Information & Computer Security

Privacy as a secondary goal problem: an experiment examining control

Thomas Hughes-Roberts

Privacy as a secondary goal problem: an experiment examining control

Thomas Hughes-Roberts
School of Science and Technology, Nottingham Trent University, Nottingham, UK

Received 22 October 2014
Revised 22 October 2014
Accepted 9 January 2015

Abstract

Purpose – The purpose of this paper is to report on results of an investigation into the impact of adding privacy salient information (defined through the theory of planned behaviour) into the user interface (UI) of a faux social network.

Design/methodology/approach – Participants were asked to create their profiles on a new social network specifically for Nottingham Trent University students by answering a series of questions that vary in the sensitivity of personal information requested. A treatment is designed that allows participants to review their answers and make amendments based on suggestions from the treatment. A dynamic privacy score that improves as amendments are made is designed to encourage privacy-oriented behaviour. Results from the treatment group are compared to a control group.

Findings – Participants within the treatment group disclosed less than those in the control with statistical significance. The more sensitive questions in particular were answered less when compared to the control, suggesting that participants were making more privacy-conscious decisions.

Practical implications – Work within this paper suggests that simple UI changes can promote more privacy-conscious behaviour. These simple changes could provide a low-cost method to ensuring the Internet safety of a wide range of users.

Originality/value – This study provides a clear definition of privacy salient UI mechanisms based on a well-established theory of behaviour and examines their potential effect on end-users through a novel experiment. Results and methods from this work can enable researchers to better understand privacy behaviour.

Keywords Information disclosure, End-user computing, Computer privacy, Human physiology

Paper type Research paper

1. Introduction

End-users of social networks routinely disclose sensitive information about themselves despite stating a high level of concern for their privacy in such services; a phenomenon known as the privacy paradox (Acquisti and Gross, 2006). Privacy itself in social networks has been described as a secondary goal problem (Bonneau *et al.*, 2009), that is, it is not considered during interaction where the focus is on achieving other goals which may well be in opposition to the idea of privacy. Users within social networks may therefore forget about the impact sharing information may have when they disclose it or fail to protect it through using appropriate privacy settings when it has been disclosed.

This paper proposes that reminding users of their privacy at the point of interaction through the user interface (UI) could produce more privacy-conscious behaviour. Specifically, this work suggests that providing a clearer method of

Information & Computer Security
Vol. 23 No. 4, 2015
pp. 382-393
© Emerald Group Publishing Limited
2056-4961
DOI 10.1108/ICS-10-2014-0068

control with an obvious focus on privacy protection will elicit pro-privacy behaviour and explores the persuasiveness of such mechanism. To accomplish this, the theory of planned behaviour (TPB) (Ajzen 1991) is used to explore the underlying motivators behind behavioural intentions and action. In particular, the salient property it defines as "Perceived Control" is used to inform a UI addition that makes it easier to identify and modify disclosed information with the view to making privacy a clearer goal of system interaction.

This paper will justify such an approach by reviewing related literature to define the privacy problem and examine the TPB as a potential method to designing solutions. The methodology section defines how the "Perceived Control" property of the TPB can be adapted into a UI element aimed at improving privacy behaviour and proposes an experiment to test the effect such an element may have on end-users. Users are asked to create their profiles on a new social network by answering personal questions where the treatment highlights more sensitive questions to allow more control over what is answered. Results from this experiment are compared to a control group and discussed using exit survey and focus group data.

2. Literature review

Literature has proposed numerous potential causes of the privacy paradox, including a low level of technical skill in users (Kolter and Pernul, 2009), a lack of awareness of privacy issues (Miller *et al.*, 2011) and the design of the social network itself (Fogg and Iizawa, 2008; Livingstone, 2008). Indeed, privacy as a concept is a highly complex issue that differs from individual to individual (Rosenblum, 2007). Given that users appear to desire a level of concern for the privacy and only they can be aware of their particular privacy context at any given time, an argument can be made that users must be empowered to enact their own privacy needs.

Privacy as a secondary goal for users has also been proposed as a potential cause of poor privacy behaviour (Bonneau *et al.*, 2009). That is, users do not consider their privacy at the point of interaction, as it may conflict with their personal reasons for using the system (or is simply not as important). For example, the aim of building social capital with friends through using Facebook would suffer if the user was privacy oriented. The lack of privacy salience embedded into the design of the social network has been cited as a potential reason for the lack of privacy consideration shown by users (Houghton and Joinson, 2010). However, what privacy salience is and what it looks like is unclear and requires further definition.

