

Trust and Privacy in Distributed Work Groups

Denise Anthony¹, Tristan Henderson², and James Kitts³

¹ denise.anthony@dartmouth.edu, Dartmouth College, Hanover, NH, USA

² tristan@cs.st-andrews.ac.uk, University of St Andrews, St Andrews, Fife, UK

³ jak2190@columbia.edu, Columbia University, New York, NY, USA

Abstract Trust plays an important role in both group cooperation and economic exchange. As new technologies emerge for communication and exchange, established mechanisms of trust are disrupted or distorted, which can lead to the breakdown of cooperation or to increasing fraud in exchange. This paper examines whether and how personal privacy information about members of distributed work groups influences individuals' cooperation and privacy behavior in the group. Specifically, we examine whether people use others' privacy settings as signals of trustworthiness that affect group cooperation. In addition, we examine how individual privacy preferences relate to trustworthy behavior. Understanding how people interact with others in online settings, in particular when they have limited information, has important implications for geographically distributed groups enabled through new information technologies. In addition, understanding how people might use information gleaned from technology usage, such as personal privacy settings, particularly in the absence of other information, has implications for understanding many potential situations that arise in pervasive computing environments.

1 Introduction

Trust plays an important role in group cooperation. During periods of broad social change, however, the basis of trust, and therefore the ability for social actors to engage in exchange and cooperation, can be disrupted. For example, during the Industrial Revolution, increased contact and interaction between unknown individuals as a result of immigration to cities amplified uncertainty regarding the reliability and trustworthiness of potential exchange partners. [1] Similarly, interaction occurred in new settings and situations, e.g., factories, in which individuals' behavior and outcomes depended on the actions of possibly unknown others. [1,2] Over time, new mechanisms were created to detect, monitor and signal the reliability/trustworthiness of social actors. [1,2] Today, as new information technologies (IT) emerge for communication and exchange that facilitate contact between unknown individuals in novel settings and situations (e.g., chat rooms, social networking websites, distributed work groups), established mechanisms of trust are disrupted or distorted, which can lead to the breakdown of cooperation or to increasing fraud in exchange. Moreover, if new mechanisms for determining or signaling trustworthiness have yet to be established, individuals may use other signals as a basis for trust; signals which may or may not be associated with reliability yet will affect interaction, exchange and cooperation.

Interest in the implications of information technology for trust crosses all of the social science disciplines. [3-11] Social scientists, for example, have explored how interpersonal trust is adapted to the digital environment [6], e.g., the reputation system of *eBay* [7] and reliability in the online encyclopedia *Wikipedia* [12]; how trust affects consumer and other behavior online [8,13-15]; and how institutional trust is managed in pervasive networks by organizations such as firms. [9,16-17] In this paper, we build on these studies to examine how personal privacy information about members of geographically distributed, virtual work groups influences cooperation and privacy behavior in the group. Specifically, we use experimental methods to examine whether people use others' personal privacy settings as signals of trustworthiness within the group that affect cooperative group behavior. In addition, we examine how individual privacy preferences relate to whether individuals cooperate in the group. Understanding how people interact with others in online settings, in particular when they have limited information, has important implications for geographically distributed groups enabled through new information technologies. Furthermore, understanding how people might use information gleaned from technology usage, such as personal privacy settings, particularly in the absence of other information, has implications for understanding many potential situations that arise in pervasive computing environments.

In section 2 we briefly discuss the social science literature on trust, as well as previous research on trust in work groups, including distributed teams facilitated with IT support. Section 3 describes the social experiment used for the study and the characteristics of the subjects. Section 4 presents results, while section 5 discusses the findings, limitations and implications.

2 Trust, Cooperation and Work Teams

Trust is a term to describe positive expectations of one actor toward another for some specific action. When we say that A trusts B, we typically mean that A trusts B to do X [18-19], and so A may take some action Y (e.g., lending \$10, sharing information, contributing to a joint project) in which A is now vulnerable to losing Y depending on the behavior of B (i.e., doing X or not). According to Edward Lorenz [20], A's behavior Y based on trust in B "consists in action that (1) increases one's vulnerability to another whose behavior is not under one's control, and (2) takes place in a situation where the penalty suffered if the trust is abused would lead one to regret the action." Snijders [21] specifies that an actor's vulnerability in trust relationships is based on uncertainty about another actor's "disposition or preferences" for cooperation, not his or her abilities to cooperate. [2, 22-23]

