Information Sharing and Privacy in Networks

Working draft, comments are welcome

Ronen Gradwohl*

Abstract

Users of social, economic, or medical networks share personal information in exchange for tangible benefits, but may be harmed by leakage and misuse of the shared information. I analyze the effect of enhancing privacy in the presence of two opposing forces: network effects and informational interdependencies. I show that two privacy enhancements—reducing the likelihood of leakage and decreasing the level of informational interdependence—have opposite effects on the volume of information sharing, and that although they always seem beneficial to non-strategic users, both privacy enhancements may backfire when users are strategic.

1 Introduction

Consumers' personal information is a valuable commodity in the booming information economy, as such data is a major driver of growth and innovation. And in a variety of contexts, ranging from online social networks to healthcare information technology, consumers willingly share their personal information in exchange for perceived benefits. However, firms' increasing use of consumers' personal information has been accompanied by the latter's growing concerns for their privacy.[1] A key question for

*Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. E-mail: r-gradwohl@kellogg.northwestern.edu.

[1] In a recent Pew survey, for example, over 90% of participants expressed such a concern (Pew Research Center, 2014).

innovation policy is thus how to balance the tradeoffs between information sharing and privacy concerns (Goldfarb and Tucker, 2012). In particular, what are the effects of various privacy policies on the level of information sharing and on consumer welfare?

One context in which the tradeoff is particularly salient is in online social networks, whose popularity has skyrocketed over the past decade. For example, Facebook's 1.5 billion users[2], who share intimate details of their lives on the platform, are increasingly demanding greater privacy for their personal information (Johnson et al., 2012; Pew Research Center, 2014). Facebook, as well as other social networking platforms, photo sharing sites, and search engines, have enacted policy modifications in response to these concerns, modifications that have affected the level of activity on the sites.

There are other contexts in which a similar tradeoff between information sharing and privacy arises. Examples include healthcare information technology, when health and genomic information is digitized and shared; and consumer loyalty programs, when shoppers share shopping habits in exchange for discounts (Rainie and Duggan, 2016).

In this paper I analyze the effects of privacy enhancements on the tradeoffs faced by privacy-concerned individuals. The main insight is that privacy enhancements directly affect welfare by lowering costs, but also indirectly affect welfare by altering users' levels of sharing. I show that different privacy enhancements may have opposite effects on the volume of information sharing, and that although they always seem beneficial to non-strategic users, privacy enhancements may backfire when users are strategic.

The observation that privacy regulation may be harmful is not new, and the burgeoning empirical and experimental literature on the topic has shown that the effects of regulation may be positive or negative, depending on the context (see the excellent survey of Acquisti et al., 2016). The theoretical literature on privacy has its roots in the work of Posner (1981) and Stigler (1980), who derive a similar conclusion in a signaling context: under stronger privacy regimes individuals can more readily hide negative traits, which may be harmful to other market participants and to social welfare. More recent theoretical work, all of which builds on the signaling model, has

[2] As of 2015. See Facebook (2016).

pointed out that the benefits or harms of privacy enhancements are highly context-dependent (Acquisti et al., 2016).

This paper's goal is to identify properties of interactions that determine the effects of various privacy regulations. The point of departure is the observation that the conception of privacy commonly studied in the theory literature—namely, as a technology for altering the signaling capabilities of individuals—misses a key dimension of privacy harm: a concern not about third parties' inferences about individuals' types, but rather about misuse of the information itself (Milberg et al., 1995; Calo, 2011). Individuals concerned about identity theft, spam, harassment, stalking, re-identification, online tracking, excessive profiling, and targeted advertising care less about whether the information leaked about them is positive or negative, and more about the fact that information has been leaked and misused in the first place.[3] This dimension of privacy concerns, highlighted in opinion surveys (Turow et al., 2009; Pew Research Center, 2014; Rainie and Duggan, 2016) and in general analyses of privacy risks associated with Facebook and other online social networks (Gross and Acquisti, 2005; Consumer Reports, 2012; Wilson et al., 2012; Liang et al., 2015), is the focus of the paper.

We frame the paper within the context of online social networks, but describe its applicability in other contexts as well. In the model, agents are not concerned about what the leaked information signals about their type, but rather about the quantity of personal information leaked. For a concrete example, consider a Facebook user posting photographs on the platform. The user enjoys this online activity, but also faces some risks: the more photos are posted, the greater the chance that they are misused by stalkers, thieves, lawyers, or employers (O'Donnell, 2016). In particular, the privacy concern is not as much about the personal traits signaled by the photos—the user can choose not to post embarrassing photos—but rather about the user's daily habits, whereabouts, appearance, and associations, which can all be deduced from the posted trove of photographs.

[3]Profiling and targeted advertising may also be beneficial to users (just as price discrimination is not necessarily harmful). However, this dimension of the associated privacy concern is less about *which* advertisements are targeted at an individual, and more about the fact that advertisers track the individual across websites (Gross and Acquisti, 2005; Turow et al., 2009; Rainie and Duggan, 2016).

As a preliminary formalization, consider a user of a social network who wishes to share some information. The user derives a benefit from sharing information on the platform, a benefit captured by an increasing function of the amount of information shared. However, there is some chance λ of information leakage and misuse—the user's account is hacked, the information he shared is seen by another user who was not the intended recipient, or the information reaches a malicious third party—in which case the user incurs a cost, captured by another increasing function of the amount of information shared. The user thus faces a tradeoff between benefit and cost. Is a privacy enhancement, in the form of lowering λ, beneficial to the user? The answer is easily seen to be positive. Interestingly, if the network owner desires to maximize the amount of information shared on the platform, then that same privacy enhancement is, in this respect, also beneficial to him.[4]

Of course, this simple example ignores the collective nature of social networks. More realistically, suppose there are many users of the network, each deciding on the amount of information to share. The interaction between users is captured by two opposing forces: First, users enjoy network effects from participation in the platform: the more information others share, the more benefit an individual user derives from his own participation. Users of Facebook, for example, benefit from the consumption of information shared with them by others (Grimmelmann, 2009; Wilson et al., 2012).

Second, a key feature of personal information on social networks is its interdependent nature: Individuals hold information not only about themselves, but also about others (boyd, 2012). If others share information about a particular user, that user may suffer from the leakage of information that he did not himself share. A simple example is photo sharing on Facebook: One can post a photo and tag another user in it, thereby sharing information about the other's facial features, dress, location, and the fact that he has a social tie with the user who posted the photo.[5]

[4] The first statement is a corollary of Claim 1, and the second of Theorem 2.

[5] Sarigol et al. (2014) study the amount of information posted about users by others, and Olteanu et al. (2014) quantify co-location data (data about joint location of a user and someone else) present, for instance, in photos. Other examples of interdependencies include information automatically shared by users about their social connections when installing third-party apps on Facebook (Biczók and Chia, 2013; Pu and Grossklags, 2014, 2015; Symeonidis et al., 2016), as well as interdependencies in genetic/genomic data (boyd, 2012; Humbert et al., 2015): sharing one's genome also reveals information about family members, including currently unborn ones!

Given these forces shaping interaction, I study two distinct privacy enhancements for online social networks: lowering the probability λ of leakage and misuse of information, and reducing the level of interdependence in users' information. On Facebook, for example, the former may correspond to a reduction in the number of other users who can view an individual's posts, limiting the ability of others to re-post one's information, or increasing the security of users' accounts (for example, using two-factor authentication). The latter privacy enhancement might be achieved by limiting the ease with which users tag others in their photos (for example, by removing the use of facial recognition technology on users' photos), or restricting the amount of information users can share about others with third-party apps.

A very different context for the model is healthcare information technology. The digitization of medical records and, more recently, their incorporation in health information exchanges, holds many potential benefits: As more patient information is digitized and shared, patients get easier access to their data and healthcare providers can more readily communicate, leading to fewer duplicate tests, lower costs, and overall more efficient care. On the other hand, the more information is shared, the greater the concern for privacy, as leaked and misused records lead to fraud and discrimination[6] (Schwartz, 1997). Possible privacy enhancements include lowering the probability of leakage by reducing access and limiting re-disclosure of electronic records (Adjerid et al., 2015), and lowering the level of interdependence by more strictly regulating access to genetic/genomic information (McGuire et al., 2008).

The main results of this paper are that although both enhancements—lowering the probability of leakage and lowering the level of interdependence—have a positive direct effect on user welfare, they have opposite effects on the volume of information sharing. The impact on the volume of sharing drives an indirect effect on user welfare, and its direction depends on the tradeoff between network effects and interdependencies. If the indirect effect is positive—when the privacy enhancement increases sharing and network effects dominate interdependencies, or when the enhancement decreases sharing and interdependencies dominate network effects—then the net ef-

[6]While the concern for discrimination is a "signaling" concern, it can also be seen as a quantitative one, as patients withhold information and may avoid medical tests (prior to knowing the results) out of the fear that these will appear on their medical records and possibly leaked (Congress, 1993).

fect of the privacy enhancement on user welfare is positive. However, if the indirect effect is negative, then the net effect depends on the structure and intensity of the interdependencies.

To better understand the change in welfare when the indirect is negative, I examine two extreme forms of interdependencies: overlapping information and additive information. Under overlapping information, the information shared by one user about another overlaps what the other shares about himself, and so contributes additional information only if he shares more of it. For example, on Facebook, profile information shared about others with third-party apps is overlapping information. Under additive information, in contrast, the information shared by one user about another is distinct from what the other has shared about himself. For example, original photos shared by two users of a social network are distinct, and so the facial-feature information they contain about a third user is additive. I show that enhancing privacy by lowering the leakage probability is beneficial when information is overlapping, but not necessarily when it is additive. In contrast, enhancing privacy by decreasing the level of interdependence is harmful if information is overlapping, but beneficial when it is additive.

The general intuition that drives these results is that under overlapping information, the interaction between users is locally similar to an interaction with no informational interdependencies. Hence, the response to small changes in λ is driven by its direct effects. Furthermore, small changes to the interdependencies have no direct effect, and so the response to those is driven by its indirect effects. Under additive information, in contrast, the indirect effect of varying the level of interdependencies is neutralized, and so only the direct effect influences utility.

Finally, I show that user-optional privacy enhancements lead to privacy "paradoxes", in which an enhancement is welfare-improving but each user optimally refrains from implementing it.

Related Literature There is by now a large and growing literature on the economics of privacy, both theoretical and empirical (see the survey of Acquisti et al., 2016, and the many references therein). The current paper is most closely related to a subset of this literature that focuses on the behavior of individuals whose privacy concerns are intrinsic—they would either like to signal a particular trait (e.g., Daughety and Reinganum, 2010), or minimize the amount of signaling altogether (e.g., Gradwohl and Smorodinsky, 2017). However, unlike this paper, nearly all of the work in the literature focuses on the signaling aspect of privacy concerns.

One thread that runs through the economics of privacy literature is that privacy enhancements can be beneficial or detrimental, depending on the setting (Acquisti et al., 2016). In particular, several papers show that enhancing privacy may backfire due to strategic behavior (see, for example, Cummings et al., 2016; Gradwohl and Smorodinsky, 2017; Gradwohl, forthcoming). In this paper we show that, even in the same setting, different privacy enhancements can have very different effects. Furthermore, we try to point out features of the setting that determine the direction of the effect.

A separate quantitative measure of the amount of information leaked is captured by the notion of differential privacy (Dwork, 2008). This quantitative measure is also related to the signaling aspect of privacy: when differential privacy is guaranteed at a certain level, it implies an upper bound on the accuracy of any Bayesian deduction made from the leaked information (Kasiviswanathan and Smith, 2008).

Another related strand of the literature consists of work on the economics of security (Anderson and Moore, 2006). While issues of interdependence also arise there (see, e.g., Kunreuther and Heal, 2003), the emphasis is quite distinct from the current paper, and consists of a focus on firms' incentives for investing in security.

In addition, this paper contributes to the growing literature on Facebook and other online social networks, a literature comprised of research from a variety of fields ranging from economics and law to psychology and sociology (see the survey of Wilson et al., 2012, and the many references therein).

Finally, our modeling of personal information is closely related to the notion of evidence in Dziuda and Gradwohl (2015): it is nonatomic and must be revealed truthfully, but its holders can refrain from revealing some or all of it. Dziuda and Gradwohl also analyze a tradeoff between the amount of information revealed and a privacy concern, but do so in the context of inter-firm communication.

2 Model and Preliminaries

We examine a family of parametrized games denoted by $\Gamma = (I, (S_i, v_i, c_i)_{i \in I})$. The first element, I, is the finite set of players, which we henceforth call *users*. Each user $i \in I$ shares information $x_i \in S_i$, and we will assume for simplicity that $S_i = [0, 1]$.[7] Let $x = (x_1, \ldots, x_{|I|})$ and $S = (S_1, \ldots, S_{|I|})$. The value of the network to user i, which captures the benefits of sharing information with its intended audience, is modeled by the function $v_i : S \mapsto \mathbb{R}_+$, which we assume to be nondecreasing as users enjoy network effects. The cost due to leakage and misuse of information, which captures the harms caused by shared information reaching an unintended recipient, is formalized by the function $c_i : S \times [0, 1) \mapsto \mathbb{R}_+$, which we also assume to be nondecreasing in S—the more information about a user is available, the greater the harm caused by misuse of this information. The second kind of element in domain of c_i is a parameter $\rho \in [0, 1)$, which is an interdependence parameter: the greater ρ, the more of i's information is contained in x_{-i}.[8] The precise interaction between c_i and ρ is explicated below, but for now we simply assume that c_i is increasing in ρ: under greater interdependence, the information x shared by users contains more information about user i, and so the corresponding cost is higher.

