A Subsidy Model for Social Media Interventions

Renita Murimi
University of Dallas
rmurimi@udallas.edu

Abstract

User activities in online social networks can be viewed as a set of actions with different payoffs. We examine how these payoffs differ with the level of anonymity afforded by a network, and the corresponding implications for the kinds of content posted by a user. Unlike in offline social networks, in online social networks, users can choose to present to their network highly curated versions of events of their real lives, which can lead to warped portrayals of a user's life events to others on the network. In this paper, we introduce AI-enabled intervention mechanisms to mitigate the impact of users' myopic behavior through curated postings that might not be entirely truthful. We evaluate these AI-interventions in terms of their impact on the user's perception of privacy and reputation on the network. We demonstrate that these mechanisms can lead to network design alternatives in which the network affordances can enhance broad user experience in social networks.

Introduction

The ability to gather information about user preferences and activities on social networking sites (SNSs) offers enormous incentives for marketing research and political campaigns. The affordances provided by the network for engaging with other users and their content lead to multiple avenues for content dissemination. For example, the network affordances on Facebook and similar other SNSs that enable a user to *like*, *share*, and comment may cause the content to go "viral", causing a significantly larger audience reaction to the content. While the affordances provided by a network allow users multiple ways to engage with content and control their privacy, it also provides tools to curate one's data. Since users choose the content they want to share on SNSs, it is not surprising that users often share curated versions of their lifestyles, that highlight the most significant aspects of their lives. It has been shown that the perceptions of other's lifestyles on social media has been linked to higher incidence of depression among users (Sidani et al, 2016) because of the impression of a "perfect" life that can be relatively easier to cultivate with the aid of network affordances on social media.

People's motivations to share on SNSs have been extensively studied (Nadkarni and Hoffman 2012; Morris, Teevan and Panovich, 2010). In (Christofides, Muise, Desmarais, 2009), the authors suggested that self-esteem, privacy and sharing tendency were all linked, where users with high self-esteem were shown to be more private, and thus, less needing of input from other users on SNS platforms. This tendency to intertwine self-esteem with feedback from other SNS users also has implications for user behavior in anonymous SNSs. Investigation of SNSs have shown that differing levels of anonymity exist, thus, the spectrum of anonymity offers people a menu of different platforms for specific networking needs. On one end of this spectrum of anonymity, lie SNSs such as Facebook, where the profile-centric architecture of the platform strongly associates a user's content to his/her profile. On the other end of the anonymity spectrum, lie anonymous networks such as Whispr where users can post content without the need for creating a profile. Studies have shown the online disinhibition effect (Suler, 2004), where people use the affordance of anonymity to act without inhibition and without consideration of social norms on anonymous networks, leads to creation of content that generally tends to feature more negative sentiment and which would not be appropriate for a non-anonymous network (Zhang and Kizilcec, 2014).

Thus, the level of anonymity afforded by a network plays a significant role when users choose to share content. In addition to causing users to consider the payoff from posting content, users are also presented with warped accounts of other's lifestyles resulting from other users' decisions to present curated versions of their lifestyles on SNSs. We ask: does the level of anonymity afforded by an SNS affect the user's decision to post content on that platform? We introduce a model that investigates this decision as a function of the following user perceptions: anonymity afforded by the SNS, privacy, user's own reputation and the reputation of the content. We introduce the idea of a subsidy-based AI-enabled intervention to incentivize users to share content of various kinds-- not just the content that boosts one's perceived reputation in the network—and thereby mitigate the creation and influence of curated user profiles. We show that the effort to create this subsidy-based intervention (computational complexity of additional AI algorithms) can be funded through a network tax imposed on users or funded by the platform (for example, a fee to use the SNS), thus creating networks that encourage users to share more truthful accounts of their lives on SNSs.