A's positive expectations about B, i.e., A's perception of B's trustworthiness, may be based on a number of different reasons, including: (1) A's past experience with B in general or on X specifically; (2) A's relationship with B; (3) A's knowledge of B's reputation from other actors; (4) A's knowledge of B's incentives to do X in response to some other third-party. Social actors may have very limited information, however, about potential exchange partners (B) and therefore have limited grounds on which to assess trustworthiness. Indeed, as noted above, during periods of rapid social change, social actors may have much greater contact with unknown others for whom they have limited information. In such cases of limited information, some characteristics or behavior may be perceived by others as signals of trustworthiness, regardless of the association between such characteristics and behavior. For example, Zucker [1] describes how some actors may perceive characteristics such as race or gender as signals of trustworthiness, without the knowledge or intention of the so-characterized actor.

3 Methods

We conducted a social experiment in which undergraduate subjects (n=110) participated in a series of simulated online groups. After responding to a recruitment advertisements and completing an Informed Consent form approved by our Institutional Review Board, undergraduate student subjects completed a brief online pre-survey measuring demographic characteristics, experience in Internet commerce, and attitudes toward taking risks. Upon completing the pre-survey, each subject was assigned a numeric identifier used to anonymously link the pre-survey responses to the experiment results and post-survey responses. Subjects brought their identification number to a laboratory session where they participated in the experiment session that lasted approximately 20 minutes.

In the laboratory session, subjects were told they were members of geographically distributed work teams (with two other individuals not known by the subjects) that were working on developing a proprietary product via a secure (password protected) online project *wiki*. In addition, the team was in competition with other project teams for the best product design, so project development information was even more closely guarded than standard proprietary designs. Given that the project *wikis* contained valuable information, team members were faced with opportunities to sell the password. Subjects made decisions about whether to sell the *wiki* password in a series of six different teams, comprised of different types of group members (described further below), with each team engaged in a competition as described. Subjects earned points based on the outcomes of each team-competition (other group members are simulated so their behavior is determined randomly); points were converted to monetary amounts that subjects were paid upon completing the study. If any team member sold the password, all team members received zero points. If the subject sold the password, s/he was awarded points according to one of two randomly assigned conditions: (1) 4 points, which was less than the potential gain from group success while protecting the password (6 points) or winning the competition (8 points); or (2) 8 points, which was greater than the potential gain from group success while protecting password (6 points) and equal to points from winning the competition (8 points). If the subject chose not to sell the password, s/he was awarded

either 0, 6 or 8 points depending on, respectively, if another group member sold the password (20% chance, randomly determined), if no member sold the password (70% chance, randomly determined), or if no member sold the password *and* the team won the competition (10% chance, randomly determined).

For each team, subjects receive two pieces of information about the other group members: (1) individual's privacy setting on a scale of 1 – 3 (1=Private 2=Moderate and 3=Open), and (2) individual's skill level: Low or High. Subjects were teamed with two (simulated) members and decisions about whether to sell the password were made simultaneously with no interaction allowed. In order to reduce the chance that feedback about outcomes would influence subjects' subsequent decisions, subjects were not told the number of teams they would participate in, or the outcomes for their teams until the end of the experiment.

4 Results

A total of 110 subjects participated in the study, with each subject making choices in 6 rounds (different team configurations) to decide whether to sell the password or not ($n=110*6 = 660$ observations).

Our central research question is whether subjects use others' privacy settings as a signal of trustworthiness, in particular when they have limited information about group members. If so, we would expect to see that subjects behave differently depending on group members' privacy settings. In the experimental setting described here, we tested whether subjects were more or less likely to cooperate (*not* sell the group password) depending on other group members' privacy settings of open, moderate or private, and controlling for other factors, including subjects' own privacy setting (explored further below), group members' skill level, and value of group cooperation (i.e., size of the incentive to sell based on whether value of selling password \geq potential group payoff).

Table 1 shows the percentage of subjects who were willing to sell the password (*not* cooperate) depending on the privacy settings of teammates, after controlling for the effect of subjects' own privacy setting, teammates' skill level, and the value of group cooperation. Subjects were significantly more likely to sell the password when paired with a teammate who was more open, and even more likely to sell when paired with two teammates who were more open. (P-values based on estimates from logistic regression of choice to sell password on teammates' privacy settings, controlling for subject's privacy setting, teammates' skill levels, and size of the incentive to sell, with robust standard errors.) That is, subjects were most likely to sell the password when paired with two open teammates, and least likely to sell the password when paired with two private teammates. These findings suggest that subjects do indeed use other's privacy settings as signals of trustworthiness in conditions of limited information, altering their behavior in ways that indicate they believe others who have more open settings are less trustworthy.

