Mining Privacy Settings to Find Optimal Privacy-Utility Tradeoffs for Social Network Services

Shumin Guo
Wright State University
3640 Colonel Glenn Hwy
Dayton, OH 45435
Email: guo.18@wright.edu

Keke Chen
Wright State University
3640 Colonel Glenn Hwy
Dayton, OH 45435
Email: keke.chen@wright.edu

***Abstract*—Privacy has been a big concern for users of social network services (SNS). On recent criticism about privacy protection, most SNS now provide fine privacy controls, allowing users to set visibility levels for almost every profile item. However, this also creates a number of difficulties for users. First, SNS providers often set most items by default to the highest visibility to improve the utility of social network, which may conflict with users' intention. It is often formidable for a user to fine-tune tens of privacy settings towards the user desired settings. Second, tuning privacy settings involves an intricate tradeoff between privacy and utility. When you turn off the visibility of one item to protect your privacy, the social utility of that item is turned off as well. It is challenging for users to make a tradeoff between privacy and utility for each privacy setting. We propose a framework for users to conveniently tune the privacy settings towards the user desired privacy level and social utilities. It mines the privacy settings of a large number of users in a SNS, e.g., Facebook, to generate latent trait models for the level of privacy concern and the level of utility preference. A tradeoff algorithm is developed for helping users find the optimal privacy settings for a specified level of privacy concern and a personalized utility preference. We crawl a large number of Facebook accounts and derive the privacy settings with a novel method. These privacy setting data are used to validate and showcase the proposed approach.**

I. Introduction

People use social network services (SNS) because SNS are effective means for getting connected with interested parties. There are a number of well-known SNS functionalities. For example, people can use SNS to organize events, make new friends, and keep in touch with old friends. Some SNS, such as LinkedIn, are also effective platforms for employers to find talents and job seekers to find jobs.

However, most users also have concerns about their privacy. The disclosing of private information can raise certain risks such as information spam [1], [2], insurance discrimination [3], and financial fraud [4]. More and more privacy incidents [5], [6] have educated people the risk of disclosing private information. As a result, social network services have provided sufficient privacy setting options for users to tune their privacy preference. For example, Facebook provides four visibility options (the account owner, friends only, friends of friends, and everyone) for 27 profile items.

However, the fine-grained privacy tuning also brings great challenges to users. First, many users may have no sufficient knowledge and patience to tune them. SNS providers often set the visibility to the highest level for most items, to maximize the utility of the social network as a whole. Users need to check item by item manually to tune the settings, which is very inconvenient. Second, when users tune privacy settings, the utility of SNS can be implicitly affected, although the users may not be aware of it. For example, hiding contact information such as email and phone number, the user may miss potential employers who want to contact the user. Therefore, privacy settings also reflect an intrinsic tradeoff between privacy and utility.

Users, who are confused with many privacy settings or unaware of the tradeoff between privacy and utility, will need an assistant tool to help them understand the settings and recommend candidate settings aligned with their utility preference. In particular, new users of SNS may not be aware of privacy risks [7], but they have much clearer intention on using the SNS, i.e., getting connected with friends, and reaching out to potential employers, etc. From the utility-centered view, it seems quite simple to set the visibility of the profile items. For example, a user looking for jobs via the social network can simply set all the related profile items visible, such as address, phone, email, and education, for potential employers evaluating and contacting her/him. However, when the user knows the potential privacy risk for each item, she/he may not be so decisive - how should an "acceptable" privacy level be specified? does item X have higher privacy risk than item Y? And which items can be selectively exposed so that the overall privacy risk is still acceptable?

Scope and Contributions. In this paper, we aim to develop a framework to help a new user to understand the potential privacy risks based on the fellow users' view, and achieve balanced and personalized privacy settings for a social network service. This approach tries to answer the following questions: can existing users' privacy settings be modeled to provide

guidance for a new user to choose desired optimal settings? Can we also derive implicit utility preference from privacy settings? How to find the optimal tradeoff between privacy and a user-specified utility?

We mine real privacy settings crawled from a large number of user accounts in Facebook to answer the above questions. Although individual user's settings might be noisy, we observed surprisingly stable patterns with a large number of users. Based on the crawled data, we employ the latent trait theory [8] to model the relationship between the privacy settings and the existing users' latent levels of privacy concern. With this model, the new users can clearly understand the relationship between each profile item and the privacy risk perceived by other users, and use this knowledge to appropriately specify his/her own privacy concern. The same set of data, with an appropriate transformation, can also be used to model the users' implicit utility preference. The utility modeling algorithm is further revised to incorporate the new user's utility preference and results in a *utility preference model* biased towards this user's preference.

Based on the privacy and utility models, we design an algorithm to find the best settings on a new user's privacy and utility requirements. Once the privacy and utility models are determined, each user's privacy settings in the training data can be mapped to a tradeoff between privacy and the personalized utility. As the new user's utility preference is already reflected in the utility model, we can easily find the tradeoff cases that meet the user's privacy concern and maximize the preferred utility. Therefore, the framework can satisfactorily address the aforementioned challenges of personalized privacy settings for a social network.

Our approach has a number of unique contributions.

- We introduce the utility aspect into privacy settings, and allow users to find optimal settings based on the tradeoff between privacy and utility.
- We develop a method to derive implicit utility preference from privacy settings, and allow the utility modeling biased towards a user-specified utility.
- The experimental studies are based on a large amount of real privacy settings (from over ten thousands of user accounts) obtained from the Facebook social network. These data will be made available to the public with appropriate data sanitization.

The rest of the paper is organized as follows. In Section II, we review the latent trait theory that we will use to model the privacy concerns and utility preference. In Section III, we describe the modeling and analysis methods. In Section IV, we focus on the experimental studies on real Facebook data. The related work is presented in Section V.

II. Preliminary

Latent trait theory (or item response theory (IRT)) [8] is widely accepted and used in psychometrics to design tests to measure human subjects' abilities (or attitudes). The basic idea is to build a mathematical model based on the subjects' answers to a questionnaire. The difficulty level of question is determined by the number of people who answer the question correctly. Intuitively, the persons who have higher ability should get more correct answers. The latent trait model tries to find the probabilistic relationship between the person's ability and his/her answers to the questions. This model can be used to understand the properties of the questions, and thus determine whether these questions can be used to properly evaluate subjects' certain trait. Latent trait theory can also be applied to understand subjects' strength of an attitude. For example, in our study, we try to model the SNS users' "level of privacy concern" or implicit utility preference as the latent trait, using the users' actual privacy settings (equivalent to the questionnaires).