Reputation, Privacy and Anonymity

An economic perspective to the study of social networks has been studied in both online and offline networks (Mayer, 2009; Eguiluz et al, 2014; Jackson and Watts, 2002; Manski, 2000; Rand and Nowak, 2013). Motives for sharing have been studied extensively. In (Parameswaran and Whinston, 2007), the authors found that online communities offered avenues to demonstrate individualism and altruism for sharing and communal benefits. In (Ligon and Schecter, 2012), the authors examined motives for sharing in an offline setting. Here, they studied why and with whom people chose to share in risky environments. The motives for sharing were ascribed to benevolence, altruism, reciprocity, and sanctions. They also found that the more connected an individual was on a network, the more valuable the social network was for the user and the higher was the reciprocity that the user demonstrated on the network. In (Krasnova et al 2010), the authors found users' motives for disclosure on social network arose primarily from the convenience that the platform afforded for creating and maintaining connections, and self-enjoyment.

The relationship between reputation and anonymity has been investigated in recent literature. In (Mui, Mohtashemi and Halberstadt, 2002), the authors propose a computational model for studying the relationship between trust and reputation. In (Bonneau, Dannezis and Anderson, 2009), the authors show that privacy in social networks is tremendously under-estimated in social networks. They showed that only a small percentage of accounts are needed to view the majority of the network, since friend-of-friend discovery makes it much easier to view the content posted by others outside one's immediate friend networks. Further, the work in (Gross and Acquisti, 2005; Acquisti, Brandimarte and Loewenstein 2015) found that social networks encouraged users to post personal and often sensitive information, despite the existence of only weak ties between a user and most of her contacts on the social network. In (Peddinti, Ross and Cappos, 2014), the authors studied the effect of anonymity on online behavior in Twitter. They found that anonymous users were more likely than identifiable users to engage in active behavior on Twitter (tweeting instead of merely lurking, following more accounts and more willing to expose their activity). Might there be a similar affect in place in SNSs? How does anonymity, or various levels of it, affect what users share in SNSs? We study user behavior in terms of the impact of their posted content on their reputation and privacy as a function of the anonymity offered by the network. To counter the myopic behavior encountered by users when they choose to boost their own payoffs by posting curated versions of their profiles, we introduce a Pigouvian subsidy-based intervention to incentivize truthful content sharing and impose a tax on the users to fund this subsidy.

Payoff Model

We now formally describe the model to study anonymity as a function of privacy and the resulting payoff from content sharing on SNSs. Consider two players p_1 and p_2. Let their payoffs from sharing content on the network be π_1 and π_2 respectively. The two commodities available for consumption are privacy (x) and self-reputation (y). Let the initial endowment of these two commodities be $\bar{x}$ units of x and $\bar{y}$ units of y.

$$\pi_1(x, y) = x^{\alpha_1} y^{\beta_1/\alpha_1} \tag{1}$$

$$\pi_2(x, y) = x^{\alpha_2} y^{\beta_2/\alpha_2} \tag{2}$$

where, $0 < \alpha_i, \beta_i$ and α_i is the perceived anonymity that player i experiences on the networks and β_i is the perceived reputation of the content that player i evaluates before posting on a network. Thus, in order to maximize the user's privacy and self-reputation on the SNS, we consider the maximization problem:

$$\lambda x^{\alpha_1} y^{\beta_1/\alpha_1} + (1-\lambda)(\bar{x}-x)^{\alpha_2}(\bar{y}-y)^{\beta_2/\alpha_2} \tag{3}$$

For $0 < \lambda < 1$, finding the first-order conditions and solving for y yields

$$y = \{(\gamma\omega^{\rho})/(1+\gamma\omega^{\rho})\}y' \tag{4}$$

where, the constant terms γ and ρ are defined as $\gamma = \alpha_1\beta_2/\alpha_2^2\,\beta_1$, and $\rho = \alpha_2/(\beta_2-\alpha_2)$ and the variable ω is defined as $(\bar{x}-x)/x$. We label the term γ as the Network Response Indicator (NRI), and denotes the user's tendency to share content on an SNS.

