Herd Privacy: Modeling the Spillover Effects of Privacy Settings on Social Networking Sites

David Masad

George Mason University, Fairfax, VA, 22030

Abstract. Concerns over privacy on social networking sites generally focus on the options available to individual users for protecting personal information on their own pages. However, similar information is often exposed on users' friends' pages in the form of public posts, comments or photographs. Users' actual privacy depends not just on their own settings but on their friends' settings as well.
In this paper I define a computational model for the exposure of personal information on a social networking site and use it to analyze the importance of aggregate privacy tendencies for individual privacy. I find that privacy settings have a spillover effect, and that exposure is further affected by network topology.

1 Background

Social networking sites (SNS) have been growing rapidly over the past decade. At present, Facebook claims 600 million users (Carlson 2011) with millions more using non-English services such as the Chinese QZone (Wauters 2009). Almost from the beginning, SNS user privacy has been a topic of both academic and popular interest. Part of this interest is driven by what Barnes (2006) described as the 'privacy paradox' which occurs when users' desire for privacy conflict with their increasing use of services which collect, store and often share their personal information.

Most SNS offer their users different privacy options. These vary from site to site; for example, Twitter only offers a choice between accounts being Public (viewable by everyone) and Protected (requiring manual approval of each user), while Facebook offers detailed settings which can be set to vary between types of content. Individuals simply interested in maximizing their privacy are unlikely to join an SNS to begin with; site users employ the available privacy settings in an attempt to achieve a desired balance between privacy and exposure.

Users are not always successful in achieving their desired level of privacy. Barnes (2006) noted that teenagers treated Facebook and other SNS as semi-private spaces without realizing the extent to which they exposed their information to parents, school administrators or future employers. Similarly, Acquisti and Gross (2006) found that a significant minority of Facebook users in their survey underestimated their actual exposure. More recently, an informal survey by boyd (2010) indicated that many Facebook users' mental models of their exposure did not match their actual exposure, though a survey of students by boyd

and Hargittai (2010) found that most had experience using Facebook's privacy settings.

Part of the difference between perceived and actual exposure may be explained by the effect of users' Friends'[1] privacy settings. In an early study of Facebook users at Carnegie Mellon University, Gross and Acquisti (2005) documented the high percentage of users who had Friends outside the university network, thereby exposing their CMU Friends to outsiders who were likely to be complete strangers (and vice versa). The boyd and Hargittai (2010) survey found that one of the most common activities on Facebook is commenting on the status updates of other users, both close friends and others; these comments then expose some degree of information to anyone who can view the Friend's page. A similar thing occurs whenever a user uploads photographs in which other users appear. Preibusch et al. (2007) even noted the tension between users desiring different degrees of privacy, and that currently SNS offer no mechanism for resolving such disputes.

Thus, it appears that at least some personal information is exposed via Friends' pages, whose privacy settings then have a spillover effect on the user's own exposure. I hypothesize that there is in fact a 'herd privacy' effect, in that each user's exposure is affected by the privacy preferences of the rest of the network.

It may be possible to analyze that effect from a large sample of user pages scraped from an SNS. However, this would risk biasing the result towards those users whose information is already the most exposed. Lampe et al. (2007) describe a similar attempt which did not find a way to retrieve information on users whose pages were hidden.

If we cannot obtain accurate real-world data, we must create our own. One way of doing so is with a computational model. By simulating the process by which personal information is exposed given different privacy preferences and network topologies, I hope to be able to estimate the magnitude of the 'herd privacy' effect in the protection or exposure of private user information.

2 Exposure Model

2.1 Informal Description

In order to estimate the importance of herd privacy, the model must capture key attributes of a social networking site as they relate to user exposure. While the model I propose is based on Facebook, it is meant to be generic and could easily be modified to represent different privacy structures.

The model treats the SNS as a network of mutual Friendships between users, all with their own pages. These users expose some personal information on their

[1] While different SNS use different terminology to refer to connected users, 'Friend' is the most common. This paper will follow boyd and Ellison (2007) and capitalize the term Friend when it refers to an SNS connection, as opposed to the colloquial and uncapitalized friend.

pages, visible to their Friends and (depending on whether their privacy setting are Closed or Open) to Friends of Friends as well. Additionally, the users expose personal information on the pages of their close friends – real-world examples of this would be comments, photographs, survey replies and even the Friendships themselves. This personal information is visible to anyone who has access to the Friends' pages.

