Peer-produced Privacy Protection

A Common-pool Approach

Vaibhav Garg
Department of Computer Science
Drexel University
Philadelphia, USA
gargv@drexel.edu

Sameer Patil
Helsinki Institute of Information Technology
Aalto University
Aalto, Finland
patil@indiana.edu

Apu Kapadia
School of Informatics and Computing
Indiana University
Bloomington, USA
kapadia@indiana.edu

L. Jean Camp
School of Informatics and Computing
Indiana University
Bloomington, USA
ljcamp@indiana.edu

Abstract— **Privacy risks have been addressed through technical solutions such as privacy-enhancing technologies (PETs) as well as regulatory measures including Do Not Track. These approaches are inherently limited as they are grounded in the paradigm of a rational end user who can determine, articulate, and manage consistent privacy preferences. This implies that self-serving efforts to implement individual privacy preferences lead to socially optimal outcomes with regard to information sharing. Consequently, solutions to specific risks are developed, and even mandated, without effective reduction in the overall harm of privacy breaches. We present a systematic framework to examine the limitations of current technical and policy solutions. To address the shortcomings of existing privacy solutions, we argue for considering information sharing to be transactions *within a community*. Outcomes of privacy management can be improved at a lower overall cost if peers, as a *community,* are empowered by appropriate technical and policy mechanisms. Designing for a community requires encouraging dialogue, enabling transparency, and supporting enforcement of community norms. In this paper we show how peer production of privacy is possible through PETs that are grounded in the notion of information as a common-pool resource and community governance.**

Keywords-privacy, computer supported collaborative work, economics.

I. INTRODUCTION

Technological advances in the past two decades, in software as well as hardware, have resulted in a great transformation in how we interact with other parties socially and commercially. While powerful computing devices, coupled with universally available Internet access, offer great benefits and conveniences, they also present a wide range of risks and problems regarding privacy.

Typically, such privacy issues and concerns are tackled via technology or regulation or some combination of the two. The technical approach involves designing Privacy-Enhancing Technologies (PETs) that counter threats to privacy posed by use of the underlying technology. For example, encryption and access control are PETs used to guard information from access by unauthorized parties. Regulation, on the other hand, utilizes measures, such as guidelines, policies, contracts, and laws. These are used to describe uses and applications of technology that are *not* permitted, even when they are technologically feasible. For instance, it is easily possible to run face-detection algorithms on any given photograph. However, Facebook has turned the feature off in Europe in response to recent European Union regulatory actions.

While these two approaches do mitigate and tackle various privacy aspects, one of their limitations is that they operate under the paradigm of individuals independently managing their own information. This paradigm involves several fundamental assumptions regarding privacy-related decision making of individuals:

- An individual has correct and complete information needed to make a decision. For instance, privacy (control) settings assume that people have necessary and sufficient *a priori* information to specify preferences that will apply to information sharing behaviors in the future.
- Individuals are able to articulate their privacy needs. For instance, specification of privacy preferences requires that a person is able to describe desires and needs that are often tacit and implicit, and thus difficult to describe explicitly.
- An individual knows about the existence of PETs and regulations regarding privacy. For instance, the onus and burden is often placed on privacy-conscious individuals, to learn whether or not it is possible for the system to serve their privacy needs and whether the system provides mechanisms or policies for this purpose.
- An individual understands how to manage the user interfaces and interactions for managing privacy. For instance, it is generally assumed that users operate with correct understanding and mental models of how the

Vaibhav Garg, Sameer Patil , Apu Kapadia, and L Jean Camp, "Peer-produced Privacy Protection: A Common-pool Approach", The IEEE International Symposium on Technology and Society (ISTAS) (Toronto, ON) 27-29 July 2013.

underlying system operates and how various privacy options affect this operation.

- Individuals are able to translate their articulated privacy requirements into preferences that can be specified via the available interface and interaction mechanisms. For instance, users are tasked with translating their privacy needs as stated in natural language into a formal "specification" that a system can parse, process, and enforce.
- An individual is able to keep track of changes in privacy desires based on changes in contextual factors, and, in turn, able to update specified privacy preferences such that they are always contextually appropriate. For instance, as privacy requirements change due to a change in context, an individual must manually update privacy preferences to match the privacy requirements of the new context.
- Individuals always make privacy decisions that achieve their privacy desires in an optimal manner. For instance, privacy-related actions and behaviors of people are assumed to achieve their stated privacy goals.

Empirical research, however, has shown that these assumptions are often not met in practice. For example, individuals frequently make decisions with incomplete and/or inaccurate information [2, 24], may not fully understand how technology could affect privacy [33], cannot completely and accurately describe their privacy needs [31], exhibit behavior inconsistent with their own stated privacy concerns [42], do not know about or utilize privacy management mechanisms [25, 12, 1], are confused by interfaces for specifying privacy preferences [35, 43], or find it difficult and burdensome to adjust preferences according to the context [16].