The latent trait model assumes there exists a function modeling the relationship between the person's ability and the probability of giving the correct answer to a question, which is often modeled as a logistic function [9] as shown in Figure 1. Let the person's trait be θ. The popular two-parameter model [8] is described as follows.

$$Pr(\theta; \alpha, \beta) = \frac{1}{1 + \exp\{-\alpha(\theta - \beta)\}} \quad (1)$$

where α and β are model parameters to be learned from the data. Figure 1 illustrates some sample distributions. X-axis represents the level of ability and Y-axis is the probability of giving the correct answer to the item. When $\theta = \beta$, $Pr(\theta) = 0.5$, thus, β indicates the level of difficulty of the question for the ability test. β can have different meanings when the model is used for different purposes (we will explain it for our purpose of modeling). The slope of the curve $Pr(\theta)$ at β is $\alpha/4$, which indicates the discrimination level of the question, and can be understood as follows. If the slope is small, the probabilities of giving correct answers, between the set of persons of ability $\theta > \beta$ and the set of ability $\theta < \beta$, have small difference. Thus, the corresponding item is not good for discriminating the persons' abilities.

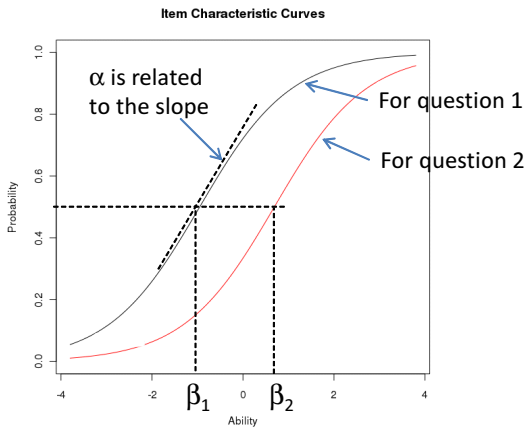

Fig. 1. Sample logistic functions for modeling the probability of giving correct answers.

It is impossible to use one question to model the latent trait. In practice, we use multiple questions (e.g., k questions) for

a number of persons (e.g., n persons), which result in a $n \times k$ grading matrix $\{s_{ij}\}$. The grading results are 0/1, indicating the correctness of the answer. This matrix is used to learn the item models: $Pr(\theta; \alpha_j, \beta_j)$, $j = 1 \ldots k$. Maximum likelihood estimation [8] is the most popular method for estimating the parameters α_j, β_j, and each person's ability θ_i from the data.

The learning process also outputs an ability level θ_i for each person in the test, which is used to assess the quality of model. In the latent trait theory, the sum $\sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j)$ is used to represent the *adjusted average score* at the ability level θ_i. Correspondingly, the answers of a specific person of ability θ_i are evaluated with the *weighted original score*: $\sum_{j=1}^{k} \alpha_j s_{ij}$, which should be close to $\sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j)$ for a well-fitting model. We will use the adjusted average score to represent the privacy rating at a certain level of privacy concern.

The goodness-of-fit measure [8] is used to evaluate the quality of the learned model, which aggregates the differences between the model prediction and actual user settings at each specific level θ_i. If the prediction error follows a Chi-square distribution with p-value < 0.05, the model fits the data well.

III. Mining Privacy Settings to Model Privacy and Utility

Our purpose of this research is two-fold. First, we want to understand the relationship between privacy settings and users' levels of privacy concern. Second, we want to study the tradeoff between privacy and utility in different privacy settings. In this section, we will focus on the modeling methods and then present the experimental results on real datasets in the next section.

A. Definitions and Notations

Each person's account has a number of profile items that can be set to different levels of visibility. These items can be email, birthdate, phone number, religious view, etc., which are denoted as $\{I_1, I_2, \ldots, I_k\}$. In Facebook, these items can be set to visible by "everyone", "friends of friends", "friends only", or "only me", as shown in Figure[1] 2. For simplicity, we assume the disclosure setting is binary by merging some settings. It is not difficult to extend the model to multiple settings [10]. Considering "friends of friends" are still strangers, we use "disclosed" to represent "everyone" or "friends of friends", and "hidden" to "friends only" or "only me". Assume there are n users. The setting for item I_j of user U_i is denoted as s_{ij}. Thus, n users will generate a $n \times k$ matrix of privacy setting.

In studying the level of privacy concern, we set "hidden" to 1 to represent the preservation of privacy, and "disclosed" to 0 to represent the loss of privacy. We call the binary setting vector $(s_{i1}, \ldots, s_{ik})$ for a specific user U_i ***a privacy configuration***. One of the problems for a new user of SNS is to find a right privacy configuration for his/her account. As disclosing or hiding one item is also directly related to the social utility

[1]Facebook's privacy control interface has been changed since we collected data.

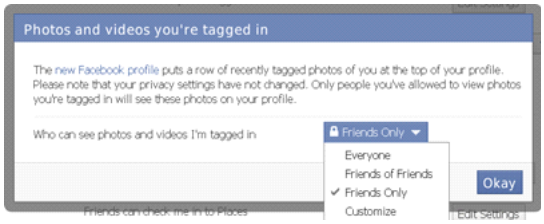

Fig. 2. A screen shot of facebook's privacy control.

of that item, privacy settings can also be used to represent users' implicit utility preference. Thus, when we study users' utility preference, we can set "disclosed" to 1 to represent preserved utility, and "hidden" to 0. Correspondingly, we call the flipped binary setting vector ***a utility configuration***. We will mine these datasets to discover the intrinsic relationships.

B. Modeling Users' Privacy Concern

Modeling users' privacy concern is the first step to understand the tradeoff between privacy and utility. This task can help understand how the privacy settings are linked to the SNS users' level of privacy concern, which will be extended to include the utility preference. Intuitively, the more profile items the user tries to hide, the higher level of privacy concern the user may have. However, items may have different weights in evaluating the privacy concern. For example, the items such as phone number have well-perceived privacy risk, and thus most users tend to hide these items. However, highly privacy conscious users may also try to hide some items that most other users think OK to disclose. As a result, two users with the same number of hidden items may have very different levels of privacy concern, and the total number of hidden items may not accurately reflect the user's privacy concern. The latent trait theory can be applied to exactly address this kind of problem.

In our study, we define the latent trait as the level of user's privacy concern, denoted as θ, and each profile item as the time in the latent trait model. Note that the latent trait is a global variable consistent for the privacy settings of different items, and each user U_i can only be described with a specific level of privacy concern θ_i. The item parameters α_j and β_j, $j = 1 \ldots k$, and the user's privacy concern θ_i, $i = 1 \ldots n$, can be jointly learned from the privacy setting data.