We assume v_i and c_i are continuous in x_i for fixed x_{-i} and in x_{-i} for fixed x_i, for every $\rho \in [0, 1)$. Γ is a family of games, each parametrized by a leakage probability $\lambda \in [0, 1]$ and an interdependence parameter $\rho \in [0, 1)$. Taken together, these determine the utility functions of users in the game $\Gamma(\lambda, \rho)$, as follows.

When users share information x, each user i derives value $v_i(x)$. There is a probability λ that this information is leaked and misused, and in this case the user suffers cost $c_i(x, \rho)$. The utility of each user $i \in I$ in the game $\Gamma(\lambda, \rho)$ is thus $u_i(\cdot, \lambda, \rho) : S \mapsto \mathbb{R}$, where

$$u_i(x, \lambda, \rho) = v_i(x) - \lambda \cdot c_i(x, \rho).^9$$

[7] The analysis in this paper can be extended to the case in which S_i is a complete lattice given appropriate supermodularity conditions. As this does not seem to add much insight, we stick with the simpler formulation.

[8] We assume that $\rho < 1$ in order to avoid trivial equilibria.

[9] Note that the leakage probability and interdependencies do not add value to the network, as v depends only on x. In Section 7.3 we extend the model to allow for the possibility that v increases with λ and ρ.

For fixed ρ (respectively, λ), denote by $\Gamma(\rho)$ (respectively, $\Gamma(\lambda)$) the subset of games of Γ in which ρ (respectively, λ) is fixed. We will also write $u_i(x, \lambda)$, $u_i(x, \rho)$, and $u_i(x)$ when ρ, λ, and both, respectively, are clear from context.

Some of our analysis relies on the dependence of c_i on ρ, which we now describe in greater detail. Suppose users share information x. First, this means user i shares x_i about himself. Second, while $j \neq i$ shares x_j about himself, some of x_j contains information about i as well. Specifically, we suppose that a ρ fraction of the information j shared is also about i, namely ρx_j.[10] The information shared by all users about a particular user i is thus captured by the vector $(x_i, (\rho x_j)_{j \neq i})$, which we denote by $x(\rho)$ or by $(x_i, \rho x_{-i})$. Note that this is a vector, rather than a number, since there may be an overlap in different users' information shared about i. The aggregated information about user i is then $\mathcal{A}^i(x(\rho))$ where $\mathcal{A}^i : [0,1]^{|I|} \mapsto \mathbb{R}_+$. Finally, the cost function c_i depends only on the information revealed about user i, and so we write $c_i(x, \rho) = \hat{c}_i(\mathcal{A}^i(\rho(x)))$, where $\hat{c}_i : \mathbb{R}_+ \mapsto \mathbb{R}_+$ is nondecreasing.

Two examples of informational interdependencies, captured by two aggregation functions $\mathcal{A}^i$, play a prominent role in the analysis: overlapping personal information and additive personal information.

Under overlapping personal information, when two users j and k share information then the shared information of user j about user i completely overlaps the shared information of user k about user i (note that we may have $i = j$). Thus, j shares "new" information about i only if he shares more information about i than k does. Formally, for overlapping information we have $\mathcal{A}^i(x(\rho)) = \max\{x_i, \max_{j \neq i}\{\rho x_j\}\}$. On Facebook, for example, data shared with third-party apps is overlapping, as the demographic information shared by a user about i overlaps the information shared by a different user about i.

Under additive personal information, in contrast, when two users $j \neq k$ share information then the shared information of user j about user i is completely distinct from the shared information of user k about user i. Formally, for additive information we have $\mathcal{A}^i(x(\rho)) = x_i + \sum_{j \neq i} \rho x_j$. On Facebook, for example, original photos shared by two users are distinct, as is thus the facial-feature information they contain about

[10] We assume the same level ρ of interdependence between any pair of users for simplicity only. An extension with heterogeneous, possibly non-linear interdependence is discussed in Appendix 7.2.

a third user.

Of course, in general, personal information need not be entirely overlapping or additive, but rather may lie somewhere on the spectrum between them or may depend on the identity of the users. For example, the genomes of two users j and k reveal information about the genome of a relative i, but if j and k are themselves blood-relatives then the information is neither additive nor overlapping. On the other hand, the information about i revealed by the genomes of j and k overlaps the information revealed by the genome of i about himself. Regardless, one assumption that we maintain throughout is that the information shared by distinct users about user i are substitutes. Formally, this means that the aggregation function $\mathcal{A}^i$ satisfies decreasing differences: for any $x_i' > x_i \in S_i$ and $x_{-i}' > x_{-i} \in S_{-i}$ it holds that

$$\mathcal{A}^i(x_i, \rho x_{-i}') - \mathcal{A}^i(x_i, \rho x_{-i}) \geq \mathcal{A}^i(x_i', \rho x_{-i}') - \mathcal{A}^i(x_i', \rho x_{-i}).$$

This is a natural assumption in this context. To see why, suppose first that user i shares x_i and the others share x_{-i}, but then that others increase the amount they share to $x_{-i}' > x_{-i}$. This increase now includes additional information about user i, namely an addition of $\mathcal{A}\left(x_i, \rho x_{-i}'\right) - \mathcal{A}\left(x_i, \rho x_{-i}\right)$. The amount of additional information, of course, depends on the overlap in information between x_i and the information that is present in x_{-i}' that was not present in x_{-i}. Next, suppose i shares $x_i' > x_i$, and once again consider the difference in information shared by others when they increase the information shared from x_{-i} to x_{-i}'. Since $x_i' > x_i$, the overlap in information between x_i' and the information that is present in $\rho x_{-i}'$ that was not present in ρx_{-i} is larger than when i shared x_i. Thus, the new informational content in the difference between $\rho x_{-i}'$ and ρx_{-i} should be smaller when i shares x_i' than when he shares x_i. This is precisely decreasing differences.

We emphasize that this does not imply that the cost associated with different users' shared information satisfies decreasing differences, only that ownership of the information does. For a more formal justification of this assumption see Appendix A.

The goal of this paper is to analyze the effect of increasing privacy, and this has two distinct meanings in this model. The first is to decrease the probability λ of leakage and misuse, and the second is to reduce the interdependencies between users, by lowering the interdependence ρ.

Before proceeding, a few notes about simplifying assumptions are in order. First, we assume that the amount of information user j shares about user i is linear in x_j: specifically, it is ρx_j. In Section 7.2 we argue that neither this linearity nor the fact that it is the same parameter for all users has any bearing on the results. Second, we assume that the probability λ of information leakage and misuse is identical for all users. Again, however, this is not necessary, and all our results are easily seen to hold when this probability is different for each user. We next discuss some more substantive assumptions.

Assumptions and their implications Our analysis rests on the following assumption about users' utilities, made entirely for tractability: that each game $\Gamma(\lambda, \rho)$ is one of strategic complements. Formally, this means that the utility functions in each game satisfy the single-crossing property in $(x_i; x_{-i})$ (see Milgrom and Shannon, 1994). Intuitively, under this property the best-reply of each user i is increasing in others' strategies. Whether it holds in practice is of course an open question, but there is some empirical evidence that is consistent with it (Moon, 2000; Burke et al., 2009; Acquisti et al., 2012).

The following theorem summarizes two useful facts about games of strategic complements.

Theorem 1 (Milgrom and Roberts, 1990) *Fix a partially-ordered set T. For every $\tau \in T$ let $G(\tau) = (I, (S_i, u_i)_{i \in I})$ be a parametrized game of strategic complements, where I is the set of users and $u_i : S \times T \mapsto \mathbb{R}$ is user i's utility function. Then:*

- *$G(\tau)$ has a maximal and a minimal pure Nash equilibrium (together called the extremal equilibria).*
- *If for every $i \in I$ the function u_i has increasing differences in $x_i \in S_i$ and $\tau \in T$ (for every fixed $x_{-i} \in S_{-i}$) then the extremal Nash equilibria are nondecreasing in τ.*

We will use pure Nash equilibrium as the solution concept. Not only do such equilibria always exist in the class of games we study, but they also have additional robustness properties—see Milgrom and Roberts (1990).

In most of the paper we will make an additional assumption: that, for every ρ, the cost c_i satisfies decreasing differences in $(x_i; x_{-i})$. Formally, for every $x'_i > x_i$ and $x'_{-i} > x_{-i}$,

$$c_i(x'_i, x_{-i}, \rho) - c_i(x_i, x_{-i}, \rho) \geq c_i(x'_i, x'_{-i}, \rho) - c_i(x_i, x'_{-i}, \rho).$$

Recall that $c_i(x, \rho) = \hat{c}_i(\mathcal{A}^i(x(\rho)))$. Then our second assumption is satisfied whenever $\hat{c}_i$ is "not too convex".[11] The meaning of "not too convex" depends on the aggregation function $\mathcal{A}^i$. For example, if information is overlapping then the assumption is satisfied for any $\hat{c}_i$. In contrast, if information is additive then the assumption is satisfied only if $\hat{c}_i$ is concave. More generally, the assumption will be satisfied for concave, as well as mildly-convex, $\hat{c}_i$.

We note that our two assumptions, that u_i satisfies single-crossing and that c_i satisfies decreasing differences, go "in the same direction". Given a function c_i that satisfies decreasing differences, a sufficient condition for u_i to satisfy single-crossing is that v_i satisfy increasing differences.

2.1 Direct effects and non-strategic users

Given a game $\Gamma(\lambda, \rho)$ and a level x of information sharing, the *direct effect* of lowering λ (respectively, ρ) is the change in users' utilities, holding ρ (respectively, λ) and x fixed. It is easy to see that the direct effects of both lowering λ and lowering ρ are positive—they lead to an increase in users' utilities.

We next consider the baseline case in which users of the social network are not strategic: they optimize the amount of information they share based on the parameters λ and ρ, viewing the information shared by others as fixed. Equivalently, such users take only the direct effects of privacy enhancements into account, ignoring the *indirect effect*—the welfare difference caused by the change in others' level of information sharing, brought about by the privacy enhancement. For the rest of this section fix a user i and an amount of information x_{-i} shared by other users.

Claim 1 *Fix ρ and λ, and let x_i be any element of* $\arg\max_{z \in S_i} u_i(z, x_{-i}, \lambda, \rho)$. *Then for any $\lambda' \leq \lambda$, $\rho' \leq \rho$, and $x'_i \in \arg\max_{z \in S_i} u_i(z, x_{-i}, \lambda', \rho')$ it holds that*

$$u_i(x'_i, x_{-i}, \lambda', \rho') \geq u_i(x_i, x_{-i}, \lambda, \rho).$$

[11] But see Section 7.4 for discussion of the opposite assumption.

In words, Claim 1 states that when users are non-strategic then they believe that enhancing privacy either by lowering λ or by lowering ρ (or both) unequivocally leads to a welfare increase. Thus, such users would be in favor of both forms of privacy enhancement.

Proof: First observe that $u_i(x_i, x_{-i}, \lambda', \rho') \geq u_i(x_i, x_{-i}, \lambda, \rho)$, since reductions in both λ and ρ lead to lower costs due to information misuse. But x_i' is an optimal response to x_{-i} under λ' and ρ', and so $u_i(x_i', x_{-i}, \lambda', \rho') \geq u_i(x_i, x_{-i}, \lambda', \rho')$. ■

3 Network Owner

In this section we analyze the effects of increasing privacy by decreasing λ or ρ on the amount of information x shared by users in equilibrium. The purpose is twofold: First, it will facilitate the analysis of user welfare in subsequent sections, as the volume of information sharing is the conduit for the indirect effect of privacy enhancements. Claim 1 states that if users are non-strategic then both privacy enhancements are welfare increasing. However, what such users miss is the indirect effect, through the amount of information shared by others. But of course this is relevant, as others' information impacts a user's utility via both network effects and interdependencies.

Second, the volume of information shared is one comparative static with which the network owner may be concerned. In particular, the network owner may wish users to increase the amount of information shared on his network. Of course, he may have additional concerns, but the basic question we seek to answer here is: if enhancing privacy were costless, would the network owner find it desirable? In this section we show that the answer depends on the form of privacy enhancement.

The first theorem is about enhancing privacy by decreasing the probability λ of information leakage.

Theorem 2 *For any ρ, the extremal equilibria of $\Gamma(\rho)$ are decreasing with λ.*

Thus, as privacy protection increases (i.e., λ decreases), the amount of personal information shared by each user increases. If such an enhancement were not too costly, then, the network owner may find it desirable.

