NC STATE UNIVERSITY

CSC 495.002 – Lecture 3
Web/Social Networks Privacy: Sharing and Disclosure

Dr. Özgür Kafalı

North Carolina State University
Department of Computer Science

Fall 2017

NC STATE UNIVERSITY

Inference

- Logic inference and how it applies to privacy
- How additional information about individuals can be inferred from known data
- How such inference can disrupt privacy or help reasoning (e.g., build connected communities)
- One medium scale study using Facebook data

Objectives

- Understand common usage scenarios of OSNs
- Identify sharing and disclosure patterns of users
  - What content types are shared?
  - Whom are they shared with?
  - How do sharing behaviors change over time?
- Determine whether shared content matches intended audience
- Understand how users mitigate privacy concerns
- Awareness of “silent listeners”, e.g., third party apps running in the background

Sharing Example

Related Problems

- Design of usable privacy tools for mitigating sharing concerns
- Warnings and nudges to proactively inform users of risky behavior
- Implement consent and opt-out mechanisms to customize privacy settings

Insider Threat

- Even intended audience may cause propagation of content in unintended ways
- Challenges:
  - Users more concerned with strangers
  - Privacy controls inadequate
  - Potential mitigation strategies: Self-censorship, removal of sensitive content

Studies

- Look at two studies that investigate
  - Users' sharing behaviors and disclosure trends
  - Privacy concerns
  - Mitigation strategies

Facebook and Privacy: It's Complicated

Facebook and Privacy: It's Complicated

Maritza Johnson
Columbia University
maritzaj@cs.columbia.edu

Serge Egelman
UC Berkeley
egelman@cs.berkeley.edu

Steven M. Bellovin
Columbia University
smb@cs.columbia.edu

ABSTRACT

We measure users' attitudes toward interpersonal privacy concerns on Facebook and measure users' strategies for reconciling their concerns with their desire to share content online. To do this, we recruited 260 Facebook users to install a Facebook application that surveyed their privacy concerns, their friend network compositions, the sensitivity of posted content, and their privacy-preserving strategies. By asking participants targeted questions about people randomly selected from their friend network and posts shared on their profiles, we were able to quantify the extent to which users trust their "friends" and the likelihood that their content was being viewed by unintended audiences. We found that while strangers are the most concerning audience, almost 95% of our participants had taken steps to mitigate those concerns. At the same time, we observed that 16.5% of participants had at least one post that they were uncomfortable sharing with a specific friend—someone who likely already had the ability to view it—and that 37% raised more general concerns with sharing their content with friends. We conclude that the current privacy controls allow users to effectively manage the outsider threat, but that they are unsuitable for mitigating concerns over the insider threat—members of the friend network who dynamically become inappropriate audiences based on the context of a post.

1. INTRODUCTION

People spend an unprecedented amount of time interacting with social network sites (SNS) and uploading large quantities of personal information [22, 19]. The dramatic growth in SNS use has created a myriad of privacy concerns. In this paper, we focus on the interpersonal privacy concerns that arise between SNS users and how they manage their concerns by expressing preferences for who should be allowed to access posted content.

Access control management is known to be a difficult problem for end-users in other domains [25, 24]. Not surprisingly, the task of correctly configuring privacy controls, a particular type of access control management task, is out of reach for many SNS users [21]. Knowing that SNS privacy settings are difficult to manage correctly, our research furthers the goal of designing a more usable mechanism, beginning with the question, *How likely are Facebook users to share content with unintended audiences, and what mitigation strategies do they use?*

We constructed an interactive Facebook application to survey 260 Facebook users about specific pieces of content that they had posted to their profiles, as well as their levels of comfort sharing content with randomly selected people from their friend networks. We observed that many participants (94.6% of 260) deny access to their profile content—posts or photos—to people outside their friend network (e.g.,

Johnson et al. Facebook and Privacy: It's Complicated. Symposium on Usable Privacy and Security (SOUPS), pages 9:1–9:15, 2012

Research Questions

- How likely are OSN users to share content with unintended audiences?
- Are users aware of the privacy risks?
- What strategies do user choose to mitigate privacy risks?