A possible approach to individual privacy is to develop techniques and solutions that attempt to eliminate, or minimize, the discrepancy between the ideal and practice. However, even if this ideal were achieved, the paradigm of an individual making decisions about privacy suffers from two additional shortcomings.

First, the individual decision-making paradigm assumes that individually optimal privacy decisions lead to privacy outcomes that are socially optimal (and desirable) for the community. As demonstrated by the prisoner's dilemma [45] and the tragedy of the commons [44], individually rational decisions can lead to Nash equilibrium that are suboptimal from a communal perspective[1]. For instance, in the case of privacy an individual may perceive the risk of privacy violation not worth the cost of privacy protection. However, an aggregate database is greater than the sum of its individual information components. Thus, the privacy risks of such aggregation of the information arguably would be greater than the sum of individual privacy violations. It is possible that the expected value of this aggregate privacy risk is greater than the aggregate cost. In hindsight it may then be rational to have invested in privacy protection.

Second, it ignores the role of the actions of others in affecting an individual's privacy, *regardless of what the individual chooses to do*. The social, professional, and business relationships that one maintains — online as well as offline — result in individuals' privacy being affected due to actions taken by those with whom they are connected. For instance, even when an individual chooses not to reveal her birthdate to a provider of an online social networking service (SNS), birthday greetings sent by one's friends via public (or private) messaging mechanisms of the SNS result in implicit disclosure of the birthdate to everyone (or at the very least the SNS provider).

We suggest that a fruitful way to address the shortcomings of focusing on individual privacy decision-making is by taking a *community*-based approach. It has already been noted that information sharing typically takes into account an imagined community [1]. In fact PETs have been criticized for not addressing trust among individuals in a community [21]. We address these limitations by considering privacy, not as a public or a private good, but a common-pool resource. Ostrom et al. [30] note that successful and sustainable community-based governance of such a shared (common-pool) resource is contingent on five conditions being met [14]:

1. Monitoring the resource must be cheap.
2. The community must have a mechanism to maintain the reputation of the users of the resource.
3. The rate of change of the resource should be relatively constant.
4. It must be possible to exclude individuals from using the resource.
5. Community members support the enforcement of community norms and therefore the monitoring required for effective enforcement.

These five conditions suggest that design for community governance requires encouraging dialogue between community members, enabling transparency of information flows, and supporting enforcement of community norms.

Ostrom et al. [30] acknowledge that in practice these five conditions are not met for any resource. However, institutional structures can be put in place to meet these requirements artificially without the construction of explicit property rights. This is also true of community management of privacy online, where technical measures can complement or substitute institutional (or policy) measures. In this paper we provide specific technical solutions that enable commons-based communal governance to achieve peer-produced privacy protection.

Section 2 begins by providing a description of Ostrom's framework. In Section 3 we discuss the limitations of the current approaches. Further, we show how Ostrom's

[1] Prisoner's Dilemma [45]: Assume there are two individuals who could go to jail for robbery. If neither individual agrees to turn state's evidence they both could be sentenced to *c* number of years; if both turn states' evidence the sentence increases to *3c* years; if only one divulges information that individual is not punished but the other's sentence is increased to *2c*. The socially optimal outcome for the two individuals here is contingent on neither individual confessing to the crime. However, individually rational decision would lead to an outcome where both individuals are worse off.

framework can be used for a systematic discovery of these limitations. Section 4 describes several technical solutions to peer-produced privacy protection grounded in Ostrom's notion of community-based governance of the information commons. Finally, Section 5 concludes with a discussion of future work.

II. A Common Pool Approach

For successful governance of the common-pool resource through local stakeholders, the five conditions discussed below should be met.

A. Resource monitoring should be inexpensive

This allows all the stakeholders to be aware of how the peers in their community are accessing and consuming the resource. In terms of privacy, this condition is rarely met. Even when information is shared willingly, it is almost impossible to observe the information flows post hoc. For example, even when Facebook controls are utilized adequately and information is shared with a specific person, there are limited, if any, options to know how frequently that information is being accessed by the specific individual.

Similarly, while Web sites differ in their privacy policies, there is little incentive for most of them to use it as a selling point. From a behavioral perspective, even a good-faith discussion of the privacy policy could create anxiety for the consumer and deter adoption [9]. Users may, for example, have higher privacy concerns when primed [18]. Similarly, technical tools that allow users to analyze how information about them is being collected and distributed by different websites are rarely available. Even when certain tools, such as Ghostery, are available they do not allow users to pool information in a manner that allows them to discriminate amongst websites based on privacy.