Figure 3 shows the probability distributions of three items among 27 items learned from our experimental datasets. The result of modeling gives a quantitative description on users' privacy concern. Although the values on X-axis have no absolute meaning, they represent a relative measurement of the level of privacy concern. The item "relationships" has the smallest β. It implies that even users with low level of privacy concern will set this item to "hidden". In other words, the privacy risk of disclosing it is well perceived by most users. In contrast, the item "current_city" is only set to "hidden" by some users who have high level of privacy concerns. These users are more concerned with their privacy than most other users, while most users may consider that disclosing this item has very low privacy risk and may bring more utility, e.g.,

getting connected with unknown people in the same city. Different from the other two, the item "networks" has a very low discrimination factor, meaning it is not an important factor for describing the level of user's privacy concern.

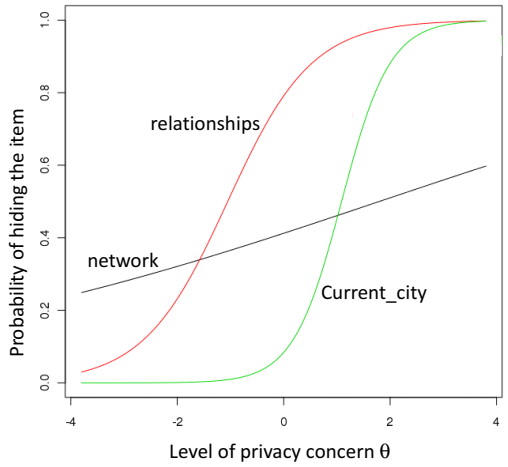

Fig. 3. β of the item model can be explained as how sensitive the item is perceived by users. The lower the β is, the more users want to hide the item.

Finding Optimal Privacy Configuration. With the latent trait model, each user U_i will get an estimated level of privacy concern, θ_i. The weighted average number of hidden items for all users at the level θ_i is uniquely defined as $p_{\theta_i} = \sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j)$, where α_i serves as the weight for the items I_i and $Pr(\theta_i; \alpha_j, \beta_j)$ is the value at θ_i on the curve of item I_i. We define this measure as *the privacy rating* at the privacy concern level θ_i. There is a relationship between the actual weighted number of hidden items and the privacy rating, according to the theory of latent trait model [8].

Theorem 1: For an ideally fitting model, the actual weighted number of hidden items $\sum_{j=1}^{k} \alpha_j s_{ij}$ for the user U_i is the privacy rating at the level of θ_i, i.e.,

$$\sum_{j=1}^{k} \alpha_j s_{ij} = \sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j) \tag{2}$$

where θ_i is the user U_i's level of privacy concern.

This theorem gives the basis for finding the optimal privacy configuration for a desired level of privacy concern. A new user can specify his/her level of privacy concern θ based on the relative measure in the range [-4,4]: we can map the level of privacy concern to the nine grades, from the lowest level of concern -4 to the highest 4. With the specific θ, Algorithm 1 can be used to find the optimal privacy configuration based on the fellow users' privacy concerns. Algorithm 1 assumes the item models $Pr(\theta; \alpha_j, \beta_j)$, $j = 1 \ldots k$ have been learned from the training data. Also, the level of privacy concern for each user in the training data has been calculated as θ_i, $i = 1 \ldots n$. The algorithm first calculates the privacy rating p_θ for the user-specified level of privacy concern θ with the item models, and then searches over the training dataset to find out the users having the similar level of privacy concern $|\theta - \theta_i| < \epsilon$, where $\epsilon > 0$ is a very small value. These users' privacy configurations form a candidate set, from which the optimal privacy configuration is found so that the weighted number of hidden items $\sum_{j=1}^{k} \alpha_j s_{ij}$ is close to the privacy rating $\sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j)$, according to Theorem 1.

Algorithm 1 OptPrivacyConfig($\theta, \{\theta_i\}, \{(\alpha_j, \beta_j)\}, S$)

Input: θ: the desired level of privacy concern; $\{\theta_i\}$: the levels of privacy concern for the users in the training data; $\{(\alpha_j, \beta_j)\}$: the parameters for the item models; S: training data
$p_\theta \leftarrow \sum_{j=1}^{k} \alpha_j Pr(\theta; \alpha_j, \beta_j)$;
$T \leftarrow$ find the set of θ_i so that $|\theta - \theta_i| < \epsilon$;
$d_{opt} \leftarrow \sum_{j=1}^{k} \alpha_j$;
OptConfig $\leftarrow \emptyset$;
for each user U_i having $\theta_i \in T$ **do**
  $p1 \leftarrow \sum_{j=1}^{k} \alpha_j s_{ij}$;
  if $d_{opt} > |p_\theta - p1|$ **then**
    $d_{opt} \leftarrow |p_\theta - p1|$;
    OptConfig $\leftarrow U_i$'s privacy configuration;
  end if
end for
return OptConfig;

C. Modeling User's Utility Preference

Note that Algorithm 1 only gives the way to find the privacy configuration for a specific level of privacy concern, without considering the factor of utility. The other important goal of our research is to understand the tradeoff between privacy and utility, which gives a comprehensive guidance for users to set the disclosure level of their private information.

Simple Utility Model. The next task is to model user's utility preference. We notice that utility and privacy is a pair of complementary factors - once one decides not to disclosure an item, any social utility with this item is disabled as well. Based on this understanding, we propose the following simple model.

We transform the data to reflect users' *implicit* consideration on utility preference. A privacy setting of "disclosed" now means preserving the utility of that item, while "hidden" means a loss of utility. They are mapped to 1 and 0, respectively, which is a flipped version of the data for privacy modeling. Let $\bar{s_{ij}}$ be such a setting used for utility modeling, i.e., $\bar{s_{ij}} = 1 - s_{ij}$.

We use the same latent trait modeling method to derive the simple model for utility preference. For clear presentation, we use different symbols to represent the utility models. Let ϕ be the level of utility preference and (λ_j, μ_j) be the pair of parameters for item I_j. The utility item model is

$$Pr(\phi; \lambda_j, \mu_j) = \frac{1}{1 + \exp\{-\lambda_j(\phi - \mu_j)\}}. \tag{3}$$

The smaller the value of μ_i is, the more people choose to set the item "disclosed". Similarly, λ_j serves as the discrimination factor as α_j does for privacy models. Because of the use of

the flipped version of data, it is easy to derive that $\lambda_j = \alpha_j$ and $\mu_j = -\beta_j$.