Exercise: Facebook Usage

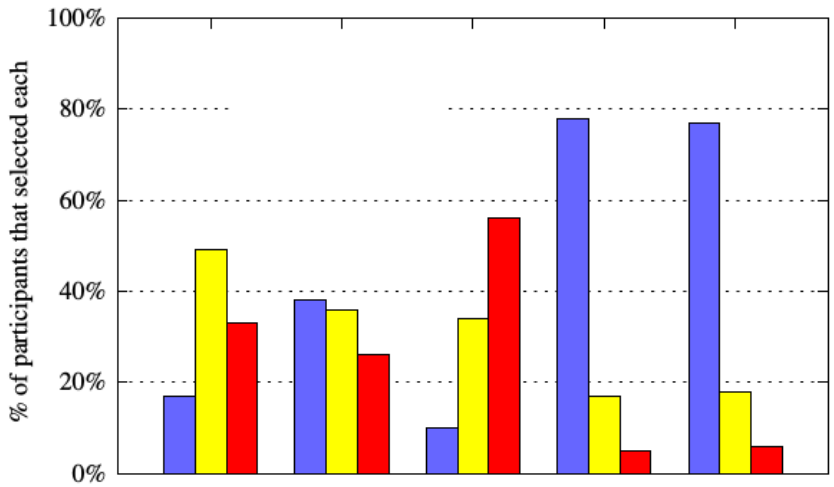

- X axis: Browse, Find friends, Add photos, Look up, View Photos
- Colors: Always, Never, Sometimes

Facebook Usage

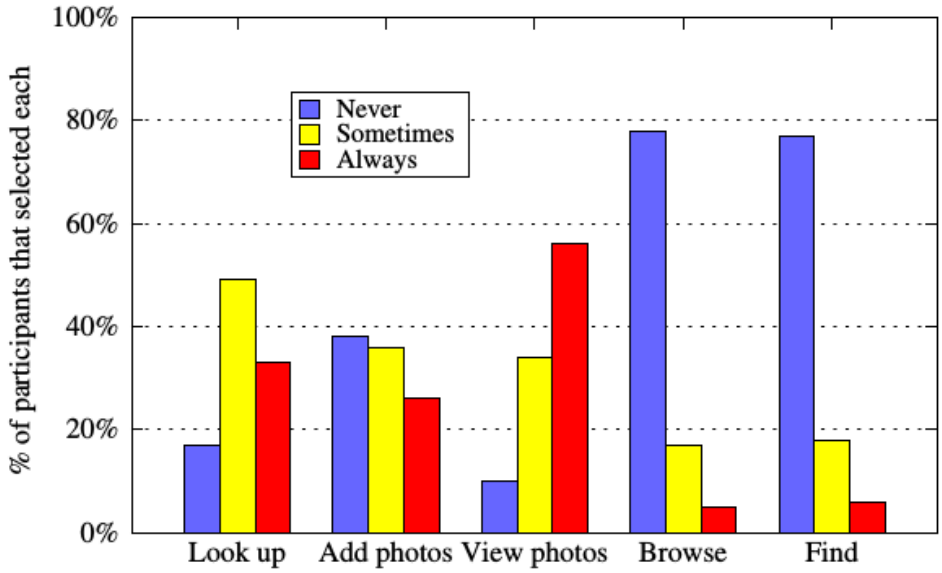

Top Privacy Concerns

- Organizational threats: Collection and use of data by OSN provider
- Lack of control over actions of other users
- Boss or acquaintance might see something embarrassing

Mitigation Strategies

- Have friends only profiles, custom friend lists (subsets of friend network)
  - Not for privacy purposes though
  - Other features, e.g., group friends who play the same game
- Curating friend network: Deny friend request, unfriend
- Delete posts, untag themselves
- Goes beyond official privacy controls provided by the OSNs
  - Multiple accounts: Maintain separate profiles, separate OSNs for different purposes
  - Ask friends to remove photos