B. Maintaining the reputation of resource users is essential

The second requirement is that of social capital; i.e. those accessing the resource should have frequent face-to-face communication to establish trust. Face-to-face communication is not always possible on the Internet. However, technical solutions to gauge social capital are readily available, e.g. reputation systems. However, these have not been incorporated as a part of PETS. Existing solutions such as TRUSTe seals are not peer produced and suffer from incentive misalignment; i.e. their customers are websites and not individuals whose privacy must be protected [3], and thus often websites with such seals provide less stringent privacy protection than those without them [27]. Peer production of reputation eliminates the cost of hidden action, as those generating the reputation rankings are the ones who are interested in using such reputation.

Peer-produced reputation systems for websites are available for information security, e.g. the Web of Trust plugin. However, similar reputation regarding privacy policies and information collection/sharing behaviors of websites are not easily available [9] and are often expensive enough for the user to be rationally indifferent [26].

C. The resource rate of change should be relatively constant

A third condition requires moderate rates of change, i.e. the resource itself, those using the resource, as well as the technological, social, and economic conditions should not change too aggressively. It is hard to argue whether or not this is true for the information commons. It is relatively easier to examine whether this condition is relevant for privacy online.

Moderate rates of change are required for physical goods to enable monitoring and reputation. Arguably, if the number of individuals using a fishery changes frequently, social capital would be hard to compute. Similarly, if those using the resource change constantly it would be difficult to monitor their usage of the resource. These can, however, be addressed by increasing the cost of creating a new identity or the opportunity cost of losing an old one.

Moderate rates of change of the resource for physical goods are also needed for reasons of sustainability. If the rate of consumption of a resource were higher than the rate with which it replenishes, then it would no longer be sustainable. This notion of sustainability is not relevant for information online.

However, in the case of privacy the cost of implementing community norms could make peer governance unsustainable. Thus, it is important to consider the cost of implementing community norms vs. that of implementing individual preferences. It has been argued that individual end-users have a limited security budget [5]. This would arguably be true for privacy as well, i.e. users would have a limited amount of resources that they would be willing to spend on implementing their privacy preferences as well as community privacy norms. In peer-produced privacy protection individuals select their peers, and therefore the nature of community norms. Thus, compliance would arguably be higher and privacy "sustainable".

D. Excluding individuals from resource use must be possible

The fourth requirement refers to exclusion, i.e. it should be possible at a relatively low cost to exclude entities from the resource. For physical resources this exclusion may be binary. For example, while members of a village on the riverbank are allowed to fish in the river, those not from the village do not have similar rights. Online, however, the choices are rarely binary. Given the contextual nature of privacy, exclusion becomes more problematic. On social networks an individual may want professional colleagues to get status updates about new publications, but not about holidays. Location privacy is particularly problematic. Even for individuals with whom one is willing to share location information, one might be concerned if that information is accessed too frequently [41].

Exclusion has been the major focus of most PETs and privacy policy. Encryption, for example, excludes everyone other than those with access to the appropriate keys to access the information. Privacy controls on Facebook similarly prevent those without appropriate permissions to access the individual's complete profile. On the policy front, efforts such as Do Not Track are initiatives that allow consumers to prevent websites from collecting information that can be used for tracking online behavior.

TABLE I. PROPERTIES OF GOODS

	Excludable	*Non-Excludable*
Rivalrous	***Private***	***Common-pool***
Non-Rivalrous	***Club/Toll***	***Public***

E. Community members must support the enforcement of community norms

The final requirement is that of enforcement, i.e. community members can identify when norms are not followed and then punish defectors through exclusion or other penalties. This is particularly difficult to do for privacy, especially with current controls. Privacy is contextual [29], however privacy preferences are typically set out of context. For example, permissions to access location information are typically given to mobile applications ahead of time, before the user can understand whether an access may reveal information that they desire to be hidden.

Simultaneously, the visibility of broken norms is low. Arguably, when norms are broken peers in a community can and do create pressure that leads to compliance. For example, Facebook, on acquiring Instagram, changed the policy on intellectual property. This was seen as Facebook breaking the foundational norm of the community, i.e. the photographs were not for commercial use. The resulting backlash from the existing Instagram community created enough negative publicity for Facebook to repeal the policy change.

Successful governance of the information commons by peers to prevent privacy violations requires that all these five requirements be met. Often in practice existing technologies only provide for a subset of these requirements. It is important to note that these five requirements are often interdependent. For example, even if enforcement mechanisms were available, if monitoring is either not possible or prohibitively expensive, enforcement would be unlikely to happen in practice. Thus, a partial fulfillment of these requirements with regards to the information commons may create the illusion of risk reduction without an actual decrease in the overall harm of privacy breaches.