If Alice and Bob are two users, Alice's total information on Bob is defined as the total of Bob's personal information available on the pages Alice can view. Note that this means that even if Alice is Friends with Bob, she may not have access to all of his personal information – some of it may be on Carol's page, which will be inaccessible to Alice if she and Carol are not Friends and Carol's privacy settings are closed.

2.2 Formal Description

The model will consist of an undirected **graph** of size N in which nodes are simulated **users** (and their personal pages) and edges are Friendship relationships between them. This means that $degree(i) \equiv$ the number of Friends user i has.

Each user page i is associated with an exposure vector $\mathbf{b_i}$ of length N, so that $b_{i,j} \in [0, 1]$ represents the personal information on user j exposed on page i **and**

$$\sum_{i=1}^{N} b_{i,j} = 1. \tag{1}$$

Each user has a **privacy setting** which determines whether the information held on their page is viewable only by their Friends or by Friends-of-Friends as well. A fraction $f \in (0, 1)$ of each user's Friends are actual **close friends** (F) whose own pages hold (and expose) some proportion of the user's personal information, such that $f \cdot degree(i)$ is an integer, and:

$$\forall j \in F_i : b_{j,i} = \frac{1 - b_{i,i}}{f \cdot degree(i)} \tag{2}$$

Each user also has a **knowledge vector $\mathbf{k_i}$** of length N so that $k_{i,j}$ represent the exposure that user i has to user j's personal information. Formally:

$$k_{i,j} = \sum_{n=1}^{N} (b_{n,j} \cdot v_{i,n}); \text{where } v_{i,n} = \begin{cases} 1 & i \text{ has visibility to } n \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

A user's overall exposure is then defined as their mean exposure to each other user; formally:

$$exposure_i = \frac{\sum_{n=1}^{N} k_{n,i}}{N - 1} \tag{4}$$

2.3 Model Parameters

Each iteration of the model will be initialized with a random graph; I use two graph-generating algorithms and compare the results. The first is the Watts-Strogatz small-world model (Watts and Strogatz 1998): its strength is that it captures the role of highly-connected overlapping clusters of Friends within the network, and its weakness is that produces a fairly homogenous distribution of degrees, potentially failing to account for the role of central, high-degree individuals within the network. The second is the Holme-Kim model (Holme and Kim 2005) which essentially adds local clustering to a conventional scale-free network; while it does allow for highly-central users, it may underestimate local community clustering.

In order to focus on the *proportion* of personal information being exposed, I assume that all users expose the same fraction of their personal information on their own pages, with the remainder exposed on their close friends' pages as described above. Thus, $b_{i,i}$ will be a constant across all users.

Both graph-generating algorithms are seeded with an **initial edge count**, which determines the number of Friends each user has. The fraction f of Friends who are information-exposing close friends will be constant, and $f \cdot degree(i)$ will be rounded to the nearest integer.

Each user's privacy setting will be picked randomly, with $p \equiv$ the probability of any user having an **Open** page. This is the key parameter to be explored in the following section.

2.4 Initial Parameters

In order to ensure that the model approximates actual SNS topology, I fix several of the parameters above. The survey conducted by Golder et al. (2007) placed the median number of Facebook Friends at 144; while the overall number of users has grown significantly since 2005, Facebook (2011) itself currently reports that the average user has 130 friends, but does not qualify this result at all. Golder et al. (2007) further note that this is within the commonly accepted range for Dunbar's number, the theoretical limit on the number of social relationships a person can maintain (Dunbar 1998). Thus, I set the seed edge count as 144. Golder et al. (2007) report that only 15.1% of Friendships featured exchanges of messages; I interpret this as a proxy for close friendship, and so set $f = 0.15$. I fix $b_{i,i} = 0.60$ arbitrarily.

I set the network size $N = 1,000$ as a compromise between a desire for a comprehensive simulation and the practical limits of available computing resources. While this size yields small-diameter networks, this is not necessarily a weakness: Barnes (2006) notes that young people in particular are most concerned with protecting their privacy from certain peers, parents or teachers – that is, individuals likely to be closer to them in the network than unknown strangers. Strictly speaking, we may imagine the model to be of a particular subgraph of an SNS, such as a high school, small college or other community we expect to be highly internally-connected compared to the network as a whole.

I set p from 0 to 1 in increments of 0.1 and repeated the model 10 times at each value for each of the two network topologies, for a total of 220 iterations.

3 Results and Discussion

Let us first examine the results aggregated to the model iteration level. Figure 1 contains two points for mean exposure from each iteration: one for users with Closed privacy, the other for users with Open privacy.

