicts. Here, we focus on the interdependent nature of app privacy: we study how the Facebook permission system affects not only the user installing an app but also his friends.

Finally, [21] is closest to our work in terms of modeling interdependence. It shows the negative externality inflicted by websites using a weak password criterion to websites with strong authentication mechanisms. In order to incorporate both positive and negative externalities, our formulation of user welfare follows [22].

7 Summary and Future Work

In this paper we have taken a first step towards defining and understanding interdependent privacy. We have demonstrated the existence of privacy interdependence through a study on the Facebook application platform and its permission system. By constructing a simple Interdependent Privacy Game, we have analyzed the externalities caused by privacy interdependence and their effect on the users' and vendor's welfare in equilibrium. We have also discussed why these inefficient equilibria can emerge, and hinted on how to design a better application installation mechanism. We hope that our paper could also inspire further theoretical and experimental research on interdependent privacy.

Future Work. We have identified three potential directions of research.

Future modeling directions. Several game-theoretical extensions of IPG can be explored in the future. These include: taking into consideration the amount and sensitivity of personal data stored in the given OSN user account; incorporating unfriending by using a coordination game with a secure outside option [23]; repeated games leading to iterated PD; an evolutionary game of privacy-conscious and thrill-seeking users; and a game model based on the decisions of friends who cooperate. A very interesting improvement would be to play the multiple person, multiple app game on real social graphs.

Mechanism design based on economic theory and usability guidelines. The standard economic literature offers two ways of dealing with negative externalities: taxing of externality-producing activities and compensation of the victim by the entity inflicting the externality. While the theory is clear, it rarely makes its way into practice due to cost-minimizing behavior, hard-to-identify externality sources or the ineffectiveness of monetary compensation [24]. A possibly more effective way involves a slow cultural change which can be implemented by educating the actors of the ecosystem about externalities. Putting it into the context of OSNs, this means incentivizing the users to learn about privacy (interdependence). While this may be useful in the long run, current Facebook privacy settings may hinder its success [10].

Another alternative solution is redesigning the system mitigating negative externalities. Regarding OSNs and Facebook in particular, this means providing the user with an intuitive, yet economically inspired app install mechanism. Such a mechanism may include giving explicit control to the user (similarly to the already existing photo tagging) and/or introducing extended control over apps with negative externality potential. Explicit control for the user could be aided by an intuitive and informative interface, such as the one presented in [25]. Stronger control over apps with negative externality could be implemented, e.g., by requiring that such apps should be able to operate both without and with “friend permissions”, defining two levels of operation. In addition, designing a method for acquiring meaningful user valuations for apps could help the platform vendor steer the system towards a favorable steady state. Note that all these ideas are very challenging to implement, since they have to satisfy strict usability requirements and be to the OSN providers’ liking.

Other online platforms with privacy interdependence. Other online systems exhibiting privacy interdependence are abundant, e.g., various mobile application platforms (Android, iOS, Windows Phone), blogs, forums, webshops and even public cloud services. Collecting measurement data on such systems and modeling their interdependent privacy aspects are important future work.

References

1. The 3 Facebook permissions you should never agree too,
<http://facecrooks.com/Internet-Safety-Privacy/the-3-facebook-app-permissions-you-should-never-agree-to.html> (last accessed: October 2012)