Social networks have also been described as a persuasive technology, influencing and changing users' behavioural habits through their use of the system (Fogg, 2009). The addition of privacy salient information could therefore persuade users to act in accordance with stronger privacy concerns. Indeed, work exploring the privacy paradox has found that providing users with "counter-arguments" adjusted their privacy perceptions and level of concern (Ur, 2014). However, this was not done through a live UI so the question is raised: can the UI remind or inform people of their privacy and improve privacy behaviour? If so, what should the UI features look like and how can they be designed to be persuasive? The act of disclosure could be considered an emergent behaviour of privacy and therefore perhaps theories of behavioural change could provide a means of defining key behavioural motivators.

One such theory of behavioural change is the TPB (Figure 1) which defines three "salient beliefs" that influence an individual's intention and behaviour (Ajzen, 1991).

Each of these aspects of salience could provide a basis for altering the UI such that privacy becomes part of the behavioural intention and action and is therefore part of the goal rather than a secondary goal. As such, this paper proposes that the UI is ideally suited to empower users to protect their privacy by adding privacy salient information and mechanisms to it, and it has been suggested that there lacks critical focus on the role of the UI (Masiello, 2009). Furthermore, behavioural psychology suggests that behaviour is a reaction environmental stimulus (Breakwell, 2006), and the UI in a social network is the environment with which users react and interact with. Wider research has found that the way in which information is presented to users and the control options surrounding can influence the amount of disclosure users exhibit (Brandimarte *et al.*, 2012).

From the TPB, *behavioural attitude* suggests that behaviour is influenced by our knowledge and perception of the consequences associated with certain acts. Hence, the UI could provide prompts to remind users of the risks of information disclosure. Indeed, an element of persuasion has been defined as suggestion (Fogg, 2003) where interventions are set to appear "at the right time" with the right information thus raising awareness of privacy issues. Furthermore, this could seek to remind or inform users of privacy in an attempt to improve their awareness: which has also been cited as a factor in contributing to the privacy paradox (John *et al.*, 2009).

Subjective norms suggest that behaviour is influenced by the thoughts and actions of those around us. The theory of social capital (Portes, 1998), for example, would see disclosure as the act of strengthening social ties with peers. The UI could also be used to deliver advice and guidance either from peers or expert users to aid in the decision-making process of privacy behaviour.

Finally, the *perceived control* aspect deals with the perception of how easy a behaviour is to perform and how easy it actually is to enact. The design of technology may make it easy to disclose private information and not so easy to protect it. Indeed, system complexity has been put forward as a cause of poor privacy behaviour (Fang and Lefevre, 2010, Brandimarte *et al.* 2012). In terms of the paradox, users may feel that the identification and protection of private information is simple, yet this perception may not relate to reality resulting in poor privacy decisions. This is somewhat in line with

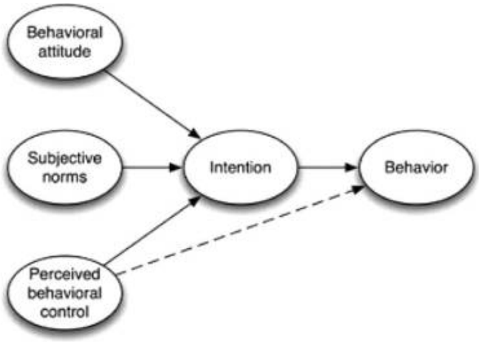

Source: Ajzen (1991)

Figure 1.
Theory of planned behaviour

tunnelling, which has been proposed as another persuasion strategy where goal-driven design aims reduce uncertainty by leading users through interaction (Fogg, 2003), i.e. the role in interaction plays in a clearly defined goal is important.

Each of these salient properties has been adapted into a UI element aimed at influencing privacy behaviour in a formal experiment. This paper presents the findings from the perceived control treatment in comparison to a control group and is described in the following section.

3. Methodology

3.1 Hypothesis and experiment

Based on the perceived control element of the TPB, an experimental treatment is designed to explore the following hypotheses:

> *H1*. A UI that aids users in identifying and controlling sensitive information will influence user behaviour and decrease the amount of sensitive information they give.