Externality Model

A user's decisions to share content that results in increase in his/her perceived value of self-reputation on the network yields a higher payoff to the user. However, this myopic behavior causes the user to project an illusory "perfect" life, a warped personality on to the feeds of friends on the network. This undesirable effect is a consequence of the externality imposed by the user's actions on the rest of the network. Let the externality experienced by user i be denoted as v_i. This externality is a consequence of the aggregate actions of other users of the network. The option to present a curated version of life events and posts has the potential to present inaccurate versions of an individual's lifestyle and personality to other users of the site. We model this misrepresentation as v_i.

Let v_i be a linear function of the self-reputation y (which, in turn depends on the perceived anonymity α and the content-reputation β). Thus, the payoff π_i is now reduced by the value of the externality v_i. For N posts, the payoffs of user p_i is given by

$$\pi_i(x, y) = (N - x - y)x^{\alpha_i}y^{\beta_i/\alpha_i} - cy \tag{5}$$

Obtaining the first-order conditions and setting them equal to zero gives us the value of x and y for user p_1 at equilibrium:

$$x^* = \left(N - \alpha_i y^{\beta_i/\alpha_i}\right)/1 + \alpha \tag{6}$$

$$y^* = (N - x)/(1 + \alpha/\beta) \tag{7}$$

Introduction of Pigouvian Subsidies to the Model

In order to mitigate the effects of the externality imposed by the user's myopic behavior, we introduce an AI-enabled intervention based on a Pigouvian subsidy. The notion of an externality and associated subsidies is a foundational concept in economics literature (Dahlman, 1979). In this section, we employ the same lens to investigate the design of better network interventions in SNSs to facilitate content-sharing on social networks to facilitate a more realistic portrayal of users' lives.

Assume that a user i's contribution of accurate/truthful posts = π_i. Let, the costs incurred due to the more accurate representation of user's content (for example, social embarrassment, potential decline in reputation) be π_i^2. Thus, the utility of player i is maximized by

$$\max_{\pi_i} \sum_{i=1}^{N} \pi_i - \pi_i^2 \tag{8}$$

The best response of user i to $\boldsymbol{\pi}_{-i}$ (Nash equilibrium) is, therefore, given by 0.5. Thus, user i's payoff at equilibrium is given by

$$N(0.5) - (0.5)^2 = (2N - 1)/4 \tag{9}$$

The total utilitarian welfare (social surplus) is given by

$$N\sum_{i=1}^{N} \pi_i - \sum_{i=1}^{N} \pi_i^2 \tag{10}$$

To maximize it, $\max\limits_{\hat{\pi}} N(N \times \hat{\pi} - \hat{\pi}^2)$. Thus, the payoff that maximizes the social surplus is given by

$$\pi_s = N/2 \tag{11}$$

Consider an AI-generated Pigouvian subsidy that incentivizes the generation of truthful posts by rewarding σ to a user for every truthful post. The truthfulness (non-conformity to traditional posting patterns) requires additional network resources (computational complexity of algorithms, bandwidth). These resources may be funded through an additional tax τ imposed on all users.

The best response function of player i is now given by

$$\max_{\hat{\pi}} \pi_1 + \pi_2 + \cdots + \pi_i(1 + \sigma) + \cdots \pi_N - \pi_i^2 - \tau\sigma \textstyle\sum_{j \neq i} \pi_j/N - 1 \tag{12}$$

Thus, the best response of player i in the presence of a subsidy and tax is given by

$$\pi_{i,\sigma}^* = (1 + \sigma)/2 \tag{13}$$

Next, we proceed to develop an intervention for the calculation of tax τ that will fund the AI-enabled Pigouvian subsidy. Thus, from equation (13), we see that the total tax to be collected from the users is given by $\sigma\tau\pi_{i,\sigma}^*$.