Examining Figure 1 immediately tells us several things. First of all, the two network topologies we are comparing produce slightly different results, though their behaviors appear extremely similar. As we would expect, users with Open pages have significantly higher exposure than users with Closed pages. Importantly, we see that the mean exposure of users with *Closed* privacy settings increases as more users have Open pages.

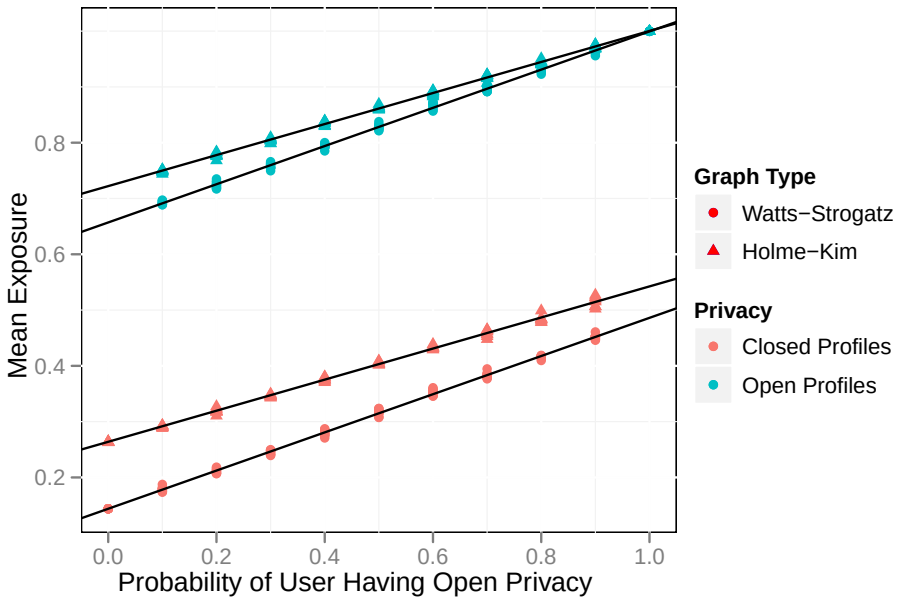

Fig. 1. Mean Exposures by Model

This bears repeating, since it gets to the heart of the research question: In this model, a major increase in users with Open pages could easily double the Closed-profile users' mean exposure, even as those user's own privacy settings remain the same. Note, furthermore, that if a user's *own* privacy settings were

to change inadvertently, their exposure would jump significantly more. In an extreme case when both were to happen simultaneously (as boyd (2010) suggests occurred following the change in Facebook's default privacy policy) exposure could increase to be nearly complete, regardless of where it began.

Now let us examine the model results at the level of the individual user. Given the strong linearity apparent in Figure 1, I will use a standard OLS linear regression, the results of which are in Table 1. In addition to each user's own settings and the overall fraction of users with Open pages, we can now examine the fraction of each user's Friends whose pages are Open. Furthermore, we can account for each user's place in the network topology; I will use closeness centrality [2] in order to capture the intuition that a user's potential exposure increases as they are Friends with more users who themselves have many Friends.

Table 1. User Exposure Linear Regressions

	Graph Type	
	Watts-Strogatz	Holme-Kim
(Intercept)	-0.383**	-0.124**
Pr(Open)	0.002	0.012*
Open Privacy	0.514**	0.458**
Fraction of Friends w/ Open Privacy	0.340**	0.266**
Closeness	0.977**	0.680**
R^2	0.99	0.98

* – significant at 5% level; ** – significant at 0.1% level

As we can see, the overall privacy in the network is actually barely significant on its own. What does matter are the privacy settings of each user's Friends; the more Friends a user has with Open settings, the more exposed that user will be.

Perhaps more importantly, this analysis shows that closeness centrality has a significant large effect on exposure – the more central users are within the network, the greater their exposure. While this may not be surprising, it has important implications. Real-world users are even less likely to be aware of their centrality than they are of their Friends' privacy settings, making it more difficult for them to estimate their exposure. Furthermore, Lampe et al. (2007) found that there is at least weak but positive correlation between the amount of information

[2] Closeness centrality is defined as the inverse of the average distance from a given node to all other nodes. Formally,

$$C(i) = \frac{N-1}{\sum_{n=1}^{N} dist(i,j)} \tag{5}$$

users post on their Facebook page and the number of Friends they have (and thus likely with their centrality as well). This would suggest that high-centrality individuals are even more exposed than this model implies.