This experiment asks users to "sign-up" to a new social network created for Nottingham Trent University students and create their profiles on it. The experiment is designed to mimic Facebook in appearance to promote the ecological validity of the experiment (Lew *et al.*, 2011) and place it in a clear context; see example in Figure 2. The red asterisks illustrate that the other questions asked of the participants are optional and that they do not have to answer.

This account creation process asks participants a series of questions that range in their potential sensitivity and privacy invasiveness similar to work by Brandimarte *et al.* (2012). In total, 33 questions are asked of the participants during this account creation process; the total amount of questions answered is used to test *H1* by comparing the results in the treatment group to the control group. These questions have a variety of input types including text boxes, drop-down menus and yes/no checkboxes (e.g. have you ever pirated media?). Participants are informed that the questions are intended to populate their profile with information and create a network of like-minded individuals, and the more they disclose, the more accurate their resulting network will be. The control group, in total, traverses three screens, the introductory screen as shown in Figure 1, the question bank that "builds" their profile (Figure 3) and, finally, a screen

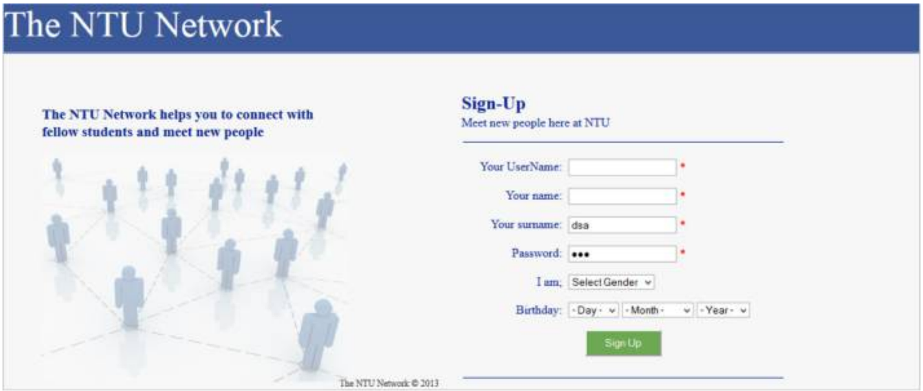

The NTU Network

The NTU Network helps you to connect with fellow students and meet new people

Sign-Up
Meet new people here at NTU

Your UserName: *
Your name: *
Your surname: dsa *
Password: ••• *
I am: Select Gender
Birthday: -Day- -Month- -Year-