III. CURRENT APPROACHES AND THEIR LIMITATIONS

In economics there are four kinds of goods: public, private, common-pool, and club. These are differentiated based on whether a good is excludable and/or rivalrous; Table 1. If individual entities can be prevented from consuming a resource it is excludable. If a good can be subtracted, i.e. one individual's consumption of a good leaves less of the good for others, the good is rivalrous. These properties are often mutable.

Canonical wisdom states that systems are sustainable, but only under a paradigm that considers system resources to be either a public or a private good. The notion of privacy as confidentiality considers information as a private good. Thus, the solutions have focused on encryption technologies. When information is considered a private good, the assumption is that information about a specific person is only relevant to them; information sharing by that person then puts only that specific individual at risk.

However, this assumption fails too easily in real life. For example, public records of genomic information about an individual are relevant to both the primary stakeholder and their relatives. In fact property rights over certain kinds of information can be hard to assign. For example, when a group photograph is taken at a friend's party, who among the group should have the rights to post the picture online? Merely providing the rights to the person who takes the photograph or the owner of the camera would lead to frequent privacy violations.

A second argument considers information to be public good. Posner, in fact, argues that given that information makes markets more efficient only those involved in unsavory activities would be invested in hiding their information, i.e., privacy is valuable only to those who have something to hide [38]. Even Posner, however, did not assume that information should be freely available. In fact for markets to be truly efficient the individual whose information is being used should be reimbursed for their resource [39]. Therein lies the idea of privacy as control; information sharing is enabled only when both parties involved have higher individual utilities post transaction. Information as a public good paradigm is then used to develop policy solutions such as DNT; the individual can choose to be tracked if the transaction is mutually beneficial or that privacy loss is adequately reimbursed by the benefits of behavioral advertising.

On the technical side control is being enabled by Privacy Enhancing Technologies. These approaches assume a rational end-user who can implement their privacy preferences, so that the user can control their information flows. This limited view of rationality is problematic, as individuals have limited control of their information flows. For example, I may choose not to share my date of birth with Facebook. However, if my friends choose to wish me Happy Birthday on the specific day there is little I can do, without incurring prohibitive transaction costs. Simultaneously, even when information is shared voluntarily with specific individuals the true "exposure" of that information is not known [41].

These and other criticisms of PETs have been made in prior literature [21]. Here we discuss two additional limitations of existing paradigms. First, solutions to privacy risks target individuals. The narrow perspective of the rationality assumption then presumes that individually rational decisions would lead to socially optimal outcomes. This is often not true as privacy risks are not individual risks but rather aggregate risks, e.g. behavioral advertising. Even when a user chooses not to be a part of a database, not only do they reveal information from refusing to participate, other inferences can still be made of them from aggregate inputs of other participants.

A second limitation of the rational actor paradigm of the end-user is that of costs. A narrow perspective of rationality may assume that individually self-serving implementations of PETs would produce socially optimal (or even socially desirable) outcomes. Of course, individually rational decisions often lead to Nash equilibriums that are suboptimal, e.g.

prisoners dilemma and on a larger scale tragedy of the commons [22]. The problem is further compounded for privacy. For a single individual the perceived costs of implementing PETs may be high, while those of privacy infringement may be perceived to be low. However, an information database is greater than the sum of its parts. Thus, the aggregate privacy loss may be greater. Arguably, then, voluntary information disclosure can be described as a tragedy of the information commons.

The typical approach to addressing the tragedy of the commons has been through public or private interventions. Let's consider private interventions. Arguably, there are (academic) incentives for private parties to preserve the common-pool or community resources both offline and online. Offline it is in the interest of the private entity to sustain the fishery or forest for long-term gain. Online incentives to provide stronger PETs have been noted both from a rational choice [8] and behavioral perspective [10]. In practice, however, the destruction of natural resources owned by private entities is widely documented [6], while online Facebook has not provided stronger and, more importantly, usable and useful privacy controls. Similarly, providing privacy information through private entities such as TRUSTe seals only creates a perception of increased privacy, which arguably lowers protection through inadequate risk compensation. (In fact it has been noted that such perception could result in more risk taking behavior [27].)

Thus, a second approach that is being tried is public interventions through the Federal Trade Commission (FTC), e.g. Do Not Track. Offline such interventions, though well meaning, have had limited success. A key constraint is that external public officials have limited knowledge about the resource, such as fisheries, compared to the granular information available to locals who are invested in sustaining the ecosystem as it provides long-term employment. For example, public efforts have managed to preserve local forests while simultaneously destroying the diversity of flora and fauna due to limited knowledge about the ecosystem [4].