We can also similarly define the *utility rating*, u_{ϕ_i}, for the ideal utility configuration at the level of utility preference ϕ_i, which is $\sum_{j=1}^{k} \lambda_j Pr(\phi_i; \lambda_j, \mu_j)$. According to Theorem 1, we have

$$\sum_{j=1}^{k} \lambda_j Pr(\phi_i; \lambda_j, \mu_j) = \sum_{j=1}^{k} \lambda_j \bar{s}_{ij} \tag{4}$$

for ideal model fitting. Because $\bar{s_{ij}} = 1 - s_{ij}$, it follows

$$\begin{aligned} &\sum_{j=1}^{n} \alpha_j Pr(\theta_i; \alpha_j, \beta_j) + \sum_{j=1}^{k} \lambda_j Pr(\phi_i; \lambda_j, \mu_j) \\ &\approx \sum_{j=1}^{k} \lambda_j + \sum_{j=1}^{k} (\alpha_j - \lambda_j) s_{ij} = \sum_{j=1}^{k} \lambda_j, \end{aligned} \tag{5}$$

which implies that the privacy and utility ratings are linearly related under the simple utility modeling. This relationship is critical for making the tradeoff between privacy and utility.

Personalized Utility Model. Users often have clear intention for using social network services, while having less knowledge on privacy risks. Among the set of well-known social utilities, a new user may want to specify his/her utility preference, and wish it is considered in utility modeling. For simplicity, we describe the utility preference as the importance weights on the items, denoted as a vector $w = (w_1, w_2, \ldots, w_k)$. For the items less relevant to the utility, the user can assign 1 to them, while for a few relevant items, the user can assign a value larger than 1 to overweight them. Note that the user can also demote certain category of utility by setting a weight less than 1 for the related items.

We develop an algorithm to incorporate these weights in learning the personalized utility model. We want to overweight those samples consistent with the target user's preference. The key is to appropriately revise the likelihood function to incorporate the weights for learning. Let's start with the formulation of the likelihood function. Since the privacy setting is a binary variable, the probability to set item I_j to $\bar{s}_{ij}$ follows the Bernoulli distribution:

$$\begin{aligned} & Pr(S = \bar{s}_{ij} | \phi_i, \lambda_j, \mu_j) \\ = \; & Pr(\phi_i; \lambda_j, \mu_j)^{\bar{s}_{ij}} (1 - Pr(\phi_i | \lambda_j, \mu_j)^{1 - \bar{s}_{ij}} \\ = \; & \frac{\exp\{-\lambda_j(\phi_i - \mu_j)(1 - \bar{s}_{ij})\}}{1 + \exp\{-\lambda_j(\phi_i - \mu_j)\}}. \end{aligned} \tag{6}$$

Often, the training of latent trait model holds an oversimplified assumption that a user's setting for an item is independent of other users and other items [8]. With this assumption, we have the likelihood function

$$\begin{aligned} L(\lambda, \mu, \phi | S) \; &= \; \prod_{i=1}^{n} \prod_{j=1}^{k} Pr(S = \bar{s}_{ij} | \phi_i, \lambda_j, \mu_j) \\ &= \; \frac{\exp\{\sum_{i=1}^{n} \sum_{j=1}^{k} (-\lambda_j(\phi_i - \mu_j)(1 - \bar{s}_{ij})\}}{\prod_{i=1}^{n} \prod_{j=1}^{k} (1 + \exp\{-\lambda_j(\phi_i - \mu_j)\})} \end{aligned}$$

where λ, μ, and ϕ represent $\{\lambda_j\}, \{\mu_j\}, \{\phi_i\}$, $j = 1 \ldots k$, $i = 1 \ldots n$, respectively. Overweighting the item I_j by w_j ($w_j > 1$) is equivalent to adding more samples of I_j setting to the training data, which can be directly reflected in the following likelihood function.

$$L(\lambda, \mu, \phi | S, w) = \frac{\exp\{\sum_{i=1}^{n} \sum_{j=1}^{k} (-\lambda_j(\phi_i - \mu_j) w_j (1 - \bar{s}_{ij})\}}{\prod_{i=1}^{n} \prod_{j=1}^{k} (1 + \exp\{-\lambda_j(\phi_i - \mu_j)\})} \tag{7}$$

We revise the original MLE learning algorithm [8] to use this likelihood function. Due to the space limitation, we will skip the details of the learning algorithm, and will show some experimental results later.

D. Tradeoff between Privacy and Personalized Utility

Note that for each user in the training data, we can derive a pair of privacy concern and utility preference (θ, ϕ), which corresponds to a pair of privacy rating and utility rating (p, u). These pairs represent the tradeoffs between the privacy and utility. We want to find the ones that satisfy the specific level of privacy concern and, meanwhile, maximizing the utility.

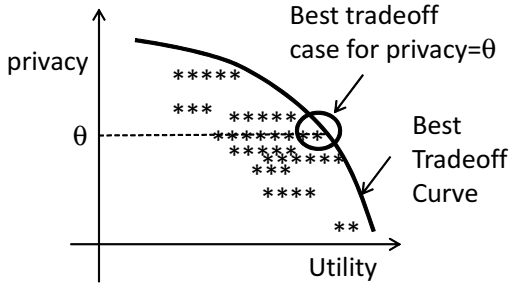

Fig. 4. Illustration of finding the best tradeoff cases.

Figure 4 illustrates a sample distribution of (p, u) pairs as we observed in experiments. A new user can specify a desired level of privacy concern θ. With the (p, u) pair distribution, we can find the pair and the corresponding privacy configuration. The following Algorithm 2 can be used to find the optimal privacy configuration with user's specification on utility preference.

IV. Experimental Studies

The experimental study will be focused on three aspects: (1) preparing the real data and study the characteristics of privacy settings, (2) using the data to train the proposed privacy models and studying the properties of the item models; and (3) understanding the tradeoff between privacy and utility with simple and weighted utility models.

A. Preparing Data

The first challenging problem is to get real privacy settings for a sufficiently large number of users. Asking users to do survey is a time consuming process, and is often limited to very small scale [11]. We need a large number of samples to reduce the model bias. We design a novel approach to

Algorithm 2 OptTradeoff(θ, S, w)

Input: θ: desired level of privacy concern; S: n users' privacy settings on k items; w: the weight vector for utility preference: $(w_1, w_2, \ldots, w_k)$.
$\{\alpha_j, \beta_j, \theta_i\} \leftarrow$ Learn the privacy model with S;
$\{\lambda_j, \mu_j, \phi_i\} \leftarrow$ Learn the personalized utility model with S and w;
for each user U_i **do**
  $p_i \leftarrow \sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j)$;
  $u_i \leftarrow \sum_{j=1}^{k} \lambda_j Pr(\phi_i; \lambda_j, \mu_j)$;
  $t_i \leftarrow (U_i, p_i, u_i, \theta_i, \phi_i)$;
end for
$T \leftarrow$find the cases t_i that $|\theta - \theta_i| < \epsilon$;
$t_{opt} \leftarrow$find the case in T that has the maximum u_i;
return t_{opt}'s corresponding user U_{opt}'s privacy configuration;

crawl and derive real privacy settings from a sufficiently large number of Facebook accounts. This approach utilizes the friends of friends (FoF) networks of Facebook accounts and derives the item privacy settings based on the visibility of the item crawled using the normal account and a fake account with no friend.