Methodology

- Qualitative survey
- Deploy as Facebook app
- Enables the use of real profile data
- Three sections:
  - General questions about Facebook usage (basis for perception vs actual risks)
  - Common scenarios with unwanted audiences (based on previous work)
  - Questions specific to user's Facebook friends and posts

Using Facebook API

- 9 randomly selected friends
- Understand whether users know who their friends are, and how much they trust them
  - What is your relationship to FRIEND-NAME?
  - How do you feel about FRIEND-NAME viewing your posts?
- Include a fictitious friend to check if user is diligent

Results: Stranger Danger

- 14% have public wall
- 7% have public photos
- 54% have public friends list
- 45% have no information accessible to strangers
- Participants with private posts had less concerns (strangers cannot access anyway)

Exercise: Likeliness of Concerns Becoming Reality

Scenario
A stranger will see an inappropriate photo or comment on your profile.
Political parties using Facebook to target you through the use of ads and data mining.
Your employer using your profile to assess your suitability for the company.
Sexual predators using Facebook to track, monitor, locate, and identify you as a potential victim.
Your university using Facebook to identify you as a university code violator.
Law enforcement using Facebook to track drug use and other illegal activities.
Thieves using Facebook to track, monitor, locate, and identify you as a potential victim.
Your employer seeing an inappropriate photo or comment on your profile.
Your employer using Facebook to monitor your conduct while you're at work.
Your employer using Facebook to monitor your conduct while you're away from work.

Likeliness of Concerns Becoming Reality

Scenario	Concerned
1. Thieves using Facebook to track, monitor, locate, and identify you as a potential victim.	68.8%
2. Your employer seeing an inappropriate photo or comment on your profile.	62.7%
3. Your employer using your profile to assess your suitability for the company.	55.0%
4. Sexual predators using Facebook to track, monitor, locate, and identify you as a potential victim.	51.9%
5. Your employer using Facebook to monitor your conduct while you're at work.	46.2%
6. Your employer using Facebook to monitor your conduct while you're away from work.	44.6%
7. A stranger will see an inappropriate photo or comment on your profile.	40.8%
8. Political parties using Facebook to target you through the use of ads and data mining.	30.4%
9. Your university using Facebook to identify you as a university code violator.	20.0%
10. Law enforcement using Facebook to track drug use and other illegal activities.	17.3%

Results: Insider Threat

- Average number of friends: 357
  - Facebook average much lower
  - Compatible with other studies
- Trust 75% of their friends
- Less trusting ($<$ 50%) users do not seem to modify privacy settings

Friend Categories

	Frequency	Participants	Comfort
School	42.6%	88%	97.0%
Socialize with	15.4%	57.3%	98.9%
Friend of a friend	12.4%	62.7%	97.0%
Coworker	11.1%	45%	96.9%
Extended family	9.4%	48.5%	95.4%
Have not met	5.3%	20%	95.2%
Immediate family	2.1%	14.2%	98.0%
Not sure	1.7%	13%	75.0%