Proof: We will show that $u_i(x, \lambda)$ has decreasing differences in x_i and λ, and so increasing differences in x_i and $-\lambda$. The result will then follow from Theorem 1.

Fix $x_{-i} \in S_{-i}$, as well as any $x_i < x_i' \in S_i$ and $\lambda < \lambda'$. Then:

$$
\begin{aligned}
[u_i(x_i', &x_{-i}, \lambda') - u_i(x_i, x_{-i}, \lambda')] - [u_i(x_i', x_{-i}, \lambda) - u_i(x_i, x_{-i}, \lambda)] \\
&= [v_i(x_i', x_{-i}) - \lambda' \cdot c_i(x_i', \rho x_{-i}) - v_i(x_i, x_{-i}) + \lambda' \cdot c_i(x_i, \rho x_{-i})] \\
&\quad - [v_i(x_i', x_{-i}) - \lambda \cdot c_i(x_i', \rho x_{-i}) - v_i(x_i, x_{-i}) + \lambda \cdot c_i(x_i, \rho x_{-i})] \\
&= \lambda'[(c_i(x_i, \rho x_{-i}) - c_i(x_i', \rho x_{-i})] - \lambda[(c_i(x_i, \rho x_{-i}) - c_i(x_i', \rho x_{-i})] \\
&= (\lambda' - \lambda)[(c_i(x_i, \rho x_{-i}) - c_i(x_i', \rho x_{-i})] \\
&\leq 0,
\end{aligned}
$$

where the inequality follows from the fact that $\lambda < \lambda'$ and the assumption that c_i is nondecreasing. ■

Next, we consider the second privacy enhancement: lowering the level of informational interdependence between users.

Theorem 3 *For any $\lambda \in [0, 1]$, the extremal equilibria of $\Gamma(\lambda)$ are increasing in ρ.*

Thus, as this form of privacy protection increases (i.e., ρ decreases), the amount of personal information shared by each user decreases. So even if such an enhancement were costless, the network owner would not find it desirable.

Proof: We will show that $u_i(x, \lambda, \rho)$ has increasing differences in x_i and ρ. The result will then follow from Theorem 1.

Fix $x_{-i} \in S_{-i}$, as well as any $x_i < x_i' \in S_i$ and $\rho < \rho'$. Then:

$$
\begin{aligned}
[u_i(x_i', &x_{-i}, \rho') - u_i(x_i, x_{-i}, \rho')] - [u_i(x_i', x_{-i}, \rho) - u_i(x_i, x_{-i}, \rho)] \\
&= [v_i(x_i', x_{-i}) - \lambda \cdot c_i(x_i', \rho' x_{-i}) - v_i(x_i, x_{-i}) + \lambda \cdot c_i(x_i, \rho' x_{-i})] \\
&\quad - [v_i(x_i', x_{-i}) - \lambda \cdot c_i(x_i', \rho x_{-i}) - v_i(x_i, x_{-i}) + \lambda \cdot c_i(x_i, \rho x_{-i})] \\
&= \lambda[c_i(x_i, \rho' x_{-i}) - c_i(x_i', \rho' x_{-i})] - \lambda[c_i(x_i, \rho x_{-i}) - c_i(x_i', \rho x_{-i})] \\
&\geq 0,
\end{aligned}
$$

where the inequality follows from the fact that $\lambda \geq 0$ and the assumption that c_i has decreasing differences in x_i and x_{-i}. ■

We noted above that the network owner may be concerned with the amount of information x shared by users, and asked: if privacy enhancements were costless, would the network owner desire their implementation? Theorems 2 and 3 show that the answer is yes for lowering λ, but no for lowering ρ.

Now, what if the network owner wishes to maximize the total amount of information shared about each user, namely $\mathcal{A}(x, \rho) = (\mathcal{A}^1(x(\rho)), \ldots, \mathcal{A}^{|I|}(x(\rho)))$? Clearly, he would still wish to lower λ, as this leads to higher x and so higher $\mathcal{A}(x, \rho)$. The case of lower ρ is a bit subtler, however, since lower ρ leads to lower x but also affects $\mathcal{A}(x, \rho)$. But note that lowering ρ leads to lower $\mathcal{A}(x, \rho)$ (holding x fixed), and so the two effects go in the same direction. Thus, a network owner who wishes to maximize $\mathcal{A}(x, \rho)$ would not want to lower ρ.

In this section we analyze the effects of the privacy enhancements on user welfare. There are two cases, differentiated by whether indirect effects are positive or negative. In this section we consider the former.

4 User welfare with positive indirect effects

The simpler case is the one in which indirect effects are positive—in which lowering λ or ρ leads to a beneficial change in the equilibrium level of information sharing—and this, together with positive direct effects, leads to a net increase in user welfare. Whether or not indirect effects are positive depends on the relationship between network effects and informational interdependencies. We consider two opposite conditions. In the first, network effects dominate the informational interdependencies:

Condition 1 *$u_i(x_i, x_{-i}, \lambda, \rho)$ is nondecreasing in x_{-i} for every $i \in I$ and $x_i \in S_i$.*

A special case of games that satisfy Condition 1 are ones in which users do not suffer from informational interdependencies:

Definition 1 (no interdependencies) *A game $\Gamma(\lambda, \rho)$ has* no interdependencies *if $c_i(x, \rho) = c_i(x_i, x'_{-i}, \rho)$ for all $x \in S$ and $x'_{-i} \in S_{-i}$. In this case we write $c_i(x, \rho) = c_i(x_i)$.*

Under Condition 1, more information sharing by others is beneficial to users, and so a privacy enhancement that leads to greater information sharing has a positive indirect effect. In particular, we have the following theorem:

Theorem 4 *Under Condition 1 the utility of every user increases as λ decreases.*

The intuition behind Theorem 4 is simple: lowering λ increases a user's utility both directly and indirectly. The direct increase follows from lower costs. The indirect one results from an increase the amount of information shared by others (by Theorem 2), which is beneficial to the user when network effects dominate interdependencies.

Proof: Fix any $\lambda' < \lambda$, and let x' and x be the maximal (minimal) equilibria of $\Gamma(\lambda', \rho)$ and $\Gamma(\lambda, \rho)$, respectively. Theorem 2 implies that $x' \geq x$. Then

$$\begin{aligned} u_i(x', \lambda') &= v_i(x') - \lambda' \cdot c_i(x'_i, x'_{-i}, \rho) \\ &\geq v_i(x_i, x'_{-i}) - \lambda' \cdot c_i(x_i, x'_{-i}, \rho) \\ &\geq v_i(x_i, x_{-i}) - \lambda' \cdot c_i(x_i, x_{-i}, \rho) \\ &\geq v_i(x) - \lambda \cdot c_i(x_i, x_{-i}, \rho) \\ &= u_i(x, \lambda), \end{aligned}$$

where the first inequality follows from the fact that x' is an equilibrium, the second from the facts that $x_{-i} < x'_{-i}$ and that $u_i(x_i, x'_{-i}, \lambda')$ is nondecreasing in x'_{-i}, and the third since $\lambda > \lambda'$. ■

A corollary of Theorem 4 is that if there are no interdependencies then lowering λ increases user welfare.

The second condition is the opposite of Condition 1, namely that informational interdependencies dominate network effects:

Condition 2 *$u_i(x_i, x_{-i}, \lambda, \rho)$ is nonincreasing in x_{-i} for every $i \in I$ and $x_i \in S_i$.*

A special case of games that satisfy Condition 2 are ones in which users do not enjoy network effects:

Definition 2 (no network effects) *A game $\Gamma(\lambda, \rho)$ has* no network effects *if $v_i(x) = v_i(x_i, x'_{-i})$ for all i, $x \in S$, and $x'_{-i} \in S_{-i}$. In this case we write $v_i(x) = v_i(x_i)$.*

Under Condition 2, less information sharing by others is beneficial to users, and so a privacy enhancement that leads to less information sharing has a positive indirect effect. In particular, we have the following theorem:

Theorem 5 *Under Condition 2 the utility of every user increases as ρ decreases.*

Note the contrast between Theorem 5 and Theorem 4. While both provide a sufficient condition for a privacy enhancement to be welfare improving, both the enhancement and the sufficient condition are different: lowering λ vs. lowering ρ, and Condition 1 vs. Condition 2.

Like Theorem 4, the intuition behind Theorem 5 is simple: lowering ρ increases a user's utility both directly and indirectly. The direct increase follows from lower costs. The indirect one results from a decrease the amount of information shared by others (by Theorem 3), which is beneficial to the user when interdependencies dominate network effects.

Proof: Fix any $\rho' < \rho$ for which $\Gamma(\lambda, \rho')$ and $\Gamma(\lambda, \rho)$ satisfy Condition 2. Let x' and x be the respective maximal (minimal) equilibria of $\Gamma(\lambda, \rho')$ and $\Gamma(\lambda, \rho)$, and observe that $x' \leq x$ by Theorem 3. Then

$$\begin{aligned}
u_i(x', \lambda, \rho') &= v_i(x') - \lambda \cdot c_i\left(x'_i, x'_{-i}, \rho'\right) \\
&\geq v_i(x_i, x'_{-i}) - \lambda \cdot c_i\left(x_i, x'_{-i}, \rho'\right) \\
&\geq v_i(x_i, x_{-i}) - \lambda \cdot c_i\left(x_i, x_{-i}, \rho'\right) \\
&\geq v_i(x) - \lambda \cdot c_i\left(x_i, x_{-i}, \rho\right) \\
&= u_i(x, \lambda, \rho),
\end{aligned}$$

where the first inequality follows from the fact that x' is an equilibrium, the second from Condition 2 and the fact that $x_{-i} \geq x'_{-i}$, and the third from the facts that $\rho > \rho'$ and c_i is nondecreasing. ■

A corollary of Theorem 5 is that if there are no network effects then lowering ρ increases user welfare.

5 Use welfare with negative indirect effects

When indirect effects are positive, the net effects of privacy enhancements are easy to determine. In contrast, when indirect effects are negative, the situation is more complicated. In order to facilitate the analysis we will examine games with some structure, particularly concave utilities. We begin with some definitions.

Definition 3 (concave games) *The family of games* $\Gamma(\rho)$ *is* concave *if for every* $i \in I$ *and* $x_{-i} \in S_{-i}$ *both* $v_i(\cdot, x_{-i})$ *and* $v_i(\cdot, x_{-i}) - c_i(\cdot, x_{-i}, \rho)$ *are concave (as functions of* x_i*).*

In particular, a game is concave if v_i is concave and c_i is convex. The main relevant property of concave games is captured by the following lemma:

Lemma 1 *In any game* $\Gamma(\lambda, \rho)$ *belonging to a concave class of games* $\Gamma(\rho)$*, the utility function* $u_i(\cdot, x_{-i}, \lambda, \rho)$ *is concave for every* $i \in I$ *and* $x_{-i} \in S_{-i}$*.*

Proof: Let $v'_i(x_i) = v_i(x_i, x_{-i})$, $c'_i(x_i) = c_i(x_i, x_{-i}, \rho)$, and $u'_i(x_i) = u_i(x_i, x_{-i}, \lambda, \rho)$. Fix some $x_i, x'_i \in S_i$ and some $\lambda \in [0, 1]$. Concavity of v'_i implies that

$$v'_i(\lambda x_i + (1 - \lambda)x'_i) \geq \lambda v'_i(x_i) + (1 - \lambda)v'_i(x'_i),$$

and so

$$(1 - \lambda) \cdot v'_i(\lambda x_i + (1 - \lambda)x'_i) \geq (1 - \lambda) \cdot (\lambda v'_i(x_i) + (1 - \lambda)v'_i(x'_i))\,. \tag{1}$$

Concavity of $v'_i - c'_i$ implies that

$$w'_i(\lambda x_i + (1 - \lambda)x'_i) \geq \lambda w'_i(x_i) + (1 - \lambda)w'_i(x'_i),$$

where $w'_i(x_i) = v'_i(x_i) - c'_i(x_i)$, and so

$$\lambda \cdot w'_i(\lambda x_i + (1 - \lambda)x'_i) \geq \lambda \cdot (\lambda w'_i(x_i) + (1 - \lambda)w'_i(x'_i))\,. \tag{2}$$

Adding inequalities (1) and (2) and yields

$$u'_i(\lambda x_i + (1 - \lambda)x'_i) \geq \lambda u'_i(x_i) + (1 - \lambda)u'_i(x'_i)$$

as desired. ■

We now analyze the effect of privacy enhancements when indirect effects are negative, and in particular distinguish between the two extreme forms of informational interdependencies: overlapping personal information and additive personal information.