The NTU Network © 2013

Figure 2. Experiment homepage

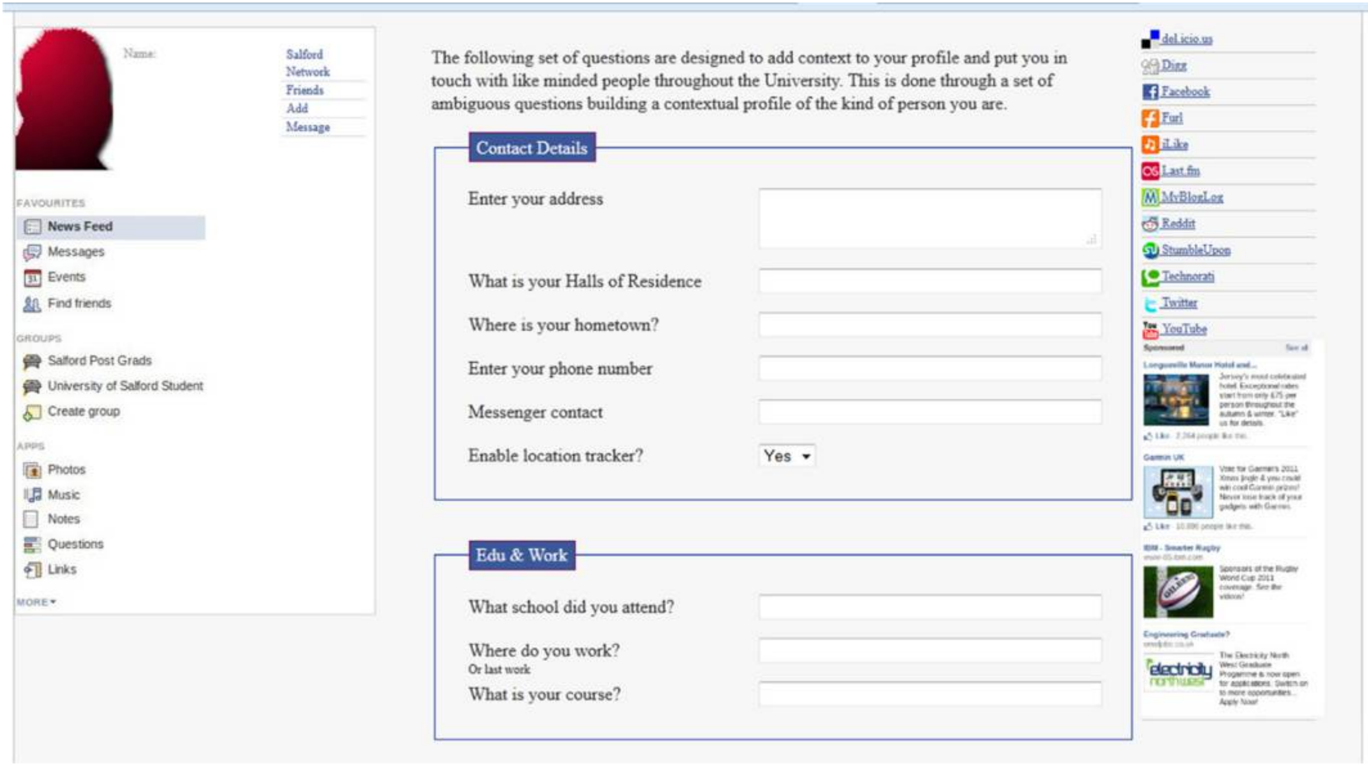

The following set of questions are designed to add context to your profile and put you in touch with like minded people throughout the University. This is done through a set of ambiguous questions building a contextual profile of the kind of person you are.

Contact Details

Enter your address

What is your Halls of Residence

Where is your hometown?

Enter your phone number

Messenger contact

Enable location tracker? Yes ▾

Edu & Work

What school did you attend?

Where do you work?
Or last work

What is your course?

Figure 3.
Profile builder page

to apply privacy settings to their new accounts (the purpose of this paper is to focus on the questions answered rather than the settings applied).

3.2 Treatment design

The treatment adds a series of review screens after each form submission from the participants. This review screen is intended to provide an opportunity to examine what data they have entered outside of the context of the social network, identify potentially sensitive information they have inputted and modify where they feel is necessary. It therefore aims to place the information submitted into a privacy-oriented context, make it easier to identify sensitive information and allow greater control over it through the review. As such, two levels of disclosure were recorded for the treatment group participants: one prior to the review offered by the treatment and one after (Figure 4).

This screen aids in the identification of potentially sensitive information that has been submitted through a variety of dynamic UI elements. Each piece of data has a rating that is hidden if the field is blank ("Delete to improve P-Score"); should a participant remove information, this rating dynamically turns off to emphasise the impact of the interaction and to demonstrate the tangible results of the control group (making risk management more obvious). This rating of sensitivity categorises the requested data into three brackets: low risk (green), medium risk (yellow) and high risk (red) (Knijnenburg *et al.*, 2013). Green category data items carry little privacy risk and include such questions dealing with favourite films, music, etc. Yellow data items have more potential impact and include questions dealing with political ideology, religion, etc. Finally, the red questions deal with potentially highly sensitive data items including address, drinking habits, etc. It is important to note that these categorisations are not clear-cut definitions of privacy invasiveness but instead are intended to provide a general guide for participants to use to inform their decisions when considering *their* privacy.