The individual payoff, $\pi_{i,\sigma}$, is given by the cost of posts minus the (Pigouvian subsidy for user i) minus the (cost for user i – tax burden of user i). Thus,

$$\pi_{i,\sigma} = N\pi_{i,\sigma}^* + \pi_{i,\sigma}^* \times \sigma - \left(\pi_{i,\sigma}^*\right)^2 - \tau\sigma\pi_{i,\sigma}^* \tag{14}$$

For all users,

$$U(\sigma) = Nu_i(\sigma) \tag{15}$$

Substitute (13) in (15) and obtaining the first-order conditions, we get the value of the optimal subsidy σ_O.

$$\sigma_o = (N - \tau)/(2\tau - 1) \tag{16}$$

Thus, the payoff for user i at optimal subsidy σ_o is given by

$$\pi_{i,o}^* = \frac{1+\sigma}{2} = \frac{N+\tau-1}{2(2\tau-1)} \tag{17}$$

Equating this to the payoff as a function of perceived anonymity α, content reputation β, self-reputation y and the privacy x, we get the underlying relationship as follows:

$$\frac{N+\tau-1}{2(2\tau-1)} = x^{\alpha}y^{\frac{\beta}{\alpha}} \tag{18}$$

Model Evaluation

The data for this work was collected through an anonymous survey administered to college students ($N = 102$). Participant information related to gender and age group was obtained. The survey asked if participants had accounts on anonymous networks and non-anonymous

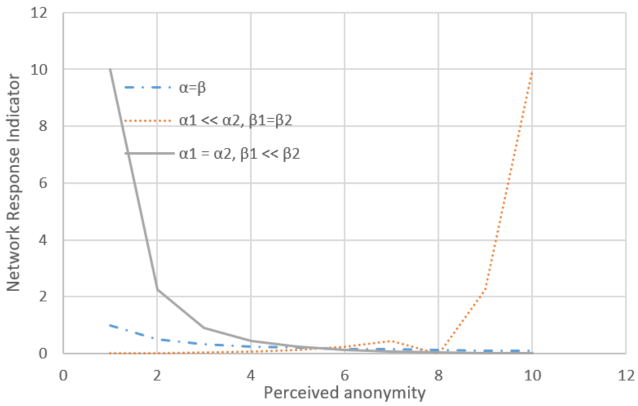

Figure 1. Plot of Network Response Indicator (γ) versus perceived anonymity.

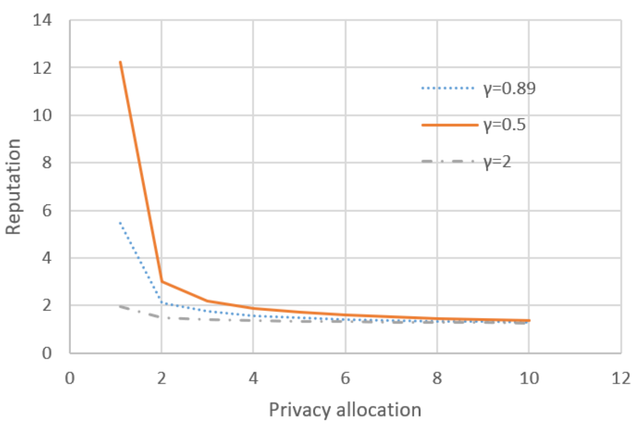

Figure 2. Plot of reputation versus privacy.

networks. Further, they were asked if participants would post content about a reputation-increasing or reputation-decreasing event on an anonymous network or a non-anonymous network. Finally, the survey asked if network size mattered in the participant's decision to post on an anonymous/non-anonymous network. Survey responses indicated that network size did not matter (98%) in deciding whether to post on anonymous or non-anonymous networks. Respondents also indicated that they preferred to share reputation-boosting events on non-anonymous networks (99.3%) and reputation-decreasing events on anonymous networks (99.7%).

For the evaluation of our model, we study the tradeoff between reputation and privacy at equilibrium and non-equilibrium conditions, first with myopic user behavior and then with an AI-generated Pigouvian subsidy and a tax to mitigate the effects of curated profiles resulting from users' decisions to boost their individual payoffs. Figure 1 plots the NRI (γ) (from equation 4) as a function of the perceived anonymity α. We observe two distinct behaviors summarized below.