A third possibility is that of considering information as a commons that is best managed by those who would be adversely effected by a privacy breach. As noted in the previous section, five requirements must be met for such an approach to be successful: monitoring, reputation, moderate rates of change, exclusion, and enforcement. This five-dimensional framework can be used to identify the failures of current privacy solutions.

We can take the example of Facebook privacy controls. Currently, these controls do not allow monitoring. For example, it is not possible for individuals to ascertain if and when their information is being accessed by their peers. (On the other hand a different social network, Orkut, did provide this option. Specifically, users could opt-in to be able to view the last five individuals who visited their profile. However, this implied that when they visited someone else's profile that information would be available to the other individual.)

There is also no mechanism to establish reputation; users cannot identify if certain peers excessively or inappropriately tend to access their information or post information about them. It may be that a member of one's friend circle repeatedly post pictures in which they tag an individual. These pictures and related tagging may be undesirable to that individual. However, for every new picture the individual must remove the tag individually.

Instead, if there was a mechanism to assign reputations to peers, one could simply influence the corresponding individual's ranking. This would provide feedback to both the offending member, who would then be encouraged to change their behavior or risk having a poor reputation, and to other members of the peer group, who would then be aware of the offending individual's deviance from the expected norm.

In terms of "rates of change" Facebook privacy controls do see frequent change. Facebook often changes both the interface and the functionality of its privacy controls. While information sharing has become increasingly automated, so that Facebook can now automatically tag individuals, similar advances in privacy controls have not been made. However, in terms of peers, the turnover rate is dependent on the users themselves, and individuals can choose to increase or decrease the number of peers they are connected to at any time.

A limited form of exclusion is possible on Facebook. Users can choose to exclude other peers that they feel should not have access to their information. However, excluding Facebook itself is not an option. Some limitations due to network effects have been discussed previously in this paper. Another key issue is that of the "right to be forgotten" [40]. Facebook does not allow users to delete their information if they so choose. Thus, Facebook can only be excluded from future data and not past data. (A different social network, Google Plus, allows users to remove their data from Google servers if they choose.)

Enforcement is possible to a certain degree on Facebook; it is possible to be friends with a certain individual as well as increase or decrease their privilege based on whether they appear to be following a certain community norm. Unfortunately, as noted before, due to lack of transparency it is unclear if someone is breaking a community norm. Ideally, it would be possible to implement a community norm itself as a privacy control. Then those that follow the norm would maintain their access. However, those who defect would suffer lower access not only to the profile of the individual who discovers the violation but for everyone in the peer community.

Thus, Ostrom's framework can be used to examine the limitations of current solutions to privacy, especially as those solutions impinge on community-based governance of the information commons. In the next section we present existing research that argues for a self-governance of the information commons and the technical solutions that enable such communities.

IV. Peer-based Approaches in Technical Implementation

Privacy functionalities in technologies, in both the consumer and the interpersonal domain, have utilized approaches that include actions of parties other than the individual whose privacy is to be managed. These other parties, peers from the community, aid privacy management in a variety of capacities. Here we discuss a model of how privacy can be instantiated by complementary design of privacy and policy.

First, there is the notion of *norms* often instantiated as self-regulation through a privacy culture. Although technology often enables violations of privacy, the social context in which the technology is embedded can act as a strong counteracting force due to prevailing norms about acceptable behavior. Due to the social costs of breaching these norms, peers in a trusted community are generally relied upon to regulate their own behavior such that it does not violate the privacy of others. As Dourish [15] notes regarding such ‘cultural’ models of privacy protection in the Media Space at Xerox PARC, “[w]ithin a small community, the result is a stable situation, comfortable and acceptable to participants, without direct need for a more technological solution.” This assertion suggests that the approach defined by Ostrom could apply to privacy management, even in a domain where the levels of expertise are quite high.

Social norms without technological solutions are unfortunately not scalable. As communities become larger the cost of monitoring through non-technological means becomes prohibitively expensive. Monitoring of privacy violations must thus be facilitated through “exposure feedback” within a community of individuals [41, 32, 17]. *Information exposure* refers to actual accesses that occur within the permissible bounds as articulated through privacy settings. Exposure feedback serves as an indicator of how entities in a community — individuals, businesses, or the government — are consuming information. Such mechanisms make the *cost of monitoring* privacy low for a community through transparency, and allow community members to assess whether such uses of information violate community norms. Such exposure mechanisms then make it easy to assess the *reputation* of peers based on the monitoring of their actions. Community members can then act as “guardians”, and respond to undesirable exposure patterns for individuals in the community.

For example, Bob may realize that Alice’s status message on Facebook is attracting more attention than was (probably) anticipated, and can temporarily restrict access to that information (until Alice is able to assess her exposure). In this scenario it is important to exclude undesirable people from the community; indeed, excludability in the community of peers is the difference between privacy and censorship in the presence of guardians. Exposure norms can be effective if there is control over data diffusion and the ability to remove information.