5.1 Overlapping personal information

Recall that when information is overlapping, when two users j and k share information then the shared information of user j about user i completely overlaps the shared information of user k about user i: $\mathcal{A}^i(x(\rho)) = \max\{x_i, \max_{j\neq i}\{\rho x_j\}\}$. We will show that, roughly speaking, in this case lowering λ is beneficial for user welfare, whereas lowering ρ is harmful. The general intuition underlying these results is that, when information is overlapping, the interaction between users is locally similar to an interaction with no informational interdependencies. Hence, the indirect effect of small changes in λ are similar to their effect when there are no interdependencies (which we know to be positive, by Theorem 4). Furthermore small changes to ρ have no direct effect, since the cost is determined by the maximum, and so the response to those is determined by its indirect effects.

Our theorems on the effect of privacy enhancements in games with overlapping information will each consist of two parts: the first claiming that there *exists* a user whose utility increases or decreases with a change in λ or ρ, and the second stating that *every* user's utility increases or decreases, under a certain property. We call this property of information sharing *self-exposing*: a user is self-exposing when he shares more information about himself than others share about him. In the following section we formalize this notion, and provide two settings in which this property is guaranteed to hold in the extremal equilibria: when the probability of leakage λ is small enough, and when there is a high level of homophily in the game.

5.1.1 Self-exposing levels of sharing

Small λ

Definition 4 (self-exposing (SE)) *User i is* self-exposing (SE) *in a game $\Gamma(\lambda, \rho)$ with level of sharing x if $x_i > \rho x_j$ for all $j \neq i$. The level of sharing x is an* SE profile *if every user is SE. Finally, $\Gamma(\lambda, \rho)$ is an* SE game *if its extremal equilibria are SE profiles.*

For our first proposition we need a slight strengthening of concavity: a class of games $\Gamma(\rho)$ is *strictly concave* if it is concave, and if $v_i(x_i, 0, \ldots, 0)$ is strictly concave in x_i, for every $i \in I$.

Proposition 1 *Fix ρ and a strictly concave class $\Gamma(\rho)$. Then there exists $\lambda_0 = \lambda_0(\Gamma(\rho))) > 0$ such that $\Gamma(\lambda, \rho) \in \Gamma(\rho)$ is an SE game for every $\lambda < \lambda_0$.*

Proof of Proposition 1: Fix a user i. We will show that there is $\lambda_0 > 0$ such that the best response of user i to others sharing 0 satisfies $x_i > \rho$, in the game $\Gamma(\lambda, \rho)$ with $\lambda < \lambda_0$. This will imply that when others share $x'_{-i} \geq 0$, the optimal level of sharing of i will be $x'_i \geq x_i$, because of strategic complements. This implies that $x'_i > \rho \geq \rho x'_j$ for all $j \neq i$ in every equilibrium of $\Gamma(\lambda, \rho)$, as claimed.

Fix λ and ρ, and denote by $u_i^0(x_i) = u_i(x_i, 0, \lambda, \rho)$, by $v_i^0(x_i) = v_i(x_i, 0)$, and by $c_i^0(x_i) = c_i(x_i, 0, \rho)$. Consider the minimal best response of user i to others sharing 0: $\mathrm{BR}_i = \min \arg\max_{x_i} u_i^0(x_i)$. If $\mathrm{BR}_i = 1$ then the claim follows immediately. We next show that, for small enough λ, it must be the case that $\mathrm{BR}_i > 0$. To see this, note that $\mathrm{BR}_i > 0$ is equivalent to

$$\frac{dv_i^0(0)}{dx_i} > \lambda \cdot \frac{dc_i^0(0)}{dx_i}.$$

The LHS must be positive due to v_i being nondecreasing and strictly convex. Thus, the inequality is satisfies whenever λ is small enough (but positive).

Thus, the best response of user i is in the interior. Let $x_i^0(\lambda)$ denote BR_i in the game $\Gamma(\lambda, \rho)$. We now show that, given that BR_i is in the interior, $x_i^0(\lambda) \to 1$ as $\lambda \to 0$. This will prove the claim.

To this end, fix some $\varepsilon > 0$. We will show that there exists $\lambda > 0$ such that $x_i^0(\lambda) > 1 - \varepsilon$. Now,

$$\begin{aligned} & x_i^0(\lambda) > 1 - \varepsilon \\ & \quad \Leftrightarrow \frac{du_i^0(1-\varepsilon)}{dx_i} > 0 \\ & \quad \Leftrightarrow \frac{dv_i^0(1-\varepsilon)}{dx_i} > \lambda \cdot \frac{dc_i^0(1-\varepsilon)}{dx_i}, \end{aligned}$$

where the first iff follows from concavity of u_i^0. Note that the LHS of the last inequality above is positive since v_i^0 is strictly convex and increasing. Thus, the inequality holds for small enough (but positive) λ. ■

Networks In more realistic settings, users derive utility only from some other users' shared information, and not all. Furthermore, users' information is interdependent with some other users' information, but not with all. We can describe such a network of interaction by two graphs, one capturing network effects and the other capturing informational interdependencies. Let $G^v(I, E^v)$ and $G^c(I, E^c)$ be two directed graphs, where the nodes in each are the users I. The idea is that there is a directed edge e_{ij} from node i to node j in the graph G^v (formally, $e_{ij} \in E^v$) if and only if v_i depends on x_j, and there is a directed edge e_{ij} from node i to node j in the graph G^c (formally, $e_{ij} \in E^c$) if and only if c_i depends on x_j. Fixing a network $N = (G^v, G^c)$, the utilities of users in the game $\Gamma(\lambda, \rho)$ on network N are

$$u_i(x, \lambda, \rho) = v_i\left(x_i, \{x_j : e_{ij} \in E^v\}\right) - \lambda \cdot c_i\left(x_i, \{x_j : e_{ij} \in E^c\}, \rho\right).$$

In the following, we restrict attention to networks with a high amount of homophily: in particular, we will require that users who are connected on the network (more specifically, on the interdependency graph) have the same utility functions (see McPherson et al., 2001, for a survey of homophily in social and other networks). In words, the following definition states that if x_j is an input to c_i—that is, if user j has information about user i—then the functions u_i and u_j are equivalent (up to a permutation of the inputs). More formally,

Definition 5 (network-symmetric) *A game $\Gamma(\lambda, \rho)$ on network $N = (G^v, G^c)$ is* network-symmetric *if for every $i \in I$ and j satisfying $e_{ij} \in E^c$ there exists a permutation $\pi : I \mapsto I$ with $\pi(i) = j$, for which $v_k \equiv v_{\pi(k)} \circ \pi$ and $c_k \equiv c_{\pi(k)} \circ \pi$.*

A trivial example of a network-symmetric game consists of symmetric games, in which both G^v and G^c are cliques, and $u_i(x_1, \ldots, x_{|I|}) = u_{\pi(i)}(x_{\pi(1)}, \ldots, x_{\pi(|I|)})$ for *every* permutation π on I.

We slightly modify the definition of self-exposing, to accommodate the presence of the network.

Definition 6 (self-exposing (SE) on a network) *User i is* self-exposing (SE) *in a game $\Gamma(\lambda, \rho)$ on network $N = (G^v, G^c)$ with level of sharing x if $x_i > \rho x_j$ for all $j \neq i$ with $e_{ij} \in E^c$. The level of sharing x is an* SE profile on network N *if every user is SE on N. Finally, $\Gamma(\lambda, \rho)$ is an* SE game on network N *if its extremal equilibria are SE profiles on N.*

Proposition 2 *Every network-symmetric game $\Gamma(\lambda, \rho)$ on network $N = (G^v, G^c)$ is an SE game on N.*

Proof of Proposition 2: Fix users $i, j \in I$ with $e_{ij} \in E^c$. We first show that in the maximal equilibrium $\overline{x}$ it holds that $\overline{x}_i \geq \overline{x}_j$, and so, in particular, $\overline{x}_i > \rho\overline{x}_j$.

Suppose towards a contradiction that $\overline{x}_i < \overline{x}_j$. Let π be the permutation with $\pi(i) = j$ guaranteed by the network-symmetry of the game. Furthermore, consider the profile x that satisfies $x_k = \overline{x}_{\pi(k)}$ for all $k \in I$. Due to network-symmetry, x is also an equilibrium, but note that in this equilibrium $x_i = \overline{x}_j$. Finally, since the set of pure Nash equilibria of games of strategic complements form a lattice (Zhou, 1994), there exists a third equilibrium x' such that $x' > \overline{x}$ and $x' \geq x$. The former inequality is strict since $x'_i \geq x_i = \overline{x}_j > \overline{x}_i$. The existence of x' contradicts the maximality $\overline{x}$.

Next, we show that in the minimal equilibrium $\underline{x}$ it holds that $\underline{x}_i \geq \underline{x}_j$, and so $\underline{x}_i > \rho\underline{x}_j$. Suppose towards a contradiction that $\underline{x}_i < \underline{x}_j$. Consider the profile x that satisfies $x_k = \underline{x}_{\pi(k)}$ for all $k \in I$. Due to network-symmetry, x is also an equilibrium, but note that in this equilibrium $x_i = \underline{x}_j$. Once again, because the set of equilibria form a lattice, there exists a third equilibrium x' such that $x' \leq \underline{x}$ and $x' \leq x$. It remains to show that $x' \neq \underline{x}$, which will imply that $x' < \underline{x}$ and so contradict the minimality of $\underline{x}$.

To see why $x' \neq \underline{x}$, consider the cycle implied by the permutation π, starting at node $i \stackrel{\text{def}}{=} i_1$, namely the set $\{i_1, \ldots, i_m\} \subseteq I$ satisfying $\pi(i_n) = i_{n+1}$ for all $n \in \{1, \ldots, m-1\}$, and $\pi(i_m) = i_1$. Since $x' \leq \underline{x}$, $x' \leq x$, and x is a permutation of $\underline{x}$, it would only be possible that $x' = \underline{x}$ if $\underline{x}_i = \underline{x}_{i_n}$ for all $n \in \{1, \ldots, m\}$. This is false, however, since, in particular, $\underline{x}_i < \underline{x}_j = \underline{x}_{i_2}$. ■

5.1.2 Decreasing λ

Our first main theorem for games with overlapping information is:

Theorem 6 *In any concave game $\Gamma(\lambda, \rho)$ with overlapping personal information there exists a user whose utility increases as λ decreases. If either $\Gamma(\lambda, \rho)$ is strictly concave with $\lambda < \lambda_0(\Gamma(\rho))$ or $\Gamma(\lambda, \rho)$ is network-symmetric, then* every *user's utility increases as λ decreases.*

The intuition behind Theorem 6 is that for a user who shares more information about himself than others share about him, the nature of overlapping information is such that he only incurs costs from information he shares. Hence, as long as he does not vary the amount of information he shares by much, it is as if there were no interdependencies. And in a game with no interdependencies lowering λ is beneficial, as demonstrated by Theorem 4. Finally, concavity of the game renders this "local" property sufficient.

Theorem 6 relies on the following lemmas:

Lemma 2 *Fix a concave game with no interdependencies, a strategy profile x, and a strategy x'_i. If $u_i(x'_i, x_{-i}) > u_i(x)$ then for every $\varepsilon > 0$ there exists x''_i such that $|x''_i - x'_i| < \varepsilon$ and such that $u_i(x''_i, x_{-i}) > u_i(x)$.*

Proof: Fix any $\delta \in (0,1)$, and let $x''_i = (1-\delta)x_i + \delta x'_i$. By Lemma 1, u_i is a concave function of x_i for every fixed x_{-i}. Thus,

$$u_i(x''_i, x_{-i}) \geq (1-\delta)u_i(x_i, x_{-i}) + \delta u_i(x'_i, x_{-i}) > u_i(x_i, x_{-i}).$$

Choosing δ sufficiently small completes the proof. ■

Lemma 3 *Fix a concave game $\Gamma(\lambda, \rho)$ with overlapping information, an equilibrium x, and an SE user i. Then for every x'_i,*

$$v_i(x) - \lambda \cdot c_i(x_i) \geq v_i(x'_i, x_{-i}) - \lambda \cdot c_i(x'_i).$$

Note that the left-hand-side of the inequality is the utility of user i, since the facts that the game has overlapping information and that x_i is SE mean that $c_i(x) = c_i(x_i)$. Now, since x is an equilibrium it must be the case that $u_i(x) \geq u_i(x'_i, x_{-i})$, but note that this does not imply the right-hand-side of the inequality in the lemma, since it is possible that $u_i(x'_i, x_{-i}) < v_i(x'_i, x_{-i}) - \lambda \cdot c_i(x'_i)$.