2. Clarke, R.: Introduction to dataveillance and information privacy, and definitions of terms (1997) (revised in 1999, 2005, 2006), <http://www.rogerclarke.com/DV/Intro.html> (last accessed: June 2012)
3. Mankiw, N.: Principles of Economics. Available Titles CourseMate Series, vol. 1. South-Western Cengage Learning (2008)
4. Facebook Help Center – App Basics, <https://www.facebook.com/help/178140838985151/> (last accessed: October 2012)
5. Chia, P.H., Yamamoto, Y., Asokan, N.: Is this app safe? A large scale study on application permissions and risk signals. In: Proceedings of the 21st International Conference on World Wide Web, WWW 2012. ACM, New York (2012)
6. Facebook Permissions Reference, <https://developers.facebook.com/docs/authentication/permissions> (last accessed: June 2012)
7. King, J., Lampinen, A., Smolen, A.: Privacy: Is there an app for that? In: Proc. of the 7th Symposium on Usable Privacy and Security, SOUPS 2011, pp. 12:1–12:20. ACM (2011)
8. David, E., Jon, K.: Networks, Crowds, and Markets: Reasoning About a Highly Connected World. Cambridge University Press, New York (2010)
9. Chia, P.H., Heiner, A.P., Asokan, N.: Use of ratings from personalized communities for trustworthy application installation. In: Aura, T., Järvinen, K., Nyberg, K. (eds.) NordSec 2010. LNCS, vol. 7127, pp. 71–88. Springer, Heidelberg (2012)
10. Johnson, M., Egelman, S., Bellovin, S.M.: Facebook and privacy: it's complicated. In: Proceedings of the Eighth Symposium on Usable Privacy and Security, SOUPS 2012, pp. 9:1–9:15. ACM, New York (2012)
11. Bonneau, J., Preibusch, S.: The privacy jungle: On the market for data protection in social networks. In: The Eighth Workshop on the Economics of Information Security, WEIS 2009 (2009)
12. Liu, Y., Gummadi, K.P., Krishnamurthy, B., Mislove, A.: Analyzing facebook privacy settings: user expectations vs. reality. In: Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, IMC 2011, pp. 61–70. ACM, New York (2011)
13. Dey, R., Jelveh, Z., Ross, K.: Facebook users have become much more private: A large-scale study. In: 2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops), pp. 346–352 (March 2012)
14. Felt, A.P., Greenwood, K., Wagner, D.: The effectiveness of application permissions. In: Proc. of the 2nd USENIX Conf. on Web Application Development, WebApps 2011. USENIX Association (2011)
15. Barrera, D., van Oorschot, P.C., Somayaji, A.: A Methodology for Empirical Analysis of Permission-Based Security Models and its Application to Android Categories and Subject Descriptors. In: Proc. of the 17th ACM Conf. on Computer and Communications Security, CCS 2010, pp. 73–84. ACM (2010)
16. Tam, J., Reeder, R.W., Schechter, S.: I'm Allowing What? Disclosing the authority applications demand of users as a condition of installation. Technical report, Microsoft Research, MSR-TR-2010-54 (2010)
17. Wang, N., Xu, H., Grossklags, J.: Third-party apps on facebook: privacy and the illusion of control. In: Proceedings of the 5th ACM Symposium on Computer Human Interaction for Management of Information Technology, CHIMIT 2011, pp. 4:1–4:10. ACM, New York (2011)

18. Zheleva, E., Getoor, L.: To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. In: Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, pp. 531–540. ACM (2009)
19. Jernigan, C., Mistree, B.F.: Gaydar: Facebook friendships expose sexual orientation. First Monday 14(10) (October 2009)
20. Hu, H., Ahn, G.J., Jorgensen, J.: Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In: Proceedings of the 27th Annual Computer Security Applications Conference, ACSAC 2011, pp. 103–112. ACM, New York (2011)
21. Preibusch, S., Bonneau, J.: The password game: negative externalities from weak password practices. In: Alpcan, T., Buttyán, L., Baras, J.S. (eds.) GameSec 2010. LNCS, vol. 6442, pp. 192–207. Springer, Heidelberg (2010)
22. Johari, R., Kumar, S.: Congestible services and network effects. In: Proceedings of the 11th ACM Conference on Electronic Commerce, EC 2010, pp. 93–94. ACM, New York (2010)
23. Goeree, J.K., Holt, C.A.: Ten little treasures of game theory and ten intuitive contradictions. Virginia Economics Online Papers 333, University of Virginia, Department of Economics (February 2000)
24. Center for the Advancement of Steady State Economy – Negative Externalities Are the Norm, <http://steadystate.org/negative-externalities/> (last accessed: October 2012)
25. Besmer, A., Lipford, H.R., Shehab, M., Cheek, G.: Social applications: exploring a more secure framework. In: Proceedings of the 5th Symposium on Usable Privacy and Security, SOUPS 2009, pp. 2:1–2:10. ACM, New York (2009)