User Privacy in Platforms

Itay P. Fainmesser Andrea Galeotti Ruslan Momot

1 Introduction

Platforms like Facebook, Google, and WeChat, store, catalog, and make available for use extensive amount of data about their users. This data helps the platforms deliver tailored advertisements as well as to customize and personalize their offerings of products and services, based on users' personal and social characteristics as inferred from their activity on the platform.

Users value the improved experience. However, as most recent reports show, personal data is also shared with and sold to third parties. Recent revelations of Facebook's data sharing practices showed that the Tech Giant was involved in data sharing with Spotify, Netflix, Amazon and others, which extended well beyond the disclosed limits or those prescribed by the privacy agreements between the platform and its users.[1] As such, irrespective of whether users disabled data sharing or not, partners of the social network had access to the content of users' personal messages (in some cases also with the right to write/delete messages), geographical location, calendar entries, contacts. This business model under which Facebook operates, is, in fact, employed by the majority of contemporary platforms. The core of their business strategy is to provide a product to users for free, while relying on heavy utilization and trade of user data for generating profits.

The prevalent use of personal data exposes users to unwelcome risks. For instance, third parties, such as Cambridge Analytica, may have goals that contrast with users' preferences and may attempt to manipulate users' behaviors using their personal data.[2] Overall, the spread of personal data increases users' vulnerability to threats and crimes ranging from financial fraud and identity theft to the misuse of their and their close ones' information.

In this paper, we build an equilibrium model of platform privacy. Users experience direct utility from using the platform, however they are also susceptible to strategic, self-interested, criminals who use the information users put on the platform. The platform, in turn, chooses the amount of data it stores and shares with business partners, and its security level. We characterize the equilibrium activity levels of users and criminals, and the resulting privacy levels, as well as how these are affected by different attributes of the platforms, such as the strength of network effects. We also investigate the extent to which a platform is able to protect users while still pursuing its own goals.

2 Model

Consider a platform which hosts a unit mass of users. Users experience direct utility from using the platform. They choose their level of activity, i.e., the amount of time they devote to different activities on a platform (e.g. sharing photos, messaging). Users benefit more from participation if other users are also more active.

Activity on the platform has a direct cost (time invested that cannot be used for activities outside the platform) and, more importantly, leads to higher exposure to risks due to personal information being revealed. For instance, users who post photos on a platform, might reveal their geographical location. Similarly, posting comments/likes or uploading original content, might reveal user's preferences or wealth. The extent to which activity puts users at risk of having their personal information exploited against them, depends on how usable

[1] "As Facebook Raised a Privacy Wall, It Carved an Opening for Tech Giants", *The New York Times*, Dec 18, 2018;
[2] "Facebook and Cambridge Analytica: What You Need to Know as Fallout Widens", *The New York Times*, 19 Mar 2018.

the data is for the purpose of individual user targeting. If the data is organized and saved in a way that allows the platform's business partners to target users accurately, it also gives that capability to a malicious third party which gets access to the data (e.g., by hacking the platform or purchasing the data with false pretenses). Alternatively, if the data is well obfuscated or garbled by the platform, this might lead to a decrease in criminal activity.[3]

Formalizing the aforementioned effects, we propose the following users' utility function:

$$U_i(a_i) = a_i b_i - \frac{1}{2}a_i^2 + \beta a_i \bar{a} - \Pr[\text{user is hacked}] \cdot a_i \xi. \tag{1}$$

Here a_i is user i's activity level, $\bar{a} = \int_j a_j$ is the average activity level on a platform, and b_i and β are parameters, the former capturing the standard benefits from activity on the platform and the latter the strength of network effects. The last term in the utility function describes the user's expected disutility from being hacked. If there is a successful attack, the user suffers a loss of $a_i\xi$ – an amount proportional to her activity level and to $\xi \in [0,1]$, which is, in turn a garbling factor set by the platform to restrict the targeting of individual consumers. For example, $\xi = 0$ corresponds to the case where the platform fully garbles users' information and any information obtained by a criminal has no use for the purpose of targeting individual users.

The probability that a user's data gets hacked is tightly connected to criminal activity levels. For simplicity, we assume that one criminal picks a mass $m < 1$ of customers to attack uniformly at random and pays a fixed cost, $\gamma \sim U[0, C]$, to set up a crime and a variable cost, $c(m)$, that is increasing and convex in the mass of users attacked. The criminal maximizes the total expected damage inflicted on users (e.g., total financial benefit it extorts from users) minus her costs, or $\bar{a}\xi - \gamma - c(m)$.

The platform attempts to increase two derivatives of users' activity levels: (1) raw users' average activity level $\bar{a}$, which the platform can leverage, for example, to display non-targeted ads, or to signal high platform valuation to potential investors, and (2) the usable information the platform accumulates on users, or $\bar{a}\xi$, which allows the platform to sell targeted advertising, as well as to directly sell the data to (unmodeled) business partners. On the other hand, the platform seeks to reduce its costs of data security, $\psi(C)$, which is increasing with the criminal's (expected) fixed cost for hacking the platform. Formally, the platform solves the following problem:

$$\max_{C,\xi} \Pi = \alpha\bar{a} - (1-\alpha)\,\bar{a}\xi - \psi(C)$$

3 Major Results

We characterize the equilibrium activity levels a_i^* of users, the equilibrium strategy of the criminal, m^*, and the platform's optimal design (security and garbling levels C^* and ξ^*). The platform can deter criminals by increasing security and/or increasing garbling levels. These two instruments may have different effects on users' privacy. Increased security leads to an increase in users' activity and in the information that can be used for targeting individual users. Surprisingly, increased garbling (i.e., lower ξ), increases users' activity but may *increase or decrease* the information that can be used for targeting individual users. We show that the optimal design of the platform depends crucially on its objective, its product quality, and the strength of network effects, and can vary over the life cycle of the firm.

[3] One of the common techniques for data garbling is the concept of differential privacy.