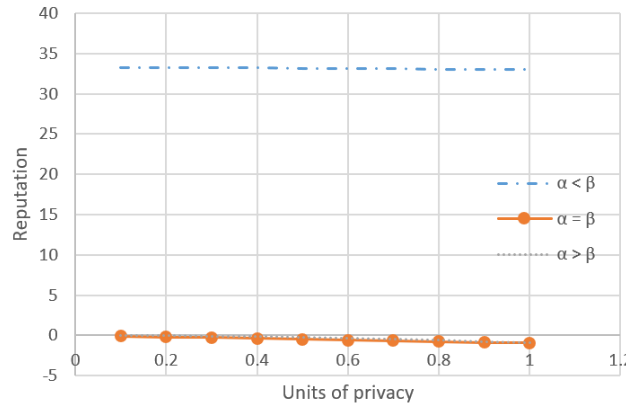

Figure 3. Plot of reputation versus privacy at equilibrium.

Case 1: ($\alpha_1 = \alpha_2, \beta_1 \ll \beta_2$) Users do not perceive a difference in anonymity between networks. Thus, they do not care for the level of anonymity afforded by the network or are unaware of the implications for their own social reputation on the network. There are layers of self-disclosure on networks and anonymity exists on a spectrum. As the perceived anonymity increases, the NRI decreases. Thus, whether the user is posting about a reputation-boosting event or a reputation-decreasing event, since the user does not care about the anonymity, there is a decrease in the tendency to post the event. There is no significant payoff for the user to engage with the network.

Case 2: ($\alpha_1 \ll \alpha_2, \beta_1 = \beta_2$) Users who do not perceive a difference in content-reputation but are concerned about the anonymity on the network. With increase in perceived anonymity, there is an increase in the tendency to post content resulting in a higher payoff for the user. This is especially significant for frequent sharers on a network, since anonymity affords them the chance to speak their minds on a digital platform more freely.

Figure 2 depicts the non-equilibrium tradeoff between reputation and privacy with varying values of the NRI. We see that with the increase in privacy, users do not share as much. In a network where users are constantly posting content, their reputation is constantly varying depending on the content that is posted. The highest value of reputation is when privacy is the least and the NRI is the highest. This suggests support for users in non-anonymous settings when the payoff is the highest.

In Figure 3, at equilibrium conditions between privacy and self-reputation, if the perceived anonymity is less than the content-reputation ($\alpha < \beta$), the user enjoys high reputation in the network. This supports users who are on a non-anonymous network, and the tradeoff between anonymity and content-reputation skews in favor of content-reputation. Thus, the user is inclined to post reputation-boosting content that in turn increases user reputation on the network. As perceived anonymity increases, the user's reputation on the

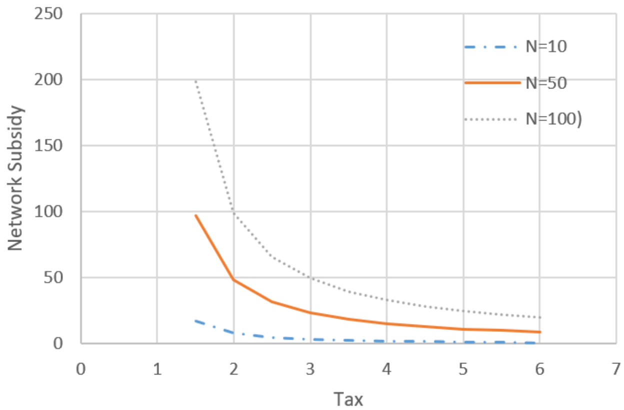

Figure 4. Plot of network subsidy versus tax.