Sharing Comfort

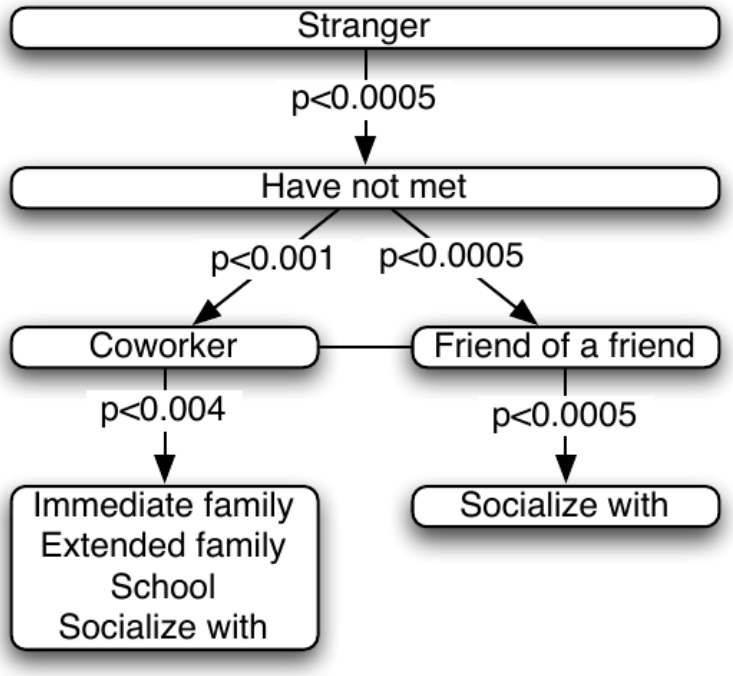

Implications

- Stranger threat mostly mitigated
- Insider threat remains
- Custom lists not utilized for managing privacy concerns

Limitations

- Hard to estimate why people choose not to participate
- Facebook app bias
- Dynamic threats: Appropriateness of the audience is highly contextual

Silent Listeners: The Evolution of Privacy and Disclosure on Facebook

Silent Listeners:
The Evolution of Privacy and
Disclosure on Facebook

Fred Stutzman*, Ralph Gross†, Alessandro Acquisti‡

Abstract. Over the past decade, social network sites have experienced dramatic growth in popularity, reaching most demographics and providing new opportunities for interaction and socialization. Through this growth, users have been challenged to manage novel privacy concerns and balance nuanced trade-offs between disclosing and withholding personal information. To date, however, no study has documented how privacy and disclosure evolved on social network sites over an extended period of time. In this manuscript we use profile data from a longitudinal panel of 5,076 Facebook users to understand how their privacy and disclosure behavior changed between 2005—the early days of the network—and 2011. Our analysis highlights three contrasting trends. First, over time Facebook users in our dataset exhibited increasingly privacy-seeking behavior, progressively decreasing the amount of personal data shared publicly with unconnected profiles in the same network. However, and second, changes implemented by Facebook near the end of the period of time under our observation arrested or in some cases inverted that trend. Third, the amount and scope of personal information that Facebook users revealed privately to other connected profiles actually increased over time—and because of that, so did disclosures to "silent listeners" on the network: Facebook itself, third-party apps, and (indirectly) advertisers. These findings highlight the tension between privacy choices as expressions of individual subjective preferences, and the role of the environment in shaping those choices.

Stutzman et al. Silent Listeners: The Evolution of Privacy and Disclosure on Facebook. Journal of Privacy and Confidentiality, 4(2):7–41, 2012

Research Questions

- How do sharing behaviors of users change over time?
- How do disclosures to “silent listeners” evolve?

Study Overview

- Understand privacy and disclosure behavior
- Longitudinal study: Over a long period of time (2005–2011)
- Dataset: 5,076 members of CMU Facebook network
- Early joiners of Facebook

Common Findings with Other Studies

- Progressively limit content to strangers
- Consistent among all profile elements
- Intended audiences not necessarily map to actual audiences
- Mitigation strategies: Self-censorship, withdrawal of content

CMU Yearly Snapshot Dataset

Time	Year	Collection Date	Total Obs.	Panel Obs.
t_0	2005	Nov 20	6380	5076
t_1	2006	Nov 29	10254	5076
t_2	2007	Nov 02	15041	5076
t_3	2008	Mar 01	15324	5076
t_4	2009	Oct 04	15024	5076
t_5	2010	Nov 12	15731	5076
t_6	2011	May 5	22124	5076

Facebook Profile Elements

Element	Disclosure Category	Note
Birthdate	Personal	FT (mm-dd-yyyy)
High School	Personal	FT
Hometown	Personal	FT
Political Affiliation	Personal	Initially DD, later FT
Instant Messenger	Contact	FT, any IM (AIM, Skype, Y!, etc)
Phone	Contact	FT, any phone (mobile or landline)
Address	Contact	FT
Looking For	Contact	DD
Interests	Interests	Initially FT, later L
Favorite Music	Interests	Initially FT, later L
Favorite Books	Interests	Initially FT, later L
Favorite Movies	Interests	Initially FT, later L