Proof of Lemma 3: Consider the game $\Gamma(\lambda, 0)$, which is the same game as $\Gamma(\lambda, \rho)$ except that there are no interdependencies (that is, $u_i(x, \lambda, 0) = v_i(x) - \lambda \cdot c_i(x_i)$). If x_i is a best response to x_{-i} in $\Gamma(\lambda, \rho)$, then I claim it must also be a best-response in $\Gamma(\lambda, 0)$. To see this, suppose towards a contradiction that there is a profitable deviation x'_i to x_{-i} in $\Gamma(\lambda, 0)$:

$$v_i(x'_i, x_{-i}) - \lambda \cdot c_i(x'_i) > v_i(x) - \lambda \cdot c_i(x_i).$$

Note that this does not immediately imply the contradiction that x'_i is a profitable deviation in $\Gamma(\lambda, \rho)$ as well, since a deviation to x'_i in $\Gamma(\lambda, \rho)$ leads to utility $v_i(x'_i, x_{-i}) - \lambda \cdot c_i\left(\max\{\, x'_i, \max_{j \neq i}\{\rho x_j\}\,\}\right)$, and this may not be higher than $v_i(x) - \lambda \cdot c_i(x_i)$. However, by Lemma 2, for any $\varepsilon > 0$ there exists x''_i such that in $\Gamma(\lambda, 0)$, x''_i is also a profitable deviation. Choose ε small enough so that $x''_i > \max_{j \neq i}\{\rho x_j\}$. This is possible since i is an SE user. This implies that in the game $\Gamma(\lambda, \rho)$,

$$
\begin{aligned}
v_i(x''_i, x_{-i}) - \lambda \cdot c_i\left(\max\{\, x''_i, \max_{j \neq i}\{\rho x_j\}\,\}\right)& \\
&= v_i(x''_i, x_{-i}) - \lambda \cdot c_i(x''_i) \\
&> v_i(x) - \lambda \cdot c_i(x_i),
\end{aligned}
$$

where the inequality follows from Lemma 2. Thus, x''_i is a profitable deviation in $\Gamma(\lambda, \rho)$, implying the contradiction that x_i is not a best response to x_{-i} in $\Gamma(\lambda, 0)$.

Hence, x_i is a best response to x_{-i} in $\Gamma(\lambda, 0)$, which implies that $u_i(x, \lambda, 0) \geq u_i(x'_i, x_{-i}, \lambda, 0)$. This is equivalent to the inequality in the statement of the lemma.

■

We now prove Theorem 6.

Proof of Theorem 6: Since $\Gamma(\lambda, \rho)$ has overlapping information, the cost function c_i depends only on the maximum of its inputs. Fix any $\lambda' < \lambda$, and let x' and x be the respective maximal (minimal) equilibria of $\Gamma(\lambda', \rho)$ and $\Gamma(\lambda, \rho)$. As before, Theorem 2 implies that $x' \geq x$. For the first part of the theorem, fix some $i \in \arg\max_j x'_j$, and note that i is an SE user under x' since $x'_i \geq x'_j$ for all $j \neq i$, and so $x'_i > \rho x'_j$ for all $j \neq i$. Then

$$
\begin{aligned}
u_i(x', \lambda') &= v_i(x') - \lambda' \cdot c_i\left(\max\{x'_i, \max_{j \neq i}\{\rho x'_j\}\}\right) \\
&= v_i(x') - \lambda' \cdot c_i(x'_i) && (3) \\
&\geq v_i(x_i, x'_{-i}) - \lambda' \cdot c_i(x_i) && (4) \\
&\geq v_i(x) - \lambda \cdot c_i\left(\max\{x_i, \max_{j \neq i}\{\rho x_j\}\}\right) && (5) \\
&= u_i(x, \lambda),
\end{aligned}
$$

where (3) follows from the choice of x'_i as maximal, (4) from Lemma 3, and (5) from the facts that $\lambda' < \lambda$, $x \leq x'$, and v_i is nondecreasing.

The second part of the theorem follows from the observations that if $\lambda < \lambda_0(\Gamma(\rho)))$ then also $\lambda' < \lambda_0(\Gamma(\rho)))$, and if $\Gamma(\lambda, \rho)$ is network-symmetric then so is $\Gamma(\lambda', \rho)$. Thus, by Propositions 1 and 2, the game $\Gamma(\lambda', \rho)$ is an SE game. Finally, in such games the analysis above holds for every user (replacing $\max_{j \neq i}$ by $\max_{j: e_{ij} \in E^c}$ if the game is network-symmetric). ■

Remark 7 The fact that all users' utilities increase as λ decreases does not hold when the resulting game $\Gamma(\lambda', \rho)$ is not SE, and in that case it is possible that lowering λ leads some users to obtain strictly lower welfare. This is demonstrated by Example 3, a concave 2-user game without network effects, in Appendix B.

Remark 8 Results analogous to Theorem 6 do not generally hold in games that are not concave. This is demonstrated by Example 4 in Appendix B, a symmetric, non-concave game with overlapping information with the property that lowering λ leads to a decrease in all users' utilities.

5.1.3 Lowering ρ

We have seen that under overlapping personal information, lowering λ is beneficial. How about lowering ρ? Here we have two sets of results, one for games that satisfy Condition 1 and one for concave games. For both, the broad conclusion is that enhancing privacy by lowering ρ is *harmful* to users.

Theorem 9 *In any game $\Gamma(\lambda, \rho)$ with overlapping personal information satisfying Condition 1 there exists a user whose utility decreases as ρ decreases. If either $\Gamma(\lambda, \rho)$ is strictly concave with $\lambda < \lambda_0(\Gamma(\rho)))$ or $\Gamma(\lambda, \rho)$ is network-symmetric, then* every *user's utility decreases as ρ decreases.*

To see the intuition, note that for for a user who shares more information about himself than others share about him, informational interdependencies do not matter, and so for him there is no indirect effect of lowering ρ. He does, however, face the indirect effect of less sharing by others, and this, by Condition 1, is harmful.

Proof: Fix any $\rho' < \rho$, and let x' and x be the respective maximal (minimal) equilibria of $\Gamma(\lambda, \rho')$ and $\Gamma(\lambda, \rho)$. As before, Theorem 3 implies that $x' \leq x$. For the first part of the theorem, fix some $i \in \arg\max_j x'_j$, and note that i is an SE user under x' since $x'_i \geq x'_j$ for all $j \neq i$, and so $x'_i > \rho' x'_j$ for all $j \neq i$. Then

$$
\begin{aligned}
u_i(x, \lambda, \rho) &= v_i(x) - \lambda \cdot c_i\left(\max\{\, x_i, \max_{j \neq i}\{\rho x_j\}\,\}\right) \\
&\geq v_i(x'_i, x_{-i}) - \lambda \cdot c_i\left(\max\{\, x'_i, \max_{j \neq i}\{\rho x_j\}\,\}\right) \qquad (6) \\
&\geq v_i(x') - \lambda \cdot c_i\left(\max\{\, x'_i, \max_{j \neq i}\{\rho' x'_j\}\,\}\right) \qquad (7) \\
&= v_i(x') - \lambda \cdot c_i(x'_i) \qquad (8) \\
&= v_i(x') - \lambda \cdot c_i\left(\max\{\, x'_i, \max_{j \neq i}\{\rho' x'_j\}\,\}\right) \\
&= u_i(x', \lambda, \rho'),
\end{aligned}
$$

where (6) follows since x is an equilibrium; (7) by Condition 1; and (8) by the choice of x'_i as maximal.

The second part of the theorem follows from the observations that if $\lambda < \lambda_0(\Gamma(\rho)))$ then also $\lambda < \lambda_0(\Gamma(\rho')))$, and if $\Gamma(\lambda, \rho)$ is network-symmetric then so is $\Gamma(\lambda, \rho')$. Thus, by Propositions 1 and 2, the game $\Gamma(\lambda', \rho)$ is an SE game. Finally, in such games the analysis above holds for every user (replacing $\max_{j \neq i}$ by $\max_{j: e_{ij} \in E^c}$ if the game is network-symmetric). ■

Our second result on lowering ρ in games with overlapping personal information concerns con cave games. Recall Theorem 6, which states that in such games, increasing privacy by lowering λ increases user welfare utility. The following theorem, in contrast, states that in such games increasing privacy by lowering ρ *decreases* user welfare. Furthermore, in SE games, user welfare is unaffected by decreases in ρ.

Theorem 10 *In any game concave $\Gamma(\lambda, \rho)$ with overlapping personal information there exists a user whose utility decreases as ρ decreases. If either $\Gamma(\lambda, \rho)$ is strictly concave with $\lambda < \lambda_0(\Gamma(\rho)))$ or $\Gamma(\lambda, \rho)$ is network-symmetric, then no user's utility is affected by a decrease in ρ.*

The intuition for Theorem 10 is similar to that of Theorem 9: for a user who shares more information about himself than others share about him, there are no direct effects from lowering ρ (for small variations). The indirect effect here lowers the benefit of network effects, which is harmful. And as before, concavity guarantees that these local effects suffice.

Proof: Fix any $\rho' < \rho$, and let x' and x be the respective maximal (minimal) equilibria of $\Gamma(\lambda, \rho')$ and $\Gamma(\lambda, \rho)$. As before, Theorem 3 implies that $x' \leq x$. For the first part of the theorem, fix some $i \in \arg\max_j x_j$, and note that i is an SE user under x since $x_i \geq x_j$ for all $j \neq i$, and so $x_i > \rho x_j$ for all $j \neq i$.

$$\begin{aligned}
u_i(x, \lambda, \rho) &= v_i(x) - \lambda \cdot c_i \left(\max\{x_i, \max_{j \neq i}\{\rho x_j\}\} \right) \\
&= v_i(x) - \lambda \cdot c_i(x_i) && (9) \\
&\geq v_i(x'_i, x_{-i}) - \lambda \cdot c_i(x'_i) && (10) \\
&\geq v_i(x') - \lambda \cdot c_i \left(\max\{x'_i, \max_{j \neq i}\{\rho' x'_j\}\} \right) && (11) \\
&= u_i(x', \lambda, \rho'),
\end{aligned}$$

where (9) follows since x_i is maximal; (10) by Lemmas 2 and 3; and (11) since v_i and c_i are nondecreasing.

For the second part of the theorem, observe that if $\lambda < \lambda_0(\Gamma(\rho))$ then also $\lambda < \lambda_0(\Gamma(\rho'))$, and if $\Gamma(\lambda, \rho)$ is network-symmetric then so is $\Gamma(\lambda, \rho')$. Thus, by Propositions 1 and 2, the game $\Gamma(\lambda', \rho)$ is an SE game. We will show that in this case, $x = x'$, implying also that $u_i(x, \lambda, \rho) = u_i(x', \lambda, \rho')$, since then both inequalities above become equalities (replacing $\max_{j \neq i}$ by $\max_{j: e_{ij} \in E^c}$ if the game is network-symmetric).

To show that $x = x'$, we will argue that any equilibrium in an SE game $\Gamma(\lambda, \rho)$ is also an equilibrium in $\Gamma(\lambda, \rho')$, and vice versa. Observe that in both $\Gamma(\lambda, \rho)$ and $\Gamma(\lambda, \rho')$, for any profile $\overline{x}$ it holds that $u_i(\overline{x}) = v_i(\overline{x}) - \lambda \cdot c_i(\overline{x}_i)$. Now, consider the profile x that is an equilibrium in $\Gamma(\lambda, \rho)$, as well as a possible deviation x''_i by some user i. By Lemmas 2 and 3, if x''_i is a profitable deviation in $\Gamma(\lambda, \rho')$, then there exists another profitable deviation x'''_i in $\Gamma(\lambda, \rho')$ that is arbitrarily close to x_i. In particular, one can choose $x'''_i > \rho x_j$ for all $j \neq i$ (or for all j with $e_{ij} \in E^c$ if the game is network-symmetric). But this implies that

$$u_i(x'''_i, x_{-i}) = v_i(x'''_i, x_{-i}) - \lambda \cdot c_i(x'''_i) > u_i(x)$$

in *both* $\Gamma(\lambda, \rho)$ and $\Gamma(\lambda, \rho')$, contradicting the assumption that x is an equilibrium in the former. Thus, an equilibrium x in $\Gamma(\lambda, \rho)$ is also an equilibrium in $\Gamma(\lambda, \rho')$.

Observing that the argument above does not depend on whether $\rho < \rho'$ or vice versa completes the proof. ■

5.2 Additive personal information

In the previous section we showed that when information is overlapping, decreasing λ is beneficial while decreasing ρ is harmful to user welfare. The situation is not as clear-cut with additive information, however. We will show that even in perhaps the simplest class of games—symmetric, concave games without network effects and with linear cost functions $\hat{c}_i$—lowering λ can be both beneficial and harmful to user welfare, depending on the magnitude of ρ. In contrast, we show that in the general class of games with additive information and linear costs lowering ρ is always *beneficial*.

Consider the following example:

Example 1 Suppose $I = \{1, 2\}$, and fix some $k \in (.5, 1)$ and $\lambda \in (1/(1+k), 1/k)$. Furthermore, for both $i \in I$ let $v_i(x_i, x_{-i}) = \log(k + x_i) - \log(k)$, and $c_i(x) = (x_i + \rho x_{2-i})$.

The conditions on λ guarantee that the solution is interior, and in equilibrium we have

$$x_i = \frac{1}{\lambda} - k$$

for each user i. Then in equilibrium,

$$\begin{aligned} u_i(x, \lambda, \rho) &= \log \frac{1}{\lambda} - \lambda(1+\rho)\left(\frac{1}{\lambda} - k\right) \\ &= \log \frac{1}{\lambda} + (1+\rho)\lambda k - (1+\rho). \end{aligned}$$

Now fix $\lambda' = 1$ and $\lambda = 1/(2k)$, and observe that $\lambda < \lambda'$. Let x and x' be the maximal (minimal) equilibria of $\Gamma(\lambda)$ and $\Gamma(\lambda')$, respectively.