Research has also shown that those who are less technically savvy often rely on technical assistance from their social network of family, friends, and colleagues [23, 36, 37, 11]. Such assistance covers privacy-affecting matters such as configuring home networks, protecting against spyware, and setting privacy preferences. All three of these examples are domains where automation can enable effective peer sharing by embedding technological expertise. Individuals bring their own context via social networks and automated settings limiting sharing empowers the social network. An individual can use information regarding aggregate community choices as guidance for making decisions. Such “social navigation” [13] approaches have been applied to inform decisions about cookie management [20], firewall rules [19], phishing sites [28], information access policies for Facebook applications [7], and privacy preference settings in Instant Messengers [34].

Here too it is important that the consequences of shared privacy settings are communicated with the community. Transparency using an exposure feedback or harm approach can convey how these settings have affected the privacy of the community and reputation mechanisms can track which members of the community provide more effective controls, resulting in the exclusion of community members who utilize settings not desirable to the group.

In all of the above cases, a person relies upon and/or utilizes the actions of others from the community to enable more enhanced and effective protection of his or her privacy.

V. Future Work and Conclusions

Designing for a community that considers information sharing as a norm and PETs as covenants [13], with and without swords, is a profoundly different paradigm than one where individuals manage their own information. A community approach to privacy assumes that there is a group of individuals who share the risk of information sharing and that no one individual can assert complete privacy. Thus, peer production recognizes information and therefore privacy as a common-pool resource.

Understanding privacy as a community good, but one that requires transparency and excludability for protection, has implications for regulation as well as technical design. The FTC has played a crucial role in ensuring that organizations comply with their privacy policies towards individuals. A natural extension of the framework for peer production would require that individuals be more empowered to limit and even remove information. Much of the technology for information control is extant; for example, facial recognition is used to suggest tags on Facebook. However, it is not used to request permission from one person about their image before another can post it. There is no technological instantiation, only the norms discussed above, and this has proven incapable of preventing promiscuous posting of problematic photographs. A peer-centered approach would allow individuals to advise each other in their interactions, and remove their digital tracks. This would be an expansion of the current transactional approach to privacy, one that would continue to focus on transparency, and expand it to include excludability, construction of community, data removal, and harm.

Privacy functionalities in technologies, in the consumer as well as the interpersonal domain, must utilize approaches that include actions of parties other than the individual whose privacy is to be managed. These other parties, these peers from the community, aid privacy management in a variety of capacities. Future work would examine policy instantiations to enable community governance for sustainable privacy. Further, these investigations would also consider technology and policy as both supplements and complements to enable peer-produced privacy protection.

The current paradigm of privacy protection considers information sharing as an individual effort. Thus, both technical solutions such as PETS and regulatory efforts such as DNT only account for and allow for individual decisions. In this paper we argue for considering information sharing to be a community interaction and thereby engendering privacy solutions that are holistic, economic, and efficient.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Award Nos. CNS-1228364 and CNS-1250367 as well as Intel through the ISTC for Secure Computing. We also thank John McCurley for his editorial help.

REFERENCES

[1] G. A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *Security Privacy, IEEE*, vol. 3, no. 1, pp. 26 –33, jan.-feb. 2005.

[2] S. Lederer, I. Hong, K. Dey, and A. Landay, "Personal privacy through understanding and action: five pitfalls for designers," *Personal Ubiquitous Comput.*, vol. 8, no. 6, pp. 440–454, Nov. 2004. [Online]. Available: http://dx.doi.org/10.1007/s00779-004-0304-9.

[3] S. Patil and A. Kobsa, "Uncovering privacy attitudes and practices in instant messaging," in *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, ser. GROUP '05. New York, NY, USA: ACM, 2005, pp. 109–112. [Online]. Available: http://doi.acm.org/10.1145/1099203.1099220.

[4] S. Patil, Y. L. Gall, A. J. Lee, , and A. Kapadia, "My privacy policy: Exploring end- user specification of free-form location access rules," in *Proceedings of the Workshop on Usable Security (USEC)*, ser. Lecture Notes in Computer Science, vol. 7398. Springer Berlin / Heidelberg, Feb. 2012, pp. 86–97.

[5] S. Spiekermann, J. Grossklags, and B. Berendt, "E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior," in *Proceedings of the 3rd ACM conference on Electronic Commerce*, ser. EC '01. New York, NY, USA: ACM, 2001, pp. 38–47. [Online]. Available: http://doi.acm.org/10.1145/501158.501163.