Concretely, the method is described as follows. First, we login to Facebook with a normal account and record all the FoF users of this login user and at the same time crawl the visible profile items of these FoF accounts. In order to reduce sampling bias, we have selected three independent user accounts for this data crawling purpose, and in the data cleaning process, FoF accounts are selected for each account such that the percentage of the overlap is around 5%. The number of FoF accounts of one Facebook user often reaches thousands. Table I shows the number of FoF accounts reached by the three normal Facebook accounts after cleaning the crawled data. Second, we login to the fake user's account that has no friend, and crawl the profile items of the recorded FoF accounts of normal users again.

	Account 1	Account 2	Account 3
# of friends	99	83	73
Original # of FoF	4104	5287	4376
Cleaned # of FoF	3798	4215	3552

TABLE I
THE NUMBER OF FRIENDS AND FRIENDS OF FRIENDS FOR THREE ACCOUNTS.

Visible to real user?	visible to fake user?	privacy setting
Y	Y	E
Y	N	FoF
N	Y	N/A
N	N	O/F

TABLE II
BASED ON THE ITEM VISIBILITY OF TWO ACCOUNTS, WE CAN DERIVE THE PRIVACY SETTING. "Y" FOR YES, "N" FOR NO AND "N/A" FOR IMPOSSIBLE.

We can derive the privacy settings of the FoF accounts based on the visibility from the real user and the fake user's accounts. At the time we crawl the data, Facebook's privacy setting for each profile item has four levels: only the account owner who can see the item (O), only the friends (F), friends of friends (FoF), and everybody in the social network (E). If the real user can see a profile item of his/her friend of friend, we can simply infer that the friend of friend has set this item's visibility to FoF or E. Otherwise, the friend of friend may set the visibility to O or F (or simply does not provide). The following inference matrix (Table II) can be used to derive the real privacy setting for the real user's FoF accounts.

Note that for items invisible to both the real user and fake user, the observed user may not even set the item values. This happens normally for items such as "political views" and "graduate school" (e.g., the user did not attend any graduate school). However, it is consistent to treat them as "hidden" cases.

According to the inference matrix, we can surely derive three levels of privacy setting: O/F, FoF and E for each item. For simplicity, we merge two of the settings to derive the binary settings: "disclosed" and "hidden". There are two possible merging schemes: (1) E $\rightarrow$ "disclosed"; O/F or FoF $\rightarrow$ "hidden"; or (2) E or FoF $\rightarrow$ "disclosed"; O/F $\rightarrow$"hidden". We consider friends-of-friends still strangers, and thus use Scheme (2) in our experiments.

We can crawl 27 profile items, including bio, birthday, family, relationships, religion views, political views, websites, photos, videos, links, notes, hometown, current city, education (high school, college, and graduate school), employers, activity, interests, phone number, sex, facebookuri, and networks. In the default settings, except for "interests", "political views", "religious views", and "phone number" set to "friends only", all other items are set to "everybody".

The collected dataset is then cleaned and analyzed to find the statistics of privacy settings. To determine how many profile items are changed for each account, we compare the derived settings with the default settings. Figure 5 shows the distribution of the number of items changed by the users crawled from Account 1. Because we consider no-value items as "hidden" items, Figure 5 shows different distribution from previous study [11]. Among 3798 friends of friends of Account 1, everyone has changed some of the profile items, while previous study shows about 20-30% users not changing any privacy setting [12]. As we have mentioned, this is because we treat no-value items as "hidden" items.

Figure 6 shows the list of items sorted by the percentage of users setting the item to "disclosed" (E or FOF). The figure shows that almost all the crawled users have disclosed "friend list", indicating users' lack of knowledge on the implications caused by this item, and on the contrary, very small percentage of users have disclosed address and phone number, possibly because their associated privacy risks are well perceived.

B. Modeling User's Privacy Concern

In this set of experiments, we want to explore the relationship between the level of user's privacy concern and the item privacy settings.

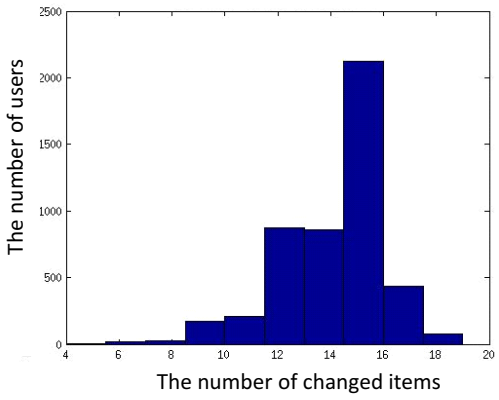

Fig. 5. Histogram of the number of changed profile items.

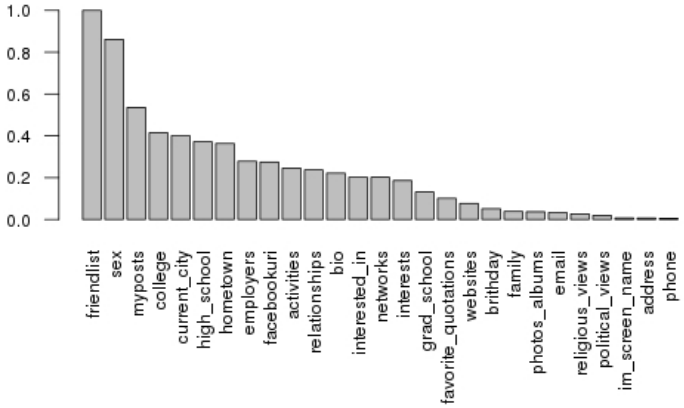

Fig. 6. The percentage of users sets the profile item to "disclosed".

The proposed modeling method (described in Section III-B) is applied to generate the item probability models. We pool the friends-of-friends records crawled from the three user accounts and perform a five-fold cross validation to check the model validity. Specifically, the data are randomly partitioned into five shares in equal size. In each round, we use four of the five shares for training and one share for testing. The testing result is evaluated by the goodness-of-fit measure [8]. The p-value is used to check whether the model fits the testing data well (p-value $<$0.05 suggests good fit). The testing results on the five folds show the models are in high quality, with all p-values $<$0.02.