- FT: Free text input, DD: Drop down list, L: Like button

Sharing Trends

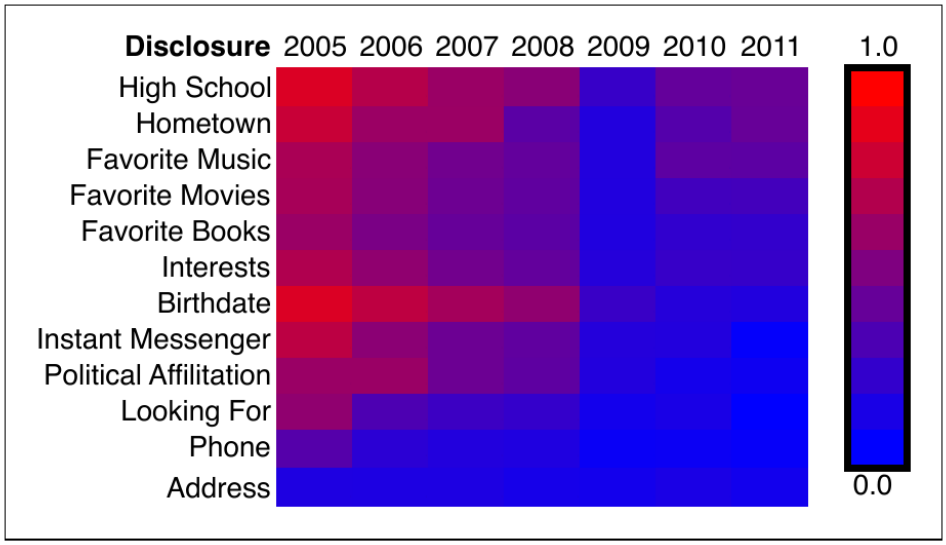

- SSN can be effectively predicted from hometown and birthdate

Disclosure: Personal Information

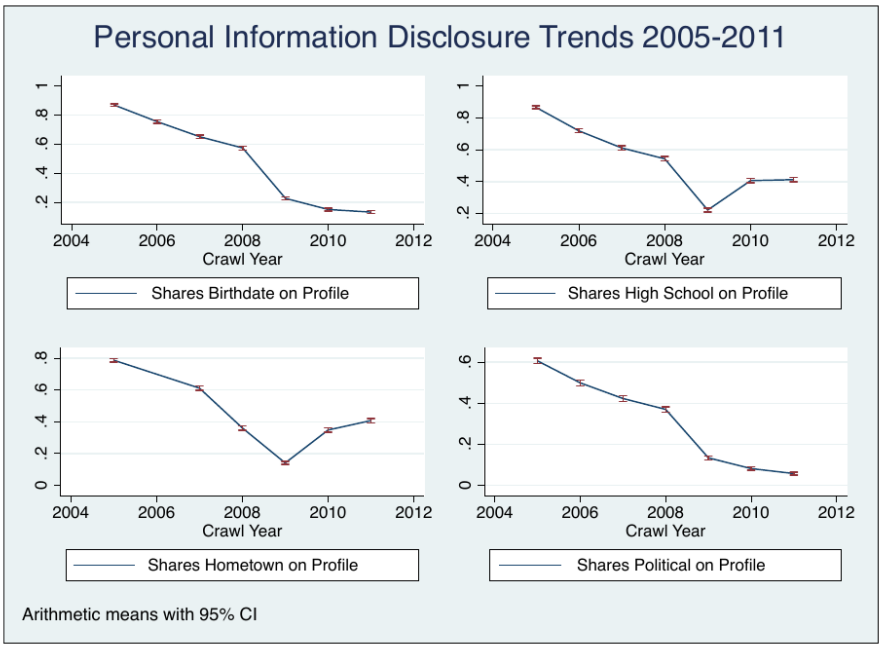

Disclosure: Interests

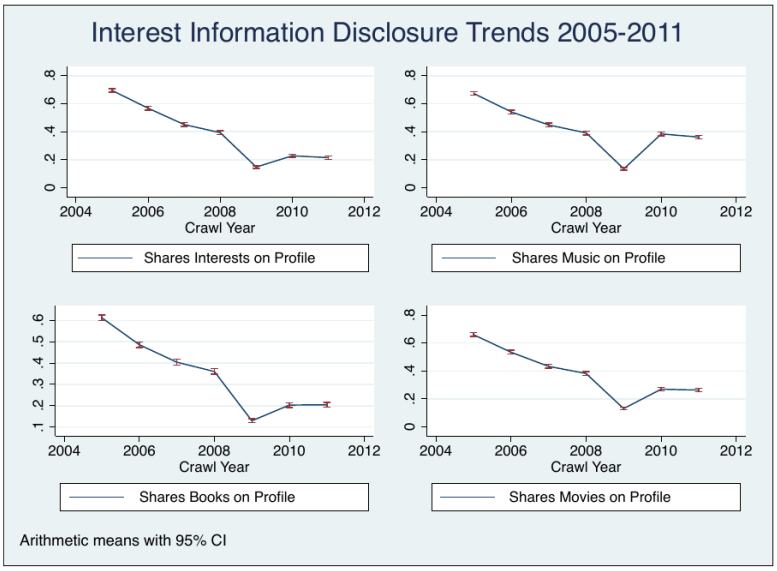

- What might have happened in 2009–2010?

Disclosure: Contact Information