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

[7] Consumer Reports Magazine, "Facebook & your privacy: Who sees the data you share on the biggest social network?" http://www.consumerreports.org/cro/magazine/2012/06/faceboyour-privacy/index.htm, June 2012.

[8] A. Acquisti and R. Gross, "Imagined communities: Awareness, information sharing, and privacy on the facebook," in *Privacy enhancing technologies*. Springer, 2006, pp. 36–58.

[9] T. Paul, D. Puscher, and T. Strufe, "Improving the usability of privacy settings in facebook," *arXiv preprint arXiv:1109.6046*, 2011.

[10] K. Strater and H. R. Lipford, "Strategies and struggles with privacy in an online social networking community," in *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction - Volume 1*, ser. BCS-HCI '08. Swinton, UK, UK: British Computer Society, 2008, pp. 111– 119. [Online]. Available: http://dl.acm.org/ citation.cfm?id=1531514.1531530.

[11] L. Fang and K. LeFevre, "Privacy wizards for social networking sites," in *Proceedings of the 19th international conference on World wide web*, ser. WWW '10. New York, NY, USA: ACM, 2010, pp. 351– 360. [Online]. Available: http://doi.acm.org/ 10.1145/1772690.1772727.

[12] S. Gürses and B. Berendt, "Pets in the surveillance society: A critical review of the potentials and limitations of the privacy as confidentiality paradigm," in *Data Protection in a Profiled World*, S. Gutwirth, Y. Poullet, and P. De Hert, Eds. Springer Netherlands, 2010, pp. 301–321. [Online]. Available: http: //dx.doi.org/10.1007/978-90-481-8865-919.

[13] E. Ostrom, J. Walker, and R. Gardner, "Covenants with and without a sword: Self- governance is possible," *The American Polit- ical Science Review*, vol. 86, no. 2, pp. 404– 417, 1992.

[14] T. Dietz, E. Ostrom, and P. Stern, "The struggle to govern the commons," *Science*, vol. 302, no. 5652, pp. 1907–1912, 2003.

[15] J. Bonneau and S. Preibusch, "The privacy jungle:on the market for data protection in social networks," in *Economics of Information Security and Privacy*, T. Moore, D. Pym, and C. Ioannidis, Eds. Springer US, 2010, pp. 121–167. [Online]. Available: http: //dx.doi.org/10.1007/978-1-4419-6967-5 8

[16] J. Gideon, L. Cranor, S. Egelman, and A. Acquisti, "Power strips, prophylactics, and privacy, oh my!" in *Proceedings of the second symposium on Usable privacy and security*, ser. SOUPS '06. New York, NY, USA: ACM, 2006, pp. 133–144. [Online]. Available: http://doi.acm.org/10.1145/1143120.1143137

[17] R. Anderson and T. Moore, "The economics of information security," *Science*, vol. 314, no. 5799, pp. 610–613, 2006.

[18] [A. Miyazaki and S. Krishnamurthy, "Internet seals of approval: Effects on online privacy policies and consumer perceptions," *Journal of Consumer Affairs*, vol. 36, no. 1, pp. 28– 49, 2002.

[19] A. McDonald and L. Cranor, "Cost of reading privacy policies, the," *ISJLP*, vol. 4, p. 543, 2008.

[20] A. Beautement, M. A. Sasse, and M. Wonham, "The compliance budget: managing security behaviour in organ- isations," in *Proceedings of the 2008 workshop on New security paradigms*, ser. NSPW '08. New York, NY, USA: ACM, 2008, pp. 47–58. [Online]. Available: http://doi.acm.org/10.1145/1595676.1595684

[21] R. Schlegel, A. Kapadia, and A. J. Lee, "Eyeing your exposure: Quantifying and con- trolling information sharing for improved privacy," in *Proceedings of the 2011 Symposium on Usable Privacy and Security (SOUPS)*, Jul. 2011.

[22] H. Nissenbaum, *Privacy in context: Technology, policy, and the integrity of social life*. Stanford Law Books, 2009.

[23] R. Posner, "The right of privacy," *Georgia Law Review*, vol. 12, no. 3, pp. 393–422, 1977.

[24] R. Posner, "The economics of privacy," *The Amer- ican economic review*, vol. 71, no. 2, pp. 405– 409, 1981.

[25] G. Hardin, "Tragedy of the commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, 1968.

[26] R. Bo¨hme, S. Koble, and T. Dresden, "On the viability of privacy-enhancing technologies in a self-regulated business-to-consumer market: Will privacy remain a luxury good," in *Workshop on the Economics of Information Security*. Pittsburgh, PA: Carnegie Mellon University, 2007.

[27] L. Brandimarte, A. Acquisti, and G. Loewen- stein, "Misplaced confidences: Privacy and the control paradox," *Social Psychological and Personality Science*, 2012.