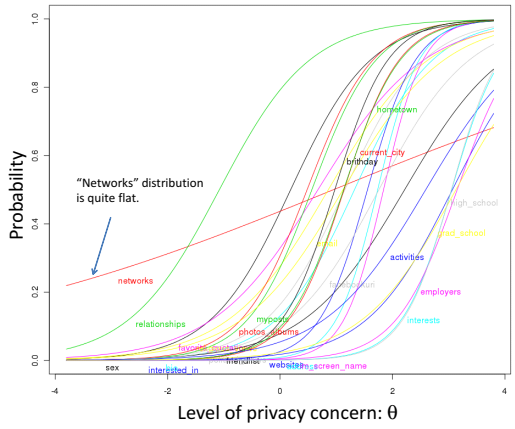

Fig. 7. The sketch of the item models for privacy modeling.

According to the latent trait theory, the items with β values in [-4, 4] are the most informative ones. Figure 7 shows the sketch of the learned item distribution models with β in [-4,4]. Small β means many users tend to hide the item - the privacy risk is highly perceived. Large β indicates only highly privacy conscious users tend to hide the item and the privacy risk of that item might not be well perceived. In the item models trained with the pooled data, we found the items "family" and "phone" have β values smaller than -4, which means most users tend to hide these items. "college" has a β value larger than 4, which makes sense since Facebook is originally designed for college students and thus "college" is normally disclosed to distinguish where users are from. We select the top items with the lowest β and the highest β in [-4, 4], as shown in Table III. "favorite_quotations" appears in the lowest five, partially because most people do not have favorite quotations. It appears very few people try to hide interests, education, and employers, as they are also the common things that people like to share with other people in real life.

The lowest five	sex, photos_album, myposts, favorite_quotations, birthday
The highest five	interests, interested_in, grad_school, high_school, employers

TABLE III
TOP ITEMS WITH THE LOWEST AND HIGHEST β IN THE RANGE [-4, 4] FOR PRIVACY MODELING.

The other parameter α serves as a discrimination factor, showing how much the item can contribute to distinguishing the level of privacy concerns. They are important in calculating the privacy rating of a privacy configuration. We found most items have α parameters around [1, 2.5]. Only "family", "college", "networks", and "phone number" have α less than 1. This makes sense because most people hide "family" and "phone number", and most people disclose "college" according to their β values. Although "networks" has an appropriate β ($\beta = 1.79$), its setting seems a random choice, resulting a low α value.

We can also learn each user's level of privacy concern θ_i from the data. Figure 8 shows the distribution of the number of users according to the level of privacy concern. There is no user with $\theta < -2$. Most users' are around the median $\theta = 0$ and in the range [-1.5, 2]. Interestingly, there is a spike around $\theta = -1.2$, and there are also very few persons having $\theta > 2$. Overall the mean θ is -0.063.

Figure 9 shows the relationship between the number of hidden items and the privacy rating. Because each user has a unique θ_i, we can plot the weighted number of hidden

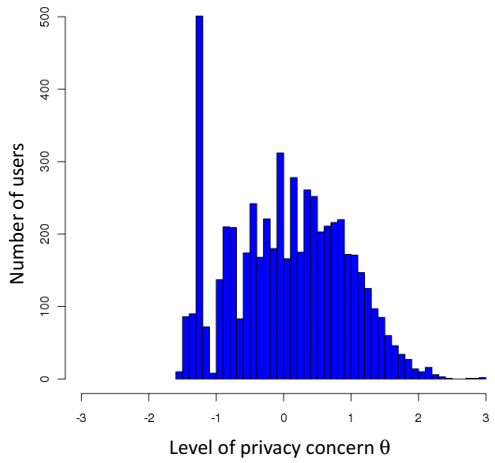

Fig. 8. The number of users at different levels of privacy concern: θ.

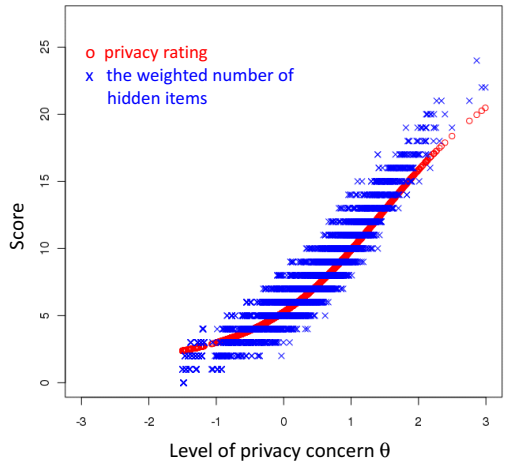

Fig. 9. Comparing the privacy rating and the weighted number of hidden items. The optimal privacy configuration can be found from the blue 'x' points that are close to the red 'o' points.

items, $(\theta_i, \sum_{j=1}^{k} \alpha_j s_{ij})$, for the user's privacy configuration. Meanwhile, for each θ_i, there is a unique privacy rating $p_{\theta_i} = \sum_{j=1}^{k} \alpha_j Pr(\theta_i; \alpha_j, \beta_j)$ (the red curve in the figure). The figure shows the weighted number of hidden items are around the corresponding privacy rating. We can also find the optimal privacy configurations around the points on or close to the curve of privacy rating.

C. Characteristics of Our Sample Data

One may have questions on the bias of our data collection method, i.e., sampling the FoF data for only three accounts is not representative enough. We argue that we only use this dataset to demonstrate our methods, not intent to generate global models for the whole Facebook community. However, the FoF data collected from different accounts indeed show some interesting characteristics, which suggest data at this scale (e.g., thousands of examples) might be sufficient to capture the global privacy-concern model.

	TestData1	TestData2	TestData3
Model1	2.624510	2.470485	2.416606
Model2	1.990507	1.725497	1.674528
Model3	2.701275	2.311214	2.122907

TABLE IV
CROSSING-MODEL TESTING RESULTS REPRESENTED WITH RMSE. MODEL i IS TRAINED WITH TESTDATA i. THE CLOSE RMSES IN EACH ROW MEANS THE CORRESPONDING MODEL HAS SIMILAR GOODNESS-OF-FIT FOR THE THREE SETS OF DATA.

	Model1	Model2	Model3	Combined
Mean of θ	-0.0625	-0.0680	-0.0521	-0.0659
Variance of θ	0.7087	0.6765	0.7031	0.6884

TABLE V
MEAN AND VARIANCE OF USERS' θ VALUES FOR DIFFERENT MODELS.