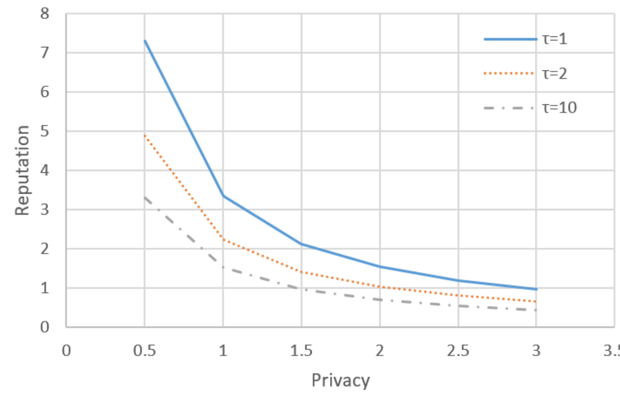

Figure 5. Plot of network subsidy versus tax.

network is dictated primarily by content-reputation. For a post of lower content-reputation, the user's reputation on the network is lowered.

While Figures 2 and 3 showcased the impact of users' myopic behavior on reputation and privacy, Figure 4 depicts the network behavior with the addition of an AI-generated Pigouvian subsidy to incentivize truthful content generation and an associated tax to fund the subsidy. From Figure 4, we see that with increase in network size, the subsidy that incentivizes users to post more truthful posts is higher. Subsequently as tax (which funds the subsidies) increases, the subsidy is higher. Figure 5 shows the tradeoff between self-reputation and privacy at equilibrium in the presence of an AI-generated Pigouvian subsidy and tax. At equilibrium, as privacy increases, self-reputation decreases since users do not post as much content. We also see that as the tax increases, reputation decreases since users are posting more truthful accounts of their lives.

Discussion

The decision to self-disclose on SNSs is a consequence of several factors. In this work, we studied how the user's perceptions of (a) anonymity on the network, (b) the content that he/she is considering disclosing, (c) the user's own reputation, and the (d) the privacy afforded by the platform affected the user's decision to post content. Specifically, we studied whether users are inclined to post reputation-increasing content on anonymous or non-anonymous networks, and whether network size plays a role in this decision to disclose on either kind of network. We found that while users' decisions to post on either kind of network were not impacted by network size, they decided to post reputation-increasing events on non-anonymous networks and reputation-decreasing events on anonymous networks. These findings are in line with several studies performed in social science and network science literature, where people perform actions in online or offline networks in line with the perceived payoff from the action. Thus, reputation-decreasing events are not disclosed on non-anonymous networks. While this myopic behavior is beneficial to the user, it also paints an inaccurate picture of one's life or personality on an SNS for other users.

Our work has proposed an AI-generated solution to incentivize truthful content generation and dissemination on SNSs. The motivating construct here is that of trust. Users are influenced by the content posted by their peers, and the presence of network effects lead to outsized influence by few accounts (for example, celebrities or influencers) over significant numbers of their followers. When users post on a SNS, in effect, they might be recommending behaviors or activities to the rest of their network. This calls into question the role of the SNSs and the participating users. Does the SNS act as a clearinghouse for recommendations? If that is the case, then users have to evaluate the amount of trust that they will place in a recommendation or opinion from a fellow user on the network. This leads to an inherent information asymmetry, where users are not able to discern the veracity of content, and consequently the amount of trust to associate with the content. From an economics points of view, this creates a scenario for an incomplete contract. SNS platforms, therefore, should internalize the costs of these incomplete contracts, because users do not have the ability to discern the veracity of content. Recent examples of SNSs such as Facebook and Twitter stepping in to flag misleading content are an important network intervention in that direction. The use of AI-enabled tools can help to design interventions that bridge the gap between the perception and veracity of content in SNSs. However, it also raises important concerns regarding the role of SNSs in supporting or limiting individual rights to freedom of expression in online environments, and the potential for algorithmic biases in AI tools to exacerbate existing issues in generation and dissemination of online content.

Limitations

Our model for AI-enabled SNS interventions has supported the need for studies about network-wide incentives that are in turn, abetted by subsidies and taxes. Our findings point to the efficacy of these interventions for rethinking the structure of existing SNSs and the nature of user interaction with these SNSs. Some limitations of our work include the modeling of anonymity and the role of individual factors in the design of interventions, and are described below.