Table 1. Percent willing to sell password by team members' privacy settings

Teammate 2: Privacy Setting	Teammate 1: Privacy Setting			P-value
	OPEN	MODERATE	PRIVATE	
OPEN	58%	45%	---	Teammate 1 P<.001
MODERATE	---	39%	---	
PRIVATE	46%	34%	24%	

Note: P-values based on estimates from logistic regression of choice to sell password on teammates' privacy settings, controlling for subject's privacy setting, teammates' skill levels, and size of the incentive to sell, with robust standard errors.

These findings raise the question as to whether individuals who are more open are indeed less trustworthy, i.e., are subjects with open privacy settings more likely to sell the group password than those with more private settings? Table 2 shows the mean differences (unadjusted) in likelihood of selling the

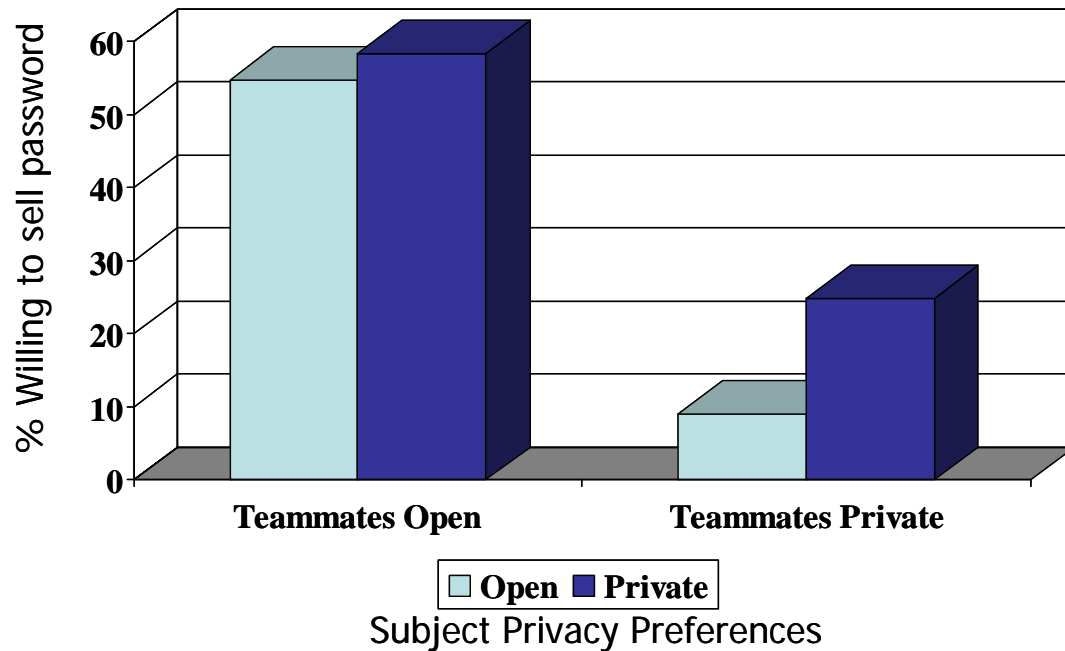
password between subjects with different privacy settings. In contrast to the findings above, subjects who are more private are somewhat *more* likely to sell the password ($p < .10$) than subjects who are open.

Table 2. Percent willing to sell password by subject's privacy setting

Subject's Privacy Setting	% willing to sell password	ANOVA
Private	48	F = 2.56 P < .10
Moderate	40	
Open	36	Bonferroni Private > Open P < .10

We examine this finding further in Figure 1, which shows the interaction between subject's privacy setting (comparing only Open with Private, suppressing Moderate category) and Teammates' privacy settings (comparing only teams in which teammates are either both Open or both Private). Consistent with the findings in Table 1, Figure 1 shows that all subjects are less likely to cooperate with open teammates than with private teammates. Surprisingly, however, when faced with two private teammates, private subjects are less likely to cooperate than open subjects. These preliminary findings suggest that though all subjects appear to cooperate more with private teammates, private subjects are least likely to actually cooperate!