4 Conclusions

This model offers strong support to the intuition that user privacy within social networking sites depends not only on the individual users' settings but on those of their Friends as well. It suggests that unexpected or poorly understood changes to default privacy settings may significantly change a user's exposure if the changes affect the user's Friends; however, changes to the network at large while a user's Friends are unaffected will not have significant consequences.

The model also highlights the role played by network topology and a user's place within it. Higher closeness centrality in particular was shown to have a strong positive effect on potential exposure. While high centrality may arise from popularity, boundary-spanners (including, for example, young people whose social identities are still in flux) are also more likely to exhibit high centrality, and hence face increased exposure.

Users on Facebook and other SNS often do not know their Friends' privacy settings, and cannot easily determine their position within the network topology. Consequently, they may not be able to properly assess their actual exposure. It is possible that increased awareness of privacy spillovers would encourage some behavioral changes, whether in the form of users modifying their privacy settings or simply becoming more aware of information they expose on pages they do not control. Allowing users to view their Friends' privacy settings, or information on their own position within the network, may also allow them to make more informed decisions concerning their privacy. Of course, these proposals raise their own set of problems and questions.

There is still much which can be done to understand SNS privacy. A dynamic expansion of this model could simulate the process by which users expose information through discrete actions over time, add and remove Friends, and adjust their privacy settings in response to external events or a desire for a specific target level of exposure.

Additional empirical data may advance our understanding even further. Data scraped or otherwise retrieved from actual SNS pages could provide more detailed information on the tendencies of users to expose personal information, as well as reveal the actual network topology within which they are doing so. This would be particularly valuable if the users' privacy settings could also be determined via surveys or technical means.

It seems that there will not be a lack of opportunity to conduct such research. Facebook and other social networking sites continue to grow, and examples of the 'privacy paradox' are likely to continue to grow right alongside them.

References

Acquisti, A., Gross, R.: Imagined communities: Awareness, information sharing, and privacy on the Facebook. In: Privacy Enhancing Technologies. p. 3658 (2006)

Barnes, S.B.: A privacy paradox: Social networking in the united states. First Monday 11(9) (Sep 2006), http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/1394/1312

boyd, d.: Making sense of privacy and publicity. In: SXSW. Austin, TX, USA (Mar 2010), http://www.danah.org/papers/talks/2010/SXSW2010.html

boyd, d., Ellison, N.: Social network sites: Definition, history, and scholarship. Journal of Computer-Mediated Communication 13(1), 210–230 (Oct 2007)

boyd, d., Hargittai, E.: Facebook privacy settings: Who cares? First Monday 15(8) (Jul 2010), http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3086/2589

Carlson, N.: Goldman: Facebook has 600 million users. Business Insider (Jan 2011), http://www.businessinsider.com/facebook-has-more-than-600-million-users-goldman-tells-clients-2011-1

Dunbar, R.: Grooming, gossip, and the evolution of language. Harvard University Press, Cambridge (Oct 1998)

Facebook: Statistics: People on facebook, http://www.facebook.com/press/info.php?statistics, [May 29, 2011]

Golder, S.A., Wilkinson, D.M., Huberman, B.A.: Rhythms of social interaction: Messaging within a massive online network. In: Steinfield, C., Pentland, B.T., Ackerman, M., Contractor, N. (eds.) Communities and Technologies 2007, pp. 41–66. Springer London, London (2007)

Gross, R., Acquisti, A.: Information revelation and privacy in online social networks. In: Proceedings of the 2005 ACM workshop on Privacy in the electronic society. p. 7180. WPES '05, ACM, New York, NY, USA (2005), ACM ID: 1102214

Holme, P., Kim, B.J.: Growing scale-free networks with tunable clustering. Physical Review E 65(2), 026107 (Jan 2002)

Lampe, C.A.C., Ellison, N., Steinfield, C.: A familiar Face(book): Profile elements as signals in an online social network. In: Proceedings of the SIGCHI conference on Human factors in computing systems. p. 435444. CHI '07, ACM, New York, NY, USA (2007), ACM ID: 1240695

Preibusch, S., Hoser, B., Gurses, S., Berendt, B.: Ubiquitous social networks opportunities and challenges for privacy-aware user modelling. In: Proceedings of Workshop on Data Mining for User Modeling. Corfu, Greece (Jun 2007)

Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. Nature 393(6684), 440–442 (Jun 1998)

Wauters, R.: China's social network QZone is big, but is it really the biggest? TechCrunch (Feb 2009), http://techcrunch.com/2009/02/24/chinas-social-network-qzone-is-big-but-is-it-really-the-biggest/