Plugging in these values of λ and λ' we get that

$$u_i(x, \lambda, \rho) = \log(2k) + \frac{1+\rho}{2} - (1+\rho),$$

and

$$u_i(x', \lambda', \rho) = (1+\rho)(k-1).$$

Then

$$
\begin{aligned}
u_i(x, \lambda, \rho) &> u_i(x', \lambda', \rho) \\
&\Leftrightarrow \log(2k) + \frac{1+\rho}{2} - (1+\rho) > (1+\rho)(k-1) \\
&\Leftrightarrow \log(2k) + \frac{1+\rho}{2} > (1+\rho)k \\
&\Leftrightarrow \log(2k) > (1+\rho)\left(k - \frac{1}{2}\right) \\
&\Leftrightarrow \rho < \frac{\log(2k)}{k - \frac{1}{2}} - 1.
\end{aligned}
$$

Now, when k approaches 1, the right-hand-side of the bound approaches roughly 0.386. In this case, when the information correlation is relatively small (i.e., below 0.386), lowering λ' to λ leads to higher user welfare. However, when the informational interdependencies are relatively high (i.e., above 0.386), lowering λ' to λ leads to lower user welfare.

The conclusion from this example is rather intuitive. Recall that the direct effect of reducing λ is positive. The indirect effect is an increase in the volume of information shred, which by assumption is harmful. Hence, if this harm is large—which happens when the level of interdependence is high—then the indirect effect dominates, and welfare decreases. Similarly, when the harm is small then the direct effect dominates and welfare increases.

In contrast with the ambiguous nature of lowering λ, the following theorem states that under additive personal information with linear costs, lowering ρ leads to higher utility to all users. This contrasts with the effect of lowering ρ in games with overlapping personal information, in which utilities either remained unchanged, or at least one user obtains lower utility.

Theorem 11 *In any game with additive information and linear costs, user welfare increases as ρ decreases.*

The main intuition is that under additive personal information, lowering ρ has no direct effect on welfare. This follows from the following lemma, which is one component of the subsequent statement that the extremal equilibria of a game with additive personal information are unaffected by changes in ρ.

Lemma 4 *If personal information is additive and cost functions $\hat{c}_i$ linear, then the extremal equilibria of $\Gamma(\lambda, \rho)$ are remain fixed as ρ varies.*

Proof: We will show that $u_i(x, \lambda, \rho)$ has both decreasing and increasing differences in x_i and ρ. The result will then follow from Theorem 1.

Fix $x_{-i} \in S_{-i}$, as well as any $x_i < x_i' \in S_i$ and $\rho < \rho'$. Then:

$$\begin{aligned}
&[u_i(x_i', x_{-i}, \rho') - u_i(x_i, x_{-i}, \rho')] - [u_i(x_i', x_{-i}, \rho) - u_i(x_i, x_{-i}, \rho)] \\
&\quad = v_i(x_i', x_{-i}) - \lambda \hat{c}_i \left(x_i' + \sum_{j \neq i} \rho' x_j \right) - v_i(x_i, x_{-i}) + \lambda \hat{c}_i \left(x_i + \sum_{j \neq i} \rho' x_j \right) \\
&\qquad - v_i(x_i', x_{-i}) + \lambda \hat{c}_i \left(x_i' + \sum_{j \neq i} \rho x_j \right) + v_i(x_i, x_{-i}) - \lambda \hat{c}_i \left(x_i + \sum_{j \neq i} \rho x_j \right) \\
&\quad = -\lambda \hat{c}_i(x_i') - \lambda \hat{c}_i \left(\sum_{j \neq i} \rho' x_j \right) + \lambda \hat{c}_i(x_i) + \lambda \hat{c}_i \left(\sum_{j \neq i} \rho' x_j \right) \\
&\qquad + \lambda \hat{c}_i(x_i') - \lambda \hat{c}_i \left(\sum_{j \neq i} \rho x_j \right) - \lambda \hat{c}_i(x_i) - \lambda \hat{c}_i \left(\sum_{j \neq i} \rho x_j \right) \\
&\quad = 0.
\end{aligned}$$

■

Theorem 11 relies on Lemma 4, which states that the extremal equilibria of $\Gamma(\lambda, \rho)$ are both nondecreasing and nonincreasing in ρ, or in other words: unchanged.

Proof of Theorem 11: Fix any $\rho' < \rho$, and let x' and x be the respective maximal (minimal) equilibria of $\Gamma(\lambda, \rho')$ and $\Gamma(\lambda, \rho)$. Lemma 4 implies that $x' = x$. Now,

$$\begin{aligned}
u_i(x', \lambda, \rho') &= v_i(x') - \lambda c_i (x', \rho') \\
&= v_i(x) - \lambda c_i (x, \rho') \\
&\geq v_i(x) - \lambda c_i (x, \rho) \\
&= u_i(x, \lambda, \rho).
\end{aligned}$$

■

6 Privacy Paradoxes

The so-called "privacy paradox" is the apparent dichotomy between consumers' stated privacy concerns and actual behavior (Norberg et al., 2007). There are different "explanations" for this paradox: for example, stated preferences may be generic attitudes about privacy, whereas behavior occurs within a particular context (see Acquisti et al. (2016) for additional discussion and references).

Here we will use an extension of the model to offer another explanation: Privacy-enhancing behavior is costly, and so individuals refrain from engaging in it. However, they may be in favor of such behavior if *all* individuals are forced to comply.

Fix a game $\Gamma(\lambda_0, \rho)$. Suppose each user i may *choose* to enhance his own privacy, so that the probability of his information being leaked is $\lambda^i < \lambda_0$, at an additional cost of $\overline{c}$. The other user's probability of leakage remains λ_0 unless he also pays the cost. Then there are games $\Gamma(\rho)$ and values $\lambda = \lambda^1 = \lambda^2 < \lambda_0$ and $\overline{c}$ in which the following hold: (i) In equilibrium, no user chooses to enhance privacy; but (ii) the utility of every user in every equilibrium of $\Gamma(\lambda, \rho)$, minus the cost $\overline{c}$, is higher than in every equilibrium of $\Gamma(\lambda_0, \rho)$.

Example 2 Consider a symmetric game with two users and no informational interdependencies, in which $v_i(x) = x_1 + x_2$ and $c_i(x) = k \cdot x_i^2/2$ for some positive k. For any (λ^1, λ^2) observe that, if the maximal equilibrium is interior, then it satisfies $x_i = 1/(\lambda^i k)$ for both users i (the equilibrium is interior whenever $\lambda^i k > 1$). Denote by $u(\lambda)$ the utility of a user in the maximal equilibrium, when the leakage probabilities are both λ. Then

$$u(\lambda) = \frac{2}{\lambda k} - \frac{\lambda k}{2} \cdot \frac{1}{(\lambda k)^2} = \frac{3}{2\lambda k}.$$

Consider a privacy enhancement that lowers both users' leakage probabilities to $\lambda < \lambda_0$. Then the utility gain under the privacy enhancement is $u(\lambda) - u(\lambda_0) = \frac{3}{2k}\left(\frac{1}{\lambda} - \frac{1}{\lambda_0}\right)$. Thus, if the cost $\overline{c} < u(\lambda) - u(\lambda_0)$, there is a welfare gain in requiring all users pay to implement the privacy enhancement.

But what if the enhancement is optional? Suppose users simultaneously decide whether or not to pay $\overline{c}$ to lower their own leakage probability, at the same time as they choose their respective levels of sharing. We will show that under the right setting of parameters, the unique equilibrium will be for both users *not* to choose the privacy enhancement. Fix a user i, and suppose the other user is sharing an amount y of information. Given a leakage probability, user i will best respond to y. What is the difference in his utility gain when going from λ_0 to $\lambda^i < \lambda_0$, given that he best-responds? We will show that this utility gain is less than $u(\lambda) - u(\lambda_0)$.

Under λ_0, user i's best response is $1/(\lambda_0 k)$, and under λ^i it is $1/(\lambda^i k)$ (assuming interior equilibria). Then the utility gain is

$$\begin{aligned} u(1/(\lambda^i k), y, \lambda, \rho) &- u(1/(\lambda_0 k), y, \lambda_0, \rho) \\ &= \frac{1}{\lambda^i k} + y - \frac{\lambda^i k}{2(\lambda^i k)^2} - \frac{1}{\lambda_0 k} - y + \frac{\lambda_0 k}{2(\lambda^i k)^2} \\ &= \frac{1}{2k}\left(\frac{1}{\lambda} - \frac{1}{\lambda_0}\right). \end{aligned}$$

Note that this is lower than $u(\lambda) - u(\lambda_0)$ since it does not contain the benefits of network effects. Thus, one can choose a value $\overline{c} < u(\lambda) - u(\lambda_0)$ under which no user will choose to pay the cost $\overline{c}$—in particular, $\overline{c} \in (k', 3k')$, where $k' = \frac{1}{2k}\left(\frac{1}{\lambda} - \frac{1}{\lambda_0}\right)$—and so the unique equilibrium is for neither user to choose the privacy enhancement and $x = (1/(\lambda_0 k), 1/(\lambda_0 k))$. However, the privacy enhancement would be welfare improving, since $\overline{c} < u(\lambda) - u(\lambda_0)$.

The "paradox" above is driven by network effects: each user's gain from a higher action and lower λ is less than $\overline{c}$. However, each user's utility gain from others' higher actions, together with gain from lower λ, is *greater* than $\overline{c}$.

The "opposite" paradox is also possible. That is, there exist games and values $\overline{c}$ in which: (i) In equilibrium, all users choose the privacy enhancement; and (ii) the utility of every user in every equilibrium of $\Gamma(\lambda_0, \rho_0)$ is higher than in every equilibrium of $\Gamma(\lambda, \rho)$, even *without* taking costs $\overline{c}$ into account.

For example, this holds if $\overline{c} = 0$, indirect effects are negative and dominate direct effects (as in Example 1 with $\rho > 0.386$), and $\rho = \rho_0$. In every equilibrium, each user chooses the privacy-enhancing policy: it is costless, and directly improves welfare. In $\Gamma(\lambda, \rho_0)$, users share more information than in $\Gamma(\lambda_0, \rho_0)$. However, because indirect effects dominate direct effects, this is harmful.

7 Extensions

7.1 Lowering the cost of information misuse

An additional privacy enhancement not considered in the current paper involves lowering the costs associated with information leakage and misuse. For example, lower costs may be implemented by insurance against identity theft. This could be formalized in our model by changing the function c_i to some c_i', where for every x it holds that $c_i'(x, \rho) \leq c_i(x, \rho)$. One might ask the same questions as we have asked: does such an enhancement increases the amount of information shared, and does it lead to a welfare increase for the users?

One special case of such a privacy enhancement is to lower costs by a fixed factor, namely $c'(x, \rho) = \delta \cdot c(x, \rho)$ for some positive $\lambda < 1$. In fact, this special case is identical to lowering λ in our model. Before the enhancement, utilities are $u_i(x, \lambda, \rho) = v_i(x) - \lambda \cdot c_i(x, \rho)$, and after lowering costs to $c' = \delta c$ they are $u_i(x, \lambda, \rho) = v_i(x) - \lambda \cdot \delta \cdot c_i(x, \rho)$. This is easily seen to be equivalent to lowering λ to $\lambda' = \delta \cdot \lambda$, and maintaining the same cost function c_i. Thus, at least in this special case, the benefits of lowering costs are the same as those of lowering the probability of information leakage.

7.2 General information overlap

One simplifying assumption we made is that the amount of information about user i that a different user j discloses is ρx_j—that is, it is linear in x_j, with the same parameter ρ for each pair of players. One simple extension is to instead have an interdependency matrix $\mathcal{R}$, with entries ρ_{ji} for every $j \neq i$. The amount of information user j discloses about i would then be $\rho_{ji} x_j$. The corresponding privacy enhancement would be to lower all the entries of $\mathcal{R}$.

Even more generally, we could allow ρ_{ji} to be an arbitrary increasing function of x_j. In this case the utility of user i would be

$$u_i(x) = v_i(x) - \lambda \cdot c_i \left(x_i, (\rho_{ji}(x_j))_{j \neq i} \right).$$

Denote by $\mathcal{R} = (\rho_{ji})_{i,j \in I}$ the vector of such functions, with the partial order that satisfies $\mathcal{R}' \prec \mathcal{R}$ if and only if $\rho_{ji}'(x_j) \leq \rho_{ji}(x_j)$ for all $i, j \in I$ and $x_j \in S_j$.

Note that the set of parameters T in Theorem 1 can be any partially-ordered set, so we can use $\mathcal{R}$ as the parameter. In particular, this implies that Theorems 2 and 3 hold in this more general model, as do all other theorems in this paper.