Anonymity: This work investigated the role of anonymity in SNSs by analyzing two kinds of SNS networks – anonymous and non-anonymous networks. However, in practice, anonymity in networks exists on a spectrum. Even within a single non-anonymous network, progressive levels of increasing anonymity can be achieved by leveraging the account settings. These settings can help achieve granular control over the entire account and individual activities. Incorporating anonymity on a spectrum instead of a binary model would result in richer insights about the role of anonymity in design SNS interventions.

Diversity of interventions: Our proposed model studies the role of four factors: user perception of anonymity, reputation, content and privacy in a user's decision to post content. Several other factors, such as the role of competition and cooperation among users on the network, the novelty of the platform and the personality traits such as introversion may also play a role in the user's decision to participate. Accordingly, the SNS interventions designed will reflect the role of these diverse factors in influencing user activities on the network. The development of more rigorous models will aid in making the connections between the factors that motivate users to participate in SNSs and effective network interventions to facilitate truthful content generation.

Future Work

Myopic behavior on the part of SNS users is inevitable due to the current design of SNSs. Although our work has focused on the self-disclosure behavior, our model can be extended to study the phenomenon of the spread of false information. This phenomenon has affected multiple aspects of our society and has impacted the spread of information concerning areas such as political events and public health campaigns. The implications of this myopic behavior are several, and call into question the trustworthiness of SNSs and undermines public trust in online information. Thus, it is imperative to redefine the structure of SNSs to allow for two key social constructs to be reliably represented: trust and privacy. There are multiple avenues for future work in this topic, including an understanding of data ownership, and accountability for moral hazards of social media participation and adverse selection.

Moral hazards: From an economics perspective, moral hazard refers to the phenomenon where users engage in behavior that is harmful to others when they are not held fully responsible for their actions. Examples of such behavior include leaving the lights on, or the irresponsible use of plastic products. Participation in SNSs creates similar moral hazards. If an influential user posts false information, will she be held responsible for the adverse consequences after the viral spread of her posts? If not, such activity sets up a moral hazard for other users of the network and the consequences might include others who are not even users of the SNS. Further research is required to determine mechanisms to mitigate moral hazards of user participation in SNSs.

Data ownership: Should users be incentivized to share their data? This begets another question about the value of our data to SNS platforms. Previous research about user tendencies has shown that users participate with varying frequencies on SNSs. Some users participate frequently, others do so moderately, while some others merely lurk (Murimi, 2016). Further, a study of SNSs has shown unique network effects in topology such as the presence of large connected components and a few islands of users. Thus, if a lurker finds utility from not participating in the network and merely observing the activity of fellow users, should her activity be rewarded according to the frequency of her participation? Additional room for investigation arises from the classification of activities themselves. Is original content to be regarded differently from reshared content, and should the rewards be prorated according to the kind of activity? SNSs have given rise to news kinds of job roles such as "influencers", where individuals gain massive followers for their activities and who in turn use this influence to make product recommendations. Determining data ownership, classifying activities and determining rewards can help to design effective interventions for the use of SNSs as a formidable tool for information dissemination.

Conclusions

SNS users exhibit myopic behavior by posting reputation-increasing events on non-anonymous networks and reputation-decreasing events on anonymous networks. This myopic behavior introduces an externality on the network, which can be mitigated through a subsidy-based intervention to incentivize more truthful content generation. The subsidy itself is funded by a tax imposed on all the users of the network. The economic perspective of an externality-mitigating subsidy and tax can be envisioned and implemented with AI models in several ways. In one such model, the tax imposed on users can be generated through a fee imposed on users to use the network. The introduction of a subsidy has significant implications in the design of network affordances that encourage users to post a variety of content. Applications of this induced subsidy could include AI algorithms for sentiment mining, bot detection and veracity that check for both content and context of user activity, and are able to dole out rewards/penalties dynamically in response. The consequent implications for network design are of growing importance for the next generation of social networks, where users are growing increasingly mindful of their privacy, and are also simultaneously disclosing a lot about themselves on online social networks.