We use the following method to understand the similarity between the three sets of FoF data. First, we train the privacy concern model for each of the three accounts. Then, we test each model with the three sets of data to find the tuple $(\theta_i, p_{\theta_i}, \sum_{j=1}^{k} \alpha_j s_{ij})$ for each user U_i in the dataset. p_{θ_i} is the privacy rating at θ_i, and $pr_i = \sum_{j=1}^{k} \alpha_j s_{ij}$ is the real privacy score according to this user's privacy configuration. If a model fits the data well, then p_{θ_i} and pr_i will be very close. Assume there are n instances in the testing dataset. We define the *root-mean-square-error* (RMSE) as $\sqrt{1/n \sum_{i=1}^{n} (pr_i - p_{\theta_i})^2}$. Table IV shows that for each model the three test datasets result in very close RMSE.

We also check the distributions of resultant users' θ values and find their mean and variance values are also very close (Table V, where "combined" means the model learned from the combined data.). These results all suggest a certain level of similarity between the three datasets. In addition, because we have carefully selected the three accounts to make their FoFs have low overlapping ($< 6\%$ overlapping for each pair of accounts), it seems that each dataset is sufficiently large to capture some shared global model. We will conduct more experiments to verify this observation.

D. Utility Modeling and Tradeoffs

Users often consider privacy settings based on an implicit tradeoff between privacy and utility. As we have discussed in Section III-C, users' privacy settings can also be used to model their utility preference. We use the same datasets with flipped values ("disclosed" is set to 1 and "hidden" to 0) to model utility preference. We train the simple utility model with this flipped dataset. As we have expected, the resultant item models satisfy the relationship $\lambda_j = \alpha_j$ and $\mu_j = -\beta_j$. Due to the similarity with the privacy model, we skip the details for the simple utility model.

We can tune the utility model towards a specific user's

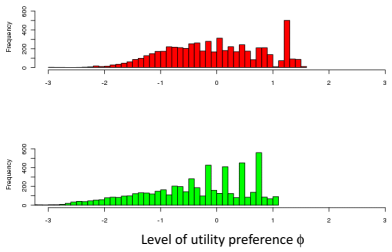

Fig. 10. Top: the the number of users at different levels of utility preference for unweighted utility model. Bottom: for weighted utility model (overweighting "phone", "email", and "employers" by a factor 1.1)

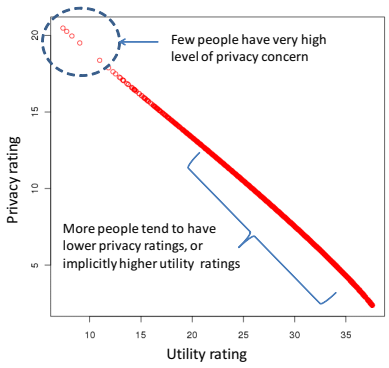

Fig. 11. The tradeoff between the privacy and utility for unweighted utility model.

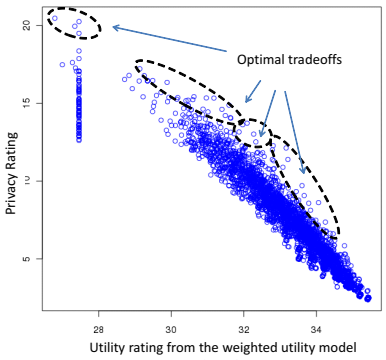

Fig. 12. Each point in the tradeoff graph represents one user. The points around the upper bound contain the optimal tradeoffs.

preference of utility. A user may specify a special use of social network services, such as job hunting. We then provide a tradeoff function that is biased to the specific utility. This biased utility model is achieved by overweighting the items related to the specified utility and learning the weighted utility model as we described in Section III-C.

In the experiment, we overweight "employers", "phone" and "email" by 1.1 to make the privacy settings friendlier to job hunting. The weight 1.1 can be understood as adding 10% more training samples that have the overweighted items disclosed. The resultant weighted model well fits the data with p-value < 0.02. How to pick weight is a complicated issue, which will be addressed by our future work. Here, we just want to show how the result looks like by introducing the weights.

By comparing the weighted model and the original model, we found the β parameters for the weighted items are slightly reduced and the α parameters are significantly increased. Figure 10 shows the change of the number of users at different levels of utility preference.

Tradeoffs. Without using the weighted scheme, we have shown that the privacy rating is approximately linearly correlated with the utility rating. Figure 11 shows the relationship derived from the real dataset, which confirms our analysis. Overall, the distribution is biased to the utility side. However, this may not imply that the users care utility more than privacy. More likely, many users are not aware of the privacy risks and thus leave the default settings of most items unchanged. As Barnes [7] mentioned, most young people care less about their privacy because they are unaware of privacy risks. Because most Facebook users are young, it explains the observed pattern. Other studies [13], [14] show that different groups of people may have different preferences over privacy and utility of social network services, which will be explored in our future work.

Figure 12 shows the relationship between privacy and the weighted utility (with the specified items overweighted by 1.1, and other items by 1), which is not linear anymore. Note that the points around the upper bound of the distribution represent the most interesting cases. The tradeoff algorithm (Algorithm 1) takes these cases and generates optimal tradeoffs for a specific level of privacy concern.

V. Related Work

Privacy leaking may result in serious attacks to the SNS users. Spamming and phishing [15], [16], [17], [1], [18], are probably the most popular attacks that utilize private information, which may cause significant financial loss of the targeted users. As more and more people are becoming SNS users, SNS providers should seriously consider the impact of SNS privacy on the whole society and make privacy setting easier to users.

Fang et al. [19] consider the difficulty of tuning privacy settings for many profile items, especially for a new SNS user. They propose to learn classification models from existing users of the SNS. The classification models try to find the correlation between the user's account features, such as age, sex, and education level, and the privacy setting for each profile item. New users can use the classification models to automatically determine their privacy settings. There are several problems with this approach. First, the classification models depend on and are biased by the training data, which do not incorporate the new user's privacy preference. Second, it is difficult to evaluate whether the classification models suggest appropriate settings, when the training data does not support high-accuracy classification models. Third, this approach does not consider a privacy setting is a tradeoff between privacy and utility.

Liu et al. [20] propose a framework for computing privacy score of users based on users' privacy settings. The latent trait model is also used to model the privacy score. Similarly, the privacy setting data is used to generate the item models. The sensitivity of an item is defined as the item model parameter β_j. The visibility concept is combined with sensitivity, which is the privacy setting s_{ij} for item I_j by user U_i. The privacy score is defined as $\sum_{j=1}^{k} \beta_j s_{ij}$, which looks like the weighted number of hidden items in our approach. However, the formula is difficult to explain with the latent trait modeling theory. In addition, this approach does not consider the utility factor as

well.