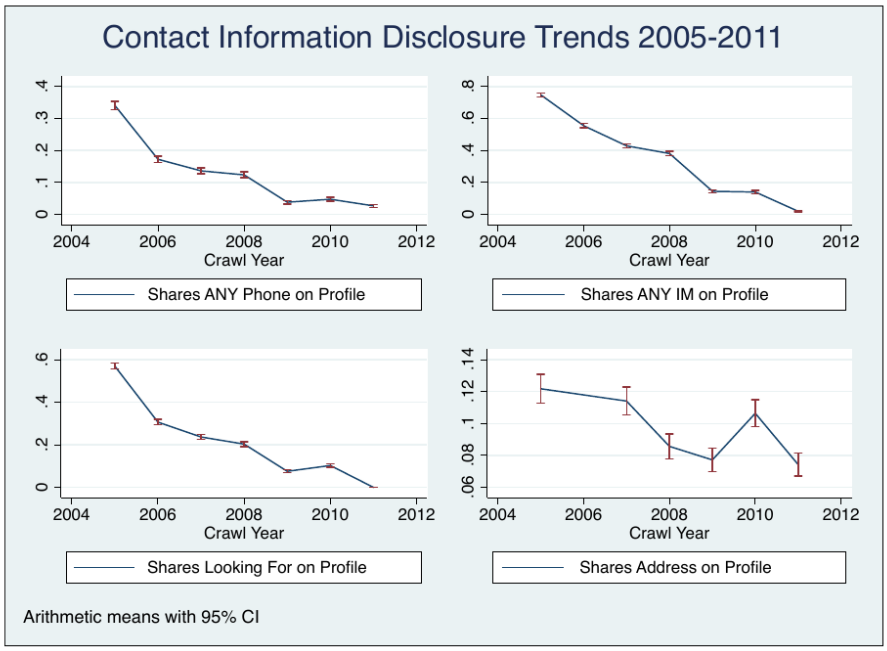

Trend Reversal

- Despite the continuous decrease, participants started sharing more after 2009
- Policy and design changes by Facebook
- Privacy Wizard: Share content with audiences of their choice
- Trend reversal apply to some profile elements
- Community pages let users connect with others who share similar interests