Figure 1. Percent willing to sell password by interaction of Subject privacy setting with teammates' privacy settings



5 Discussion and Conclusion

Trust in groups is affected by the privacy preferences of other members of the group; those with more private settings are more likely to be viewed as trustworthy, and therefore to be trusted. Users' own privacy preferences also appear to matter for trust; more private users appear to be less trustworthy with regard to protecting group privacy and/or to be more distrustful of others than are users with more open settings.

Users will use privacy settings as "signals" of trustworthy behavior in groups, but those signals are not necessarily accurately associated with trustworthy behavior. Managing privacy in online groups or the "commons" enabled by new information technologies may be more difficult than expected, as it does not appear to be a simple aggregate of individual preferences or behavior. In short, social dynamics and social context are likely to matter as much (or more) for ensuring privacy and security than technology alone.

Acknowledgements This research was supported by the Institute for Security Technology Studies at Dartmouth College under grant 2005-DD-BX-1091 awarded by the US Bureau of Justice Assistance. Points of view or opinions in this document are those of the authors and do not represent the official position or policies of the US Department of Justice or of other sponsors. We thank Clare Fortune-Agan, Linda Lomelino and Sara del Nido for research assistance.

References

1. Zucker, LG (1986) Production of Trust: Institutional Sources of Economic Structure, 1840-1920. In Staw BM and Cummings LL (eds) *Research in Organizational Behavior*, volume 8. JAI Press Inc., Greenwich, CT.
2. Shapiro, SP (1987) The Social Control of Impersonal Trust. *Am J Sociology* 93:623-58.

3. Baye, M (2002) Special Issue on The economics of the Internet and e-commerce. *Advances in Applied Microeconomics*. Volume 11.
4. Camp, LJ (2000) *Trust and Risk in Internet Commerce*. The MIT Press, Cambridge, MA.
5. Falcone R, Singh M, Tan YH (2001) *Trust in Cyber-societies*. Springer, Berlin.
6. Friedman E, Resnick P (2001) The Social Cost of Cheap Pseudonyms. *J Econ & Management Str* 10:173-199.
7. Kollock P, (1999) The Production of Trust in Online Markets. *Advance in Group Processes* 16:99-123.
8. Lunn R, Suman M (2002) Experience and Trust in Online Shopping. In: Wellman B, Haythornthwaite C (eds) *The Internet in Everyday Life*. Blackwell, Oxford UK.
9. Osterwalder D, (2001) Trust through evaluation and certification? *Soc Sc Comp Rev* 19:32-46.
10. Sambamurthy V, Jarvenpaa S, (eds) (2002). Special Issue on 'Trust in the Digital Economy.' *J Strategic Inf Sys* 11: 183-346.
11. Shapiro C, Varian H, (1999). *Information Rules. A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston.
12. Anthony D, Smith SW, , Williamson T, (Forthcoming) Reputation and Reliability in Collective Goods: The case of the online encyclopedia Wikipedia. *Rationality & Society*.
13. Belanger F, Hiller J, Smith W, (2002) Trustworthiness in electronic commerce: the role of privacy, security and site attributes. *J Strategic Info Sys* 11: 245-270.
14. Castells M, (2001) *The Internet Galaxy*. Oxford University Press, Oxford.
15. Wellman B, Haythornthwaite C, (eds) (2002) *The Internet in Everyday Life*. Blackwell, Oxford.
16. Anthony D, Lewis E, Who do you call in a crisis? Reliability and Capability for Trusted Communication. Conference paper at the International Sunbelt Social Networks Conference XXIV, May 2004, Slovenia.
17. Knights D, Noble F, Vurdubakis T, Willmott H, (2001) Chasing Shadows: Control, Virtuality and the Production of Trust. *Org Studies* 22:311-336.
18. Hardin R (1991)
19. Hardin R (2001)
20. Lorenz E (1988)
21. Snijders C (1996) *Trust and Commitments*. Interuniversity Center for Social Science Theory and Methodology, Gronigen, Netherlands.
22. Luhmann N (1988) Familiarity, Confidence, Trust: Problems and Alternatives. In: Gambetta D (ed) *Trust Making and Breaking Cooperative Relations*. Basil Blackwell, New York.
23. Gambetta D (1988) Can We Trust Trust? In: Gambetta D (ed) *Trust Making and Breaking Cooperative Relations*. Basil Blackwell, New York.