[28] F. Berkes, "Fishermen and 'The tragedy of the commons'," *Environmental Conservation*, vol. 12, no. 03, pp. 199–206, September 1985.

[29] W. Ascher, "Communities and sustainable forestry in developing countries," Duke University, Tech. Rep., 1995.

[30] J. Rosen, "The right to be forgotten," *Stanford Law Review Online*, vol. 64, p. 88, 2012.

[31] P. Dourish, "Culture and control in a media space," in *Proceedings of the third confer- ence on European Conference on Computer- Supported Cooperative Work*. Kluwer Aca- demic Publishers, 1993, pp. 125–137.

[32] S. Patil and A. Kapadia, "Are you exposed? conveying information exposure (extended abstract)," in *Proceedings of The 2012 ACM Conference on Computer Supported Cooper- ative Work Companion (CSCW)*, Feb. 2012, pp. 191–194.

[33] Y. L. Gall, A. J. Lee, and A. Kapadia, "PlexC: A policy language for exposure control," in *Proceedings of The 17th ACM Symposium on Access Control Models and Technologies (SACMAT)*, Jun. 2012, pp. 219–228.

[34] S. Kiesler, B. Zdaniuk, V. Lundmark, and R. Kraut, “Troubles with the internet: the dynamics of help at home,” *Hum.- Comput. Interact.*, vol. 15, no. 4, pp. 323– 351, Dec. 2000. [Online]. Available: http://dx.doi.org/10.1207/S15327051HCI1504 2.

[35] E. S. Poole, M. Chetty, R. E. Grinter, and W. K. Edwards, “More than meets the eye: transforming the user experience of home network management,” in *Proceedings of the 7th ACM conference on Designing interactive systems*, ser. DIS ’08. New York, NY, USA: ACM, 2008, pp. 455–464. [Online]. Available: http://doi.acm.org/10.1145/1394445.1394494.

[36] E. S. Poole, M. Chetty, T. Morgan, R. E. Grinter, and W. K. Edwards, “Computer help at home: methods and motivations for informal technical support,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI ’09. New York, NY, USA: ACM, 2009, pp. 739– 748. [Online]. Available: http://doi.acm.org/10.1145/1518701.1518816.

[37] L. J. Camp, “Reliable, usable signaling to defeat masquerade attacks,,” in *Workshop on the Economics of Information Security*, Cambridge, UK, Jun. 2006.

[38] A. Dieberger, P. Dourish, K. Ho ̈o ̈k, P. Resnick, and A. Wexelblat, “Social navigation: techniques for building more usable systems,” *interactions*, vol. 7, no. 6, pp. 36–45, Nov. 2000. [Online]. Available: http://doi.acm.org/10.1145/352580.352587.

[39] J. Goecks and E. Mynatt, “Supporting privacy management via community experience and expertise,” *Communities and Technologies*, pp. 397–417, 2005.

[40] J. Goecks, W. K. Edwards, and E. D. Mynatt, “Challenges in supporting end-user privacy and security management with social navigation,” in *Proceedings of the 5th Sym- posium on Usable Privacy and Security*, ser. SOUPS ’09. New York, NY, USA: ACM, 2009, pp. 5:1–5:12. [Online]. Available: http://doi.acm.org/10.1145/1572532.1572539.

[41] T. Moore and R. Clayton, “Evaluating the wisdom of crowds in assessing phishing websites,” in *Financial Cryptography*, ser. Lecture Notes in Computer Science, G. Tsudik, Ed., vol. 5143. Springer, 2008, pp. 16–30. [Online]. Available: http://lyle.smu.edu/~tylerm/fc08.pdf.

[42] A. Besmer, J. Watson, and H. R. Lipford, “The impact of social navigation on privacy policy configuration,” in *Proceedings of the Sixth Symposium on Usable Privacy and Security*, ser. SOUPS ’10. New York, NY, USA: ACM, 2010, pp. 7:1–7:10. [Online]. Available: http://doi.acm.org/10.1145/1837110.1837120.

[43] S. Patil, X. Page, and A. Kobsa, “With a little help from my friends: can social navigation inform interpersonal privacy preferences?,” in *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, ser. CSCW ’11. New York, NY, USA: ACM, 2011, pp. 391–394. [Online]. Available: http://doi.acm.org/10.1145/1958824.1958885.

[44] G. Hardin, “Tragedy of the Commons,” *Science,* vol. 162, num. 3859, pp. 1243-1248, 1968.

[45] D.M. Kreps, P. Milgrom, J. Roberts, and R. Wilson, “Rational cooperation in the finitely repeated prisoners' dilemma,” *Journal of Economic Theory,* vol. 27, num. 2, pp. 245-252, 1982.