7.3 Added value for greater leakage and interdependence

In this section we consider the effect of extending our model to incorporate a possible benefit to greater interdependence or greater leakage, in addition to the cost discussed thus far. Observe first that the privacy paradox of Section 6 is actually a special case: when λ is larger there is a greater chance of leakage, but there is also the benefit of not having to pay the cost $\overline{c}$ for the privacy enhancement. As we have seen in that section, this can lead to a so-called "privacy paradox" in which the privacy enhancement is welfare enhancing, but users prefer not to implement it. Thus, when there is a benefit to greater leakage or interdependence, we can expect such "paradoxes" to arise.

For the rest of this section we examine the robustness of our results to extensions of the model that allow for a more general way of incorporating a benefit to greater λ and ρ.

Benefits of greater leakage While we have thought of leakage of personal information as necessarily harmful, there are settings where a greater probability of leakage—caused by greater exposure to others—has some benefits. For example, on Facebook there may be value to making one's shared photos public. In order to incorporate this benefit, define the class of games Γ as before, but now let v_i be a function of both x and λ. To capture the *benefit* of greater exposure, suppose that v_i is increasing with λ.

A first observation in the extended model is that all the results about lowering ρ still hold, as they take λ to be fixed. Additionally, Theorem 2 holds as long as u_i satisfies decreasing differences in $(x_i; \lambda)$. This holds, for example, if v_i satisfies decreasing differences in $(x_i; \lambda)$. The extension to the proof is straightforward.

Next, consider the following condition (in the spirit of Conditions 1 and 2) on the tradeoff between the benefit of greater exposure, captured by v_i, and the cost associated with this leakage, captured by $\lambda \cdot c_i$: For every $i \in I$ and $x \in S$, increasing λ leads to a decrease in $u_i(x, \lambda)$. That is, while there is a benefit to increased λ, the cost associated with this increase is always greater. In this case it is straightforward to see that Theorem 6 will continue to hold (in particular, inequality (5) would continue to hold.

Benefits of interdependence While the interdependence of information leads to increased cost due to information leakage, there are some settings in which such interdependence may be beneficial. For example, users of a social network may benefit from the interdependence amongst their peers, as they thus obtain more information: a user may observe not only the information x_i shared by peer i, but also the information ρx_j shared by peer j about i. Furthermore, a user may also benefit from others sharing information about him—for example, being tagged in someone else's photo is often desirable.

To take such benefits of interdependence into account, extend our model to incorporate ρ into the v_i function. Thus, define the class of games Γ as before, but now let v_i be a function of both x and ρ. To capture the *benefit* of interdependence, suppose that v_i is increasing with ρ.

As above, a first observation in the extended model is that Theorem 2 still holds, as do all the results that apply to lowering λ. For other results the situation is a bit different, and in general they require additional assumptions in order to hold. First, Theorem 3 holds as long as v_i has increasing differences in x_i and ρ. The extension of the proof is straightforward.

Next, consider the following condition (in the spirit of Conditions 1 and 2) on the tradeoff between the benefit of greater interdependence, captured by v_i, and the cost associated with interdependence, captured by c_i: For every $i \in I$ and $x \in S$, increasing ρ leads to an increase in $u_i(x, \rho)$. If this holds, then Theorem 9 and the first part of Theorem 10 will go through. Similarly, if increasing ρ leads to a decrease in $u_i(x, \rho)$, then Theorem 11 goes through.

Finally, whether or not the second part of Theorem 10 holds depends on more structural assumptions about the dependence of v_i on ρ. In particular, lowering ρ now has two opposing direct effects: decreasing c_i as before, but now also decreasing v_i. This implies that the direct effect of the privacy enhancement is no longer as beneficial.

7.4 Convex $\hat{c}_i$

Throughout the paper we assumed that the cost function c_i satisfies decreasing differences in $(x_i; x_{-i})$. This assumption is the driving force behind Theorem 3, that the extremal equilibria are increasing in ρ. Furthermore, it is straightforward to see that if c_i instead satisfies *increasing* differences, then the conclusion of Theorem 3 would be reversed: the extremal equilibria would be decreasing in ρ.

Now, recall that the severity of our assumption depends on the structure of the informational interdependencies. In particular, under overlapping personal information it is always satisfied, whereas under additive personal information it is satisfied only if $\hat{c}_i$ is concave. So what can be said when $\hat{c}_i$ is convex?

One interesting result is that lowering ρ can have opposite effects on the volume of information sharing, depending on the structure of information. In particular, if information is overlapping then lowering ρ will *decrease* the volume of sharing, as per Theorem 3. In contrast, however, if information is additive then a convex $\hat{c}_i$ implies that c_i satisfies increasing differences in $(x_i; x_{-i})$, and so lowering ρ leads to an *increase* in the volume of sharing.

8 Conclusion

This paper presents a new model of information sharing subject to privacy concerns, whose focus is on the quantitative aspect of the information shared. Using this model, I analyzed two distinct privacy enhancements—lowering the probability λ of information leakage and misuse, and lowering the level ρ of informational interdependencies amongst users—and showed that they can have opposite effects. The first enhancement is welfare improving when network effects dominate interdependencies, and the second when interdependencies dominate network effects. However, it is also possible that a privacy-diminishing policy is beneficial to all parties: for example, *increasing* ρ is beneficial to the network owner, as it increases the volume of information shared, and at the same time it may also improve user welfare (for example, Theorem 9 illustrates such a case for concave, network-symmetric games).

In the paper I analyzed the welfare of users, and one factor that impacts the utility of the network owner (the volume of information shared). However, a full welfare analysis might also take into account the third party to whom information has been leaked. Appendix C contains a preliminary version of such an analysis, with the main conclusion that, irrespective of network effects and informational interdependencies, the impact of lowering λ on the third party is ambiguous.

Appendix

A Aggregation Function

In this section we provide some formal justification for the assumption that the aggregation function $\mathcal{A}^i$ satisfies decreasing differences: that for every $x' > x$ it holds that

$$\mathcal{A}^i\left(x'_i, \rho x'_{-i}\right) - \mathcal{A}^i\left(x'_i, \rho x_{-i}\right) \leq \mathcal{A}^i\left(x_i, \rho x'_{-i}\right) - \mathcal{A}^i\left(x_i, \rho x_{-i}\right).$$

For the formal justification we will consider a variant in which there are finitely many pieces of information about each user, different users may hold different pieces of information, and users reveal information in some a priori fixed order. We will see that in this setup the decreasing differences assumption holds, and then one may view our model as the limit case that has a continuum of pieces of information.

For simplicity, assume there are two users, i and j. Consider user i, and suppose there are m total pieces of information about i, and denote them by $M = \{1, \ldots, m\}$. Let $P \subseteq M$ and $Q \subseteq M$ be the pieces of information held by i and j, respectively, about i. Finally, fix orders on the elements of P and Q, denoted by $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$. The interpretation is that if user i reveals k pieces of information then they are $p_1, \ldots, p_k$. Similarly, if user j reveals k pieces of information then they are $q_1, \ldots, q_k$. Finally, if i reveals k pieces of information and j reveals ℓ pieces of information, then the total number of pieces of information revealed about i is $\mathcal{A}^i(\{p_1, \ldots, p_k\}, \{q_1, \ldots, q_\ell\}) = |\{p_1, \ldots, p_k\} \cup \{q_1, \ldots, q_\ell\}|$.

Fix $(k', \ell') \gg (k, \ell)$. Then decreasing differences of $\mathcal{A}^i$ amounts to

$$\begin{aligned} &|\{p_1, \ldots, p_{k'}\} \cup \{q_1, \ldots, q_{\ell'}\}| - |\{p_1, \ldots, p_{k'}\} \cup \{q_1, \ldots, q_\ell\}| \\ &\quad \leq |\{p_1, \ldots, p_k\} \cup \{q_1, \ldots, q_{\ell'}\}| - |\{p_1, \ldots, p_k\} \cup \{q_1, \ldots, q_\ell\}|. \end{aligned}$$

This inequality is equivalent to

$$\begin{aligned} &\ell' - \ell - |\{p_1, \ldots, p_{k'}\} \cap \{q_{\ell+1}, \ldots, q_{\ell'}\}| \leq \ell' - \ell - |\{p_1, \ldots, p_k\} \cap \{q_{\ell+1}, \ldots, q_{\ell'}\}| \\ &\quad \Leftrightarrow |\{p_1, \ldots, p_{k'}\} \cap \{q_{\ell+1}, \ldots, q_{\ell'}\}| \geq |\{p_1, \ldots, p_k\} \cap \{q_{\ell+1}, \ldots, q_{\ell'}\}|, \end{aligned}$$

which follows from the assumption that $k' \geq k$.

B Counter-Examples

Example 3 Suppose $I = \{1, 2\}$, and fix $\rho_{12} = \rho_{21} = \rho$. Let $v_1(x) = \log(1 + x_1)$, $c_1(x) = C \cdot (\max\{x_1, \rho x_2\})$ for some $C > 0$, $v_2(x) = 0$, and $c_2(x) = \max\{x_2, \rho x_1\}$. Since $v_2(x) = 0$, user 2 would like to minimize the amount of information shared about him (by himself and by user 1). However, in an interior solution, user 1 does share some information. In particular, for any $\lambda \in (1/(2C), 1/C)$, the optimal amount of information shared by user 1 satisfies $v_1'(x) = \lambda \cdot c_1'(x)$, yielding the solution

$$x_1 = \frac{1}{\lambda C} - 1.$$

Given that user 1 shares x_1, the amount shared by user 2 is $x_2 \in [0, \rho x_1]$. The utility of user 2 in equilibrium (for any optimal choice of x_2) is thus

$$u_2(x, \lambda) = -\lambda\rho\left(\frac{1}{\lambda C} - 1\right) = \lambda\rho - \frac{\rho}{C}.$$

Now, fix any $\lambda, \lambda' \in (1/(2C), 1/C)$ with $\lambda < \lambda'$, and let x and x' be the maximal (minimal) equilibria of $\Gamma(\lambda)$ and $\Gamma(\lambda')$, respectively. By Theorem 6, there must be a user whose utility is higher in $\Gamma(\lambda)$, and, as we will see momentarily, this user must be user 1. The utility of user 2, however, decreases as λ decreases:

$$u_2(x, \lambda) = \lambda\rho - \frac{\rho}{C} < \lambda'\rho - \frac{\rho}{C} = u_2(x', \lambda').$$

Example 4 Suppose $I = \{1, 2\}$, and fix $\rho_{12} = \rho_{21} = \rho$. For both users i let $v_i(x) = x_i$. Denoting by $\overline{z}_i = \max\{x_i, \rho x_j\}$, let $c_i(x_i, \rho x_j) = 0$ if $\overline{z}_i \leq .5$ and $c_i(x_i, \rho x_j) = \sin(3\overline{z}_i - 1.5)$ otherwise. Note that both v_i and c_i are increasing, and satisfy the increasing/decreasing differences assumptions.

Fix $\rho = 0.8$, and consider the leakage probabilities $\lambda = 1$ and $\lambda' = 0.9$. Consider the symmetric profiles (x, x), which are potential symmetric equilibria. Figure 1 plots the utilities of a user associated with such profiles (with x on the horizontal axis). For both λ and λ', there are two potential equilibria: one at $(0.5, 0.5)$ and one at $(1, 1)$. Under λ only $(0.5, 0.5)$ is an equilibrium, yielding a utility of 0.5 to each user. The profile $(1, 1)$ is not an equilibrium here, since each user can gain by a unilateral deviation to 0.8. Under λ', however, both $(0.5, 0.5)$ and $(1, 1)$ are equilibria. At the

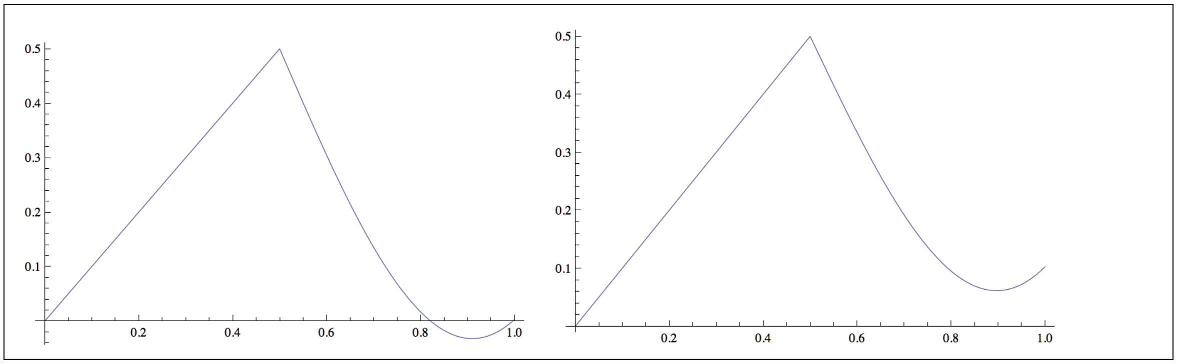

Figure 1: Utilities of users using strategies (x, x) under λ and λ'.

maximal equilibrium the utility of each user is $1 - 0.9 \cdot \sin(3 - 1.5) \approx 0.1$. Thus, lowering the probability of leakage from 1 to 0.9 yields lower welfare at the maximal equilibrium.[12]

Example 5 Suppose $I = \{1, 2\}$, and fix $\rho_{12} = \rho_{21} = \rho$. Let $v_1(x) = \log(1 + x_1) + \log(1 + x_2)$, $c_1(x) = C \cdot (\max\{x_1, \rho x_2\})$ for some $C > 0$, $v_2(x) = 0$, and $c_2(x) = \max\{x_2, \rho x_1\}$. Since $v_2(x) = 0$, user 2 would like to minimize the amount of information shared about him (by himself and by user 1). However, in an interior solution, user 1 does share some information. In particular, for any $\lambda \in (1/(2C), 1/C)$, the optimal amount of information shared by user 1 satisfies $v_i'(x) = \lambda \cdot c_i'(x)$, yielding the solution

$$x_1 = \frac{1}{\lambda C} - 1.$$

Given that user 1 shares x_1, the amount shared by user 2 is $x_2 \in [0, \rho x_1]$. Let us focus on the maximal equilibrium, in which $x_2 = \rho x_1$.