Yabing Liu et al. [11] study privacy settings of SNS users by using survey data. The survey dataset is collected from 200 Facebook users recruited via Amazon Mechanical Turk. Their results show that most users use default settings provided by the SNS, and users have trouble to correct the default privacy settings, which supports the importance of our study.

SNS users' attitudes to the privacy issues are studied by Barnes [7]. The author observes the privacy paradox problem of online social network: adult people are concerned about the invasion of privacy while teenage people freely give up their personal information. The main reason for the paradox is that teens are usually inexperienced and thus unaware of the problems. Our study shows that overall people care more about utility of Facebook than the privacy issues, which is probably due to the fact that most users of Facebook are young people.

The tradeoff between utility and privacy is also studied by Norberg et al. [13]. Their survey data reveals that people indeed provide significantly greater amounts of personal information than they expected. They find that perception of risks has a negative effect on peoples' information sharing behavior, while on the other hand perceived trust or benefit doesn't have a clear positive effect on people's sharing behavior. Utz et al. [14] also studied the personal factors such as personal impression, narcissism and their relation to users' social network information sharing behaviors.

VI. Conclusions and Future Work

In this paper, we propose a framework to address the tradeoff between privacy and utility for users' privacy settings of a social network service. We use the latent trait theory to model users' levels of privacy concern and levels of utility preference. We develop an algorithm to incorporate users' personalized utility preference in learning the utility model. Finally, a tradeoff algorithm is developed for users to find the optimal privacy configuration based on the desired privacy level and utility preference. The whole approach is validated and showcased with a large real dataset crawled from Facebook.

The future work will be conducted along several directions. First, we will continue to work on the problem of learning weighted utility model. In this paper, we did not answer how to set item weights for a user's utility preference, which appears intricate. Further study is needed to find (or learn) the appropriate weighting scheme for any specific type of utility, so that users do not need to worry about the utility weights. Second, the community structure in social networks can be used to group users. We believe that this structure may result in different privacy concern models. We would like to study this factor and incorporate it into our analysis framework.

References

[1] B. Markines, C. Cattuto, and F. Menczer, "Social spam detection," in *Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web*, ser. AIRWeb '09. New York, NY, USA: ACM, 2009, pp. 41–48. [Online]. Available: http://doi.acm.org/10.1145/1531914.1531924

[2] H. Gao, J. Hu, C. Wilson, Z. Li, Y. Chen, and B. Y. Zhao, "Detecting and characterizing social spam campaigns," in *Proceedings of the 10th annual conference on Internet measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 35–47. [Online]. Available: http://doi.acm.org/10.1145/1879141.1879147

[3] J. Williams, "Social networking applications in health care: threats to the privacy and security of health information," in *Proceedings of the 2010 ICSE Workshop on Software Engineering in Health Care*, ser. SEHC '10. New York, NY, USA: ACM, 2010, pp. 39–49. [Online]. Available: http://doi.acm.org/10.1145/1809085.1809091

[4] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer, "Social phishing," *Commun. ACM*, vol. 50, no. 10, pp. 94–100, 2007. [Online]. Available: http://markus-jakobsson.com/papers/jakobsson-commacm07.pdf

[5] B.Schiffman, "Facebook ceo apologizes, lets users turn off beacon," 2007. [Online]. Available: http://www.economist.com/node/15351002

[6] B. Slattery, "Google hit with lawsuit over google buzz," http://www.pcworld.com/article/189712/google_hit_with_lawsuit_over_google_buzz.html, Feb. 2010.

[7] S. B. Barnes, "A privacy paradox: Social networking in the United States," *First Monday*, vol. 11, no. 9, sep 2006. [Online]. Available: http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/1394

[8] R. J. de Ayala, *The Theory and Practice of Item Response Theory*. The Guilford Press, 2008.

[9] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer-Verlag, 2001.

[10] N. Longford, *Random Coefficient Models*. Oxford: Clarendon Press, 1993.

[11] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, "Analyzing facebook privacy settings: user expectations vs. reality," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, ser. IMC '11. New York, NY, USA: ACM, 2011, pp. 61–70. [Online]. Available: http://doi.acm.org/10.1145/2068816.2068823

[12] R. Gross and A. Acquisti, "Information revelation and privacy in online social networks," in *WPES '05: Proceedings of the 2005 ACM workshop on Privacy in the electronic society*. New York, NY, USA: ACM, 2005, pp. 71–80.

[13] P. A. Norberg, D. R. Horne, and D. A. Horne, "The privacy paradox: personal information disclosure intentions versus behaviors," *Journal of Consumer Affairs*, vol. 41, no. 1, 2007.

[14] S. Utz and N. C. Kramer, "The privacy paradox on social network sites revisited: The role of individual characteristics and group norms," *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, vol. 3, no. 2, 2009.

[15] C. Grier, K. Thomas, V. Paxson, and M. Zhang, "@spam: the underground on 140 characters or less," in *Proceedings of the 17th ACM conference on Computer and communications security*, ser. CCS '10. New York, NY, USA: ACM, 2010, pp. 27–37. [Online]. Available: http://doi.acm.org/10.1145/1866307.1866311

[16] M. Huber, M. Mulazzani, E. Weippl, G. Kitzler, and S. Goluch, "Exploiting social networking sites for spam," in *Proceedings of the 17th ACM conference on Computer and communications security*, ser. CCS '10. New York, NY, USA: ACM, 2010, pp. 693–695. [Online]. Available: http://doi.acm.org/10.1145/1866307.1866400

[17] M. Huber, M. Mulazzani, S. Schrittwieser, and E. Weippl, "Cheap and automated socio-technical attacks based on social networking sites," in *Proceedings of the 3rd ACM workshop on Artificial intelligence and security*, ser. AISec '10. New York, NY, USA: ACM, 2010, pp. 61–64. [Online]. Available: http://doi.acm.org/10.1145/1866423.1866435

[18] F. Benevenuto, T. Rodrigues, V. Almeida, J. Almeida, C. Zhang, and K. Ross, "Identifying video spammers in online social networks," in *Proceedings of the 4th international workshop on Adversarial information retrieval on the web*, ser. AIRWeb '08. New York, NY, USA: ACM, 2008, pp. 45–52. [Online]. Available: http://doi.acm.org/10.1145/1451983.1451996

[19] L. Fang and K. LeFevre, "Privacy wizards for social networking sites," in *Proceedings of the 19th international conference on World wide web*, ser. WWW '10. New York, NY, USA: ACM, 2010, pp. 351–360.

[20] K. Liu and E. Terzi, "A framework for computing the privacy scores of users in online social networks," *ACM Trans. Knowl. Discov. Data*, vol. 5, pp. 6:1–6:30, December 2010. [Online]. Available: http://doi.acm.org/10.1145/1870096.1870102