References

Clancey, W. J. 1984. Classification Problem Solving. In *Proceedings of the Fourth National Conference on Artificial Intelligence*, 49-54. Menlo Park, Calif.: AAAI Press.

Acquisti, A., Brandimarte, L., and Loewenstein, G. 2015. Privacy and Human Behavior in The Age of Information. *Science*, 347(6221), 509-514.

Christofides, E., Muise, A., and Desmarais, S. 2009. Information Disclosure and Control on Facebook: Are They Two Sides of the same coin or two different processes? *Cyberpsychology & Behavior*, 12(3), 341-345.

Dahlman, C. J. 1979. The Problem of Externality. *The Journal of Law and Economics*, 22(1), 141-162.

Bonneau, J., Anderson, J., and Danezis, G. 2009. Prying Data Out of a Social Network. In *Proceedings of the International Conference on Advances Social Network Analysis and Mining (ASONAM '09)*, 249-254.

Eguíluz, V. M., Zimmermann, M. G., Cela-Conde, C. J., and Miguel, M. S. 2005. Cooperation and The Emergence of Role Differentiation in the Dynamics of Social Networks. *American Journal of Sociology*, 110(4), 977-1008.

Gross, R., and Acquisti, A. 2005. Information Revelation and Privacy in Online Social Networks. In *Proceedings of the ACM workshop on Privacy in the Electronic Society*, 71-80.

Jackson, M. O., and Watts, A. 1998. The Evolution of Social and Economic Networks. Journal of Economic Theory, 106(2), 265-295.

Krasnova, H., Spiekermann, S., Koroleva, K., and Hildebrand, T. 2010. Online Social Networks: Why We Disclose. *Journal of Information Technology*, 25(2), 109-125.

Ligon, E., and Schechter, L. 2012. Motives for Sharing in Social Networks. *Journal of Development Economics*, 99(1), 13-26.

Manski, C. F. 2000. Economic analysis of social interactions. *Journal of Economic perspectives*, 14(3), 115-136.

Mayer, A. 2009. Online Social Networks in Economics. *Decision Support Systems*, 47(3), 169-184.

Morris, M. R., Teevan, J., and Panovich, K. 2010. What Do People Ask Their Social Networks, And Why? A Survey Study of Status Message Q&A Behavior. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, 1739-1748.

Murimi, R. 2016. An Analysis of Trimming in Digital Social Networks. In *Proceedings of the AAAI Workshop on Incentives and Trust in e-Communities*.

Mui, L., Mohtashemi, M., and Halberstadt, A. 2002. A Computational Model of Trust and Reputation. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS)*, 2431-2439.

Nadkarni, A., and Hofmann, S. G. 2012. Why Do People Use Facebook? *Personality and Individual Differences*, 52(3), 243-249.

Parameswaran, M., and Whinston, A. B. 2007. Research Issues in Social Computing. *Journal of the Association for Information Systems*, 8(6), 336.

Peddinti, S. T., Ross, K. W., and Cappos, J. 2014. On the Internet, Nobody Knows You're a Dog: A Twitter Case Study of Anonymity in Social Networks. In *Proceedings of the Second ACM Conference on Online social networks*, pp. 83-94.

Rand, D. G., and Nowak, M. A. 2013. Human cooperation. *Trends in Cognitive Sciences*, 17(8), 413-425.

Sidani, J. E., Shensa, A., Radovic, A., Miller, E., Colditz, J. B., Hoffman, B. L., ... and Primack, B. A. 2016. Association Between Social Media Use and Depression Among US Young Adults. *Depression and Anxiety*, 33(4), 323-331.

Suler, J. 2004. The Online Disinhibition Effect. *Cyberpsychology and Behavior*, 7(3), 321-326.

Zhang, K. and Kizilcec, R. F. 2014. Anonymity in Social Media: Effects of Content Controversiality and Social Endorsement on Sharing Behavior. In *Proceedings of the International Conference on Web and Social Media (ICWSM '14)*.