Observe that user 1's utility in this equilibrium is

$$u_1(x) = \log\left(1 + \frac{1}{\lambda C} - 1\right) + \log\left(1 + \rho\left(\frac{1}{\lambda C} - 1\right)\right),$$

which is strictly increasing with ρ. In other words, lower ρ implies strictly lower utility.

Furthermore, the utility of user 2 in this equilibrium is

$$u_2(x, \lambda) = -b\rho\left(\frac{1}{\lambda C} - 1\right) = \lambda\rho - \frac{\rho}{C}.$$

[12] Note, however, that welfare in the minimal equilibrium increases as λ decreases.

Note that this is strictly decreasing with ρ, since $\lambda < 1/C$, and so a lower ρ yields strictly higher utility for user 2.

C Hacker

In order to analyze the full effect on welfare, we may wish to consider the welfare of the party to whom the information is leaked. We refer to that party here as the hacker. The utility of the hacker depends on the probability of information leakage, as well as amount of information obtained. Of course, if the hacker is sufficiently risk averse, then lowering the probability of leakage λ (say, by strengthening security and thus increasing the cost to the hacker of obtaining the information) will be harmful to him. However, what if the hacker is not so risk averse? In an extreme case, suppose the hacker is risk neutral, and his utility is an increasing function of the expected amount of information leaked, namely $\lambda \cdot (\sum_i x_i)$. Does this quantity increase or decrease as λ decreases?

Suppose $c_i(x) = c_i(x_i)$, and that there are no network effects. In this case, user i must simply balance the benefit $v_i(x_i)$ with the cost $\lambda \cdot c_i(x_i)$. Denote by x_i^* the optimal amount of information shared by user i. We will provide two examples: in one, decreasing λ increases λx_i^*, and in the other it decreases it. Hence, this privacy enhancement has an ambiguous effect on the hacker *even without network effects and interdependent information.*

Example 6 Consider the game $\Gamma(\lambda, \rho)$ defined as follows: fix constants $k \in (0, 1)$ and $C > 0$, and for every user $i \in I$ set $v_i(x_i) = x_i^k$ and $c_i(x_i) = C \cdot x_i$. If the solution is interior then the optimal level of sharing satisfies $v_i'(x_i) = (\lambda C x_i)'$, which implies that $k x_i^{k-1} = \lambda C$.

Now, fix some $\lambda < \lambda'$, and let x_i and x_i' be the corresponding optimal levels of sharing. Suppose the solutions are interior—namely, that $x_i, x_i' \in (0, 1)$—which holds whenever $k < \lambda C$. Observe that

$$\begin{aligned}
&\lambda x_i > b' x_i' \\
&\quad \Leftrightarrow \frac{x_i \cdot v_i'(x_i)}{C} > \frac{x_i' \cdot v_i'(x_i')}{C} \\
&\quad \Leftrightarrow x_i \cdot v_i'(x_i) > x_i' \cdot v_i'(x_i') \\
&\quad \Leftrightarrow \frac{x_i}{x_i'} > \frac{v_i'(x_i')}{v_i'(x_i)}.
\end{aligned}$$

Furthermore,

$$\frac{v_i'(x_i')}{v_i'(x_i)} = \left(\frac{x_i'}{x_i}\right)^{k-1} = \left(\frac{x_i}{x_i'}\right)^{1-k} < \frac{x_i}{x_i'}$$

since $x_i > x_i'$ by Theorem 2. Thus, $\lambda x_i > \lambda' x_i'$, and so a decrease in λ leads to higher λx_i.

We now consider an example in which a decrease in λ leads to a corresponding decrease in λx_i.

Example 7 Consider the game $\Gamma(\lambda, \rho)$, which will use numbers $0 < a_1 < a_2 < 1$ and $m_1 > m_2 > m_3 > 0$ to be chosen later. Set

$$v_i(x_i) = \begin{cases} m_1 x_i & \text{if } x_i \in [0, a_1], \\ m_2(x_i - a_1) + m_1 a_1 & \text{if } x_i \in (a_1, a_2], \\ m_3(x_i - a_2) + m_1 a_1 + m_2(a_2 - a_1) & \text{if } x_i \in (a_2, 1]. \end{cases}$$

Thus, v_i is concave and piecewise linear. Additionally, set $c_i(x) = C x_i$ for some $C > 0$.

Now, for any λ, in an interior solution $v_i'(x_i) = \lambda \cdot c_i'(x_i)$. Choose $\lambda = m_3/C$ and $b' = m_2/C$, and note that $\lambda < \lambda'$. Furthermore, the solutions of $\Gamma(\lambda, \rho)$ consist of all $x_i \in [a_2, 1]$, and those of $\Gamma(\lambda', \rho)$ consist of all $x_i' \in [a_1, a_2]$. Thus, for any choice of solutions x_i and x_i' it holds that

$$\lambda x_i \leq \lambda = \frac{m_3}{C}$$

and

$$\lambda' x_i' \geq \lambda' a_1 = \frac{m_2}{C} \cdot a_1.$$

We thus obtain that

$$\lambda x_i < \lambda' x_i'$$

whenever

$$\frac{m_3}{C} < \frac{m_2}{C} \cdot a_1,$$

and this last inequality holds for sufficiently small $m_3 > 0$.

References

Acquisti, A., John, L. K. and Loewenstein, G. (2012). The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, **49** 160–174.

Acquisti, A., Taylor, C. R. and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, **54** 442–492.

Adjerid, I., Acquisti, A., Telang, R., Padman, R. and Adler-Milstein, J. (2015). The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, **62** 1042–1063.

Anderson, R. and Moore, T. (2006). The economics of information security. *Science*, **314** 610–613.

Biczók, G. and Chia, P. H. (2013). Interdependent privacy: Let me share your data. In *Financial Cryptography and Data Security*. Springer, 338–353.

boyd, d. (2012). Networked privacy. *Surveillance & Society*, **10** 348.

Burke, M., Marlow, C. and Lento, T. (2009). Feed me: motivating newcomer contribution in social network sites. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 945–954.

Calo, R. (2011). Boundaries of privacy harm, the. *Ind. LJ*, **86** 1131.

Congress, U. (1993). Office of technology assessment: Protecting privacy in computerized medical information. *Washington, DC: US Govt. Printing Office* 44.

Consumer Reports (2012). Facebook & your privacy. [accessed July 12, 2016], URL `http://www.consumerreports.org/cro/magazine/2012/06/facebook-your-privacy/index.htm`.

Cummings, R., Ligett, K., Pai, M. M. and Roth, A. (2016). The strange case of privacy in equilibrium models. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 659–659.

DAUGHETY, A. F. and REINGANUM, J. F. (2010). Public goods, social pressure, and the choice between privacy and publicity. *American Economic Journal: Microeconomics*, **2** 191–221.

DWORK, C. (2008). Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*. Springer, 1–19.

DZIUDA, W. and GRADWOHL, R. (2015). Achieving cooperation under privacy concerns. *American Economic Journal: Microeconomics*, **7** 142–173.

FACEBOOK (2016). Number of monthly active Facebook users worldwide as of 1st quarter 2016 (in millions). In Statista – The Statistics Portal. Accessed July 1, 2016, URL `http://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/`.

GOLDFARB, A. and TUCKER, C. (2012). Privacy and innovation. *Innovation policy and the economy*, **12** 65–90.

GRADWOHL, R. (forthcoming). Voting in the limelight. *Economic Theory. Extended abstract in Proceedings of the 2017 ACM Conference on Economics and Computation.*

GRADWOHL, R. and SMORODINSKY, R. (2017). Perception games and privacy. *Games and Economic Behavior*, **104** 293–308.

GRIMMELMANN, J. (2009). Saving facebook. *Iowa L. Rev.*, **94** 1137–1793.

GROSS, R. and ACQUISTI, A. (2005). Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on privacy in the electronic society*. ACM, 71–80.

HUMBERT, M., AYDAY, E., HUBAUX, J.-P. and TELENTI, A. (2015). On non-cooperative genomic privacy. In *Financial Cryptography and Data Security*. Springer, 407–426.

JOHNSON, M., EGELMAN, S. and BELLOVIN, S. M. (2012). Facebook and privacy: it's complicated. In *Proceedings of the eighth symposium on usable privacy and security*. ACM, 9.

Kasiviswanathan, S. P. and Smith, A. (2008). On the 'semantics' of differential privacy: A Bayesian formulation. *arXiv preprint arXiv:0803.3946*.

Kunreuther, H. and Heal, G. (2003). Interdependent security. *Journal of risk and uncertainty*, **26** 231–249.

Liang, K., Liu, J. K., Lu, R. and Wong, D. S. (2015). Privacy concerns for photo sharing in online social networks. *Internet Computing, IEEE*, **19** 58–63.

McGuire, A. L., Fisher, R., Cusenza, P., Hudson, K., Rothstein, M. A., McGraw, D., Matteson, S., Glaser, J. and Henley, D. E. (2008). Confidentiality, privacy, and security of genetic and genomic test information in electronic health records: points to consider. *Genetics in Medicine*, **10** 495–499.

McPherson, M., Smith-Lovin, L. and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, **27** 415–444.

Milberg, S. J., Burke, S. J., Smith, H. J. and Kallman, E. A. (1995). Values, personal information privacy, and regulatory approaches. *Communications of the ACM*, **38** 65–74.

Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, **58** 1255–1277.

Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, **62** 157–80.

Moon, Y. (2000). Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of consumer research*, **26** 323–339.

Norberg, P. A., Horne, D. R. and Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, **41** 100–126.

O'Donnell, A. (2016). The dangers of Facebook oversharing. [accessed July 13, 2016], URL `http://netsecurity.about.com/od/securityadvisorie1/a/The-Dangers-Of-Facebook-Oversharing.htm`.

Olteanu, A.-M., Huguenin, K., Shokri, R. and Hubaux, J.-P. (2014). Quantifying the effect of co-location information on location privacy. In *Privacy Enhancing Technologies*. Springer, 184–203.

Pew Research Center (2014). Public perceptions of privacy and security in the post-Snowden era. *Pew Research Internet Project*. URL `http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/`.

Posner, R. A. (1981). The economics of privacy. *The American Economic Review* 405–409.

Pu, Y. and Grossklags, J. (2014). An economic model and simulation results of app adoption decisions on networks with interdependent privacy consequences. In *Decision and Game Theory for Security*. Springer, 246–265.

Pu, Y. and Grossklags, J. (2015). Towards a model on the factors influencing social app users valuation of interdependent privacy. *Proceedings on Privacy Enhancing Technologies*, **2016** 61–81.

Rainie, L. and Duggan, M. (2016). Privacy and information sharing. *Pew Research Center, Jan*, **14**. URL `http://www.pewinternet.org/2016/01/14/2016/Privacy-and-Information-Sharing/`.

Sarigol, E., Garcia, D. and Schweitzer, F. (2014). Online privacy as a collective phenomenon. In *Proceedings of the second ACM conference on online social networks*. ACM, 95–106.

Schwartz, P. M. (1997). Privacy and the economics of personal health care information. *Tex. L. Rev.*, **76** 1.

Stigler, G. J. (1980). An introduction to privacy in economics and politics. *The Journal of Legal Studies* 623–644.

Symeonidis, I., Shirazi, F., Biczók, G., Pérez-Solà, C. and Preneel, B. (2016). Collateral damage of Facebook apps: Friends, providers, and privacy interdependence. In *IFIP International Information Security and Privacy Conference*. Springer, 194–208.

Turow, J., King, J., Hoofnagle, C. J., Bleakley, A. and Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it. *Available at SSRN 1478214*.

Wilson, R. E., Gosling, S. D. and Graham, L. T. (2012). A review of Facebook research in the social sciences. *Perspectives on psychological science*, **7** 203–220.

Zhou, L. (1994). The set of nash equilibria of a supermodular game is a complete lattice. *Games and economic behavior*, **7** 295–300.