Community Pages

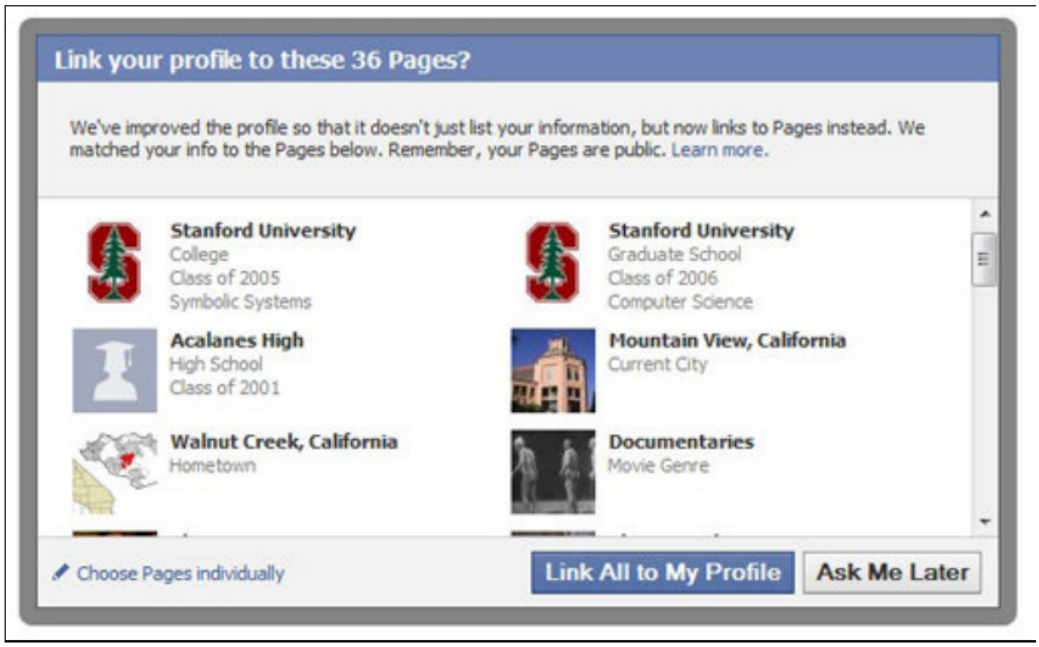

Silent Listeners

- While public disclosures decreased, private sharing of content increased
- This, in turn, increases disclosures to “silent listeners”
  - Facebook itself
  - Third party apps
  - Advertisers
- Disclosure without awareness or explicit consent
- Users underestimate their audience: They can only guess 27% of their true audience

Facebook Apps

AppName	Type	Data
ChefVille	Game	Food preferences
TripAdvisor	Travel	Trip recs./history/checkins
Yahoo! Social Bar	News/Social Reader	Yahoo activity reported
Instagram	Photo	Photo sharing/check-ins (FB)
Microsoft Live	Utilities/Communication	Social/search engine
Bing	Utilities	Social/search engine
Spotify	Music/Entertainment	Music choices
Scribd	Utilities	Interests for reading/publishing
SchoolFeed	Online Communications	Connects users to others
Between You and Me	Dating	Dating
MyPad for iPad	FB for iPad	Recreates FB for an iPad
Skype	Utilities	Communications/networking
FourSquare	Utilities	Aggregates Check-ins

Limitations

- Not a random sample of Facebook users
- Based on profile elements available in 2005

Girls Around Me

- News article: http://www.cultofmac.com/157641/this-creepy-app-isnt-just-stalking-women-without-their-knowledge-its-a-wake-up-call-about-facebook-privacy/
- Links are also on the course website

Things to Look For

- Root cause: What went wrong?
- If it was not intentional, what was the original aim?
- Affected parties
- Implications and similar problems
- Mitigation (using methods we have seen): Prevention, detection, recovery

- Take 10 minutes to look at the incident on your own

- Now discuss with your neighbor
- Also take a look at the summary report: https://drive.google.com/a/ncsu.edu/file/d/0B3m-I0YVAv0EX19nWkRvaGdEVTg/view