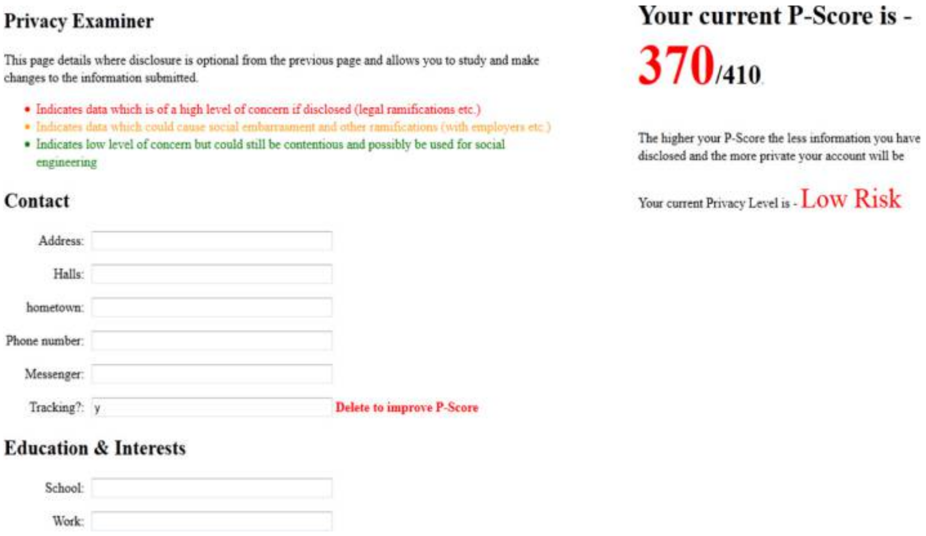

Privacy Examiner

This page details where disclosure is optional from the previous page and allows you to study and make changes to the information submitted.

- Indicates data which is of a high level of concern if disclosed (legal ramifications etc.)
- Indicates data which could cause social embarrassment and other ramifications (with employers etc.)
- Indicates low level of concern but could still be contentious and possibly be used for social engineering

Contact

Address:

Halls:

hometown:

Phone number:

Messenger:

Tracking?: y **Delete to improve P-Score**

Education & Interests

School:

Work:

Your current P-Score is -

370/410

The higher your P-Score the less information you have disclosed and the more private your account will be

Your current Privacy Level is - Low Risk

Figure 4.
Salient review treatment

A live "P-Score" is provided to assess the level of privacy risk there potentially is based on how questions have been filled in. These dynamic elements aim to not only aid participants in identifying their sensitive information but also to promote increased control over it through a demonstration of the impact of interaction. Furthermore, the goal of privacy is made the centre of interaction to determine the degree to which participants may be affected when considering extra information during their account creation. As participants will interact with two screens in this treatment group, two measures of answered questions are provided: before and after review of information.

It is important to note that this experiment methodology will only record instances of disclosure in terms of the presence of data and not the integrity of the data, i.e. participants may choose to lie as a means of protecting their privacy. However, the presence of any information (truth or not) can still be used to form an online identity and perceived by others when viewed in a myriad of ways. Furthermore, research has suggested that users of social networks tend to "stretch the truth" rather than outright lie (Amichai-Hamburger *et al.*, 2010).

Following the experiment, participants in the treatment group have access to an exit survey and took part in a focus group to assess their perceptions of the treatment. This exit survey aims to explore the degree to which participants felt the treatment was useful and specifically provides a number of statements for participants to state their level of agreement on a five-point Likert scale (strongly agree to strongly disagree):

- I found the privacy information helpful.
- The privacy information helped to select what to fill in.
- I believe the privacy information would be beneficial in the long run.
- I acted differently due to its presence.

Participants were sampled from Nottingham Trent University's information systems course (convenience sampling) and were approached in scheduled lab sessions. They were asked if they would like to sign-up to a new social network specifically for the University. They were then randomly assigned to either the control or a treatment group. In total, 20 (16 male and 4 female) participants were gained for the control group and 21 (17 male and 4 female) for the perceived control treatment group. It is noted that the participants were predominantly male as a result of the sampling technique used and as such may not be considered representative of a social network systems population as a whole.

4. Results

Table I provides an overview of the results from the experiment in terms of the total amount of disclosure (per cent of questions answered) exhibited in the experiment

Table I.
Total disclosure in groups

Group	No. of participants	Total % of questions answered
Control	20	82
PC1	21	74
PC2	21	48

groups (PC1 and 2 represent the perceived control group before and after data review, respectively).

An initial review of these results shows that there is a percentage decrease when compared to the control for PC2 (after data review). Indeed, this is a statistically significant result with a Mann–Whitney U $p < 0.0001$. The reduction for PC1, however, is not statistically significant ($p = 0.244$); this would suggest that the treatment influenced participants to review and amend their submitted data. If a participant is considering their privacy, then it is reasonable to assume that disclosure will be the least in more sensitive data categories, as the treatment should aid in identifying which these are (a breakdown of which is in Table II).

It would appear here that upon review of their data, participants did disclose less in the more sensitive data categories with statistical significance compare to the control (Table III).

Again, this table would suggest that the greatest effect of the treatment appeared in the more sensitive data categories; although, the green category was also reduced with statistical significance. *H1*, therefore, would appear to be true based on these results. However, further exploration is required to examine if participants are enacting their own privacy desires or those that they feel are persuaded by the system. Literature has noted that user tend to forsake sub-goals in pursuit of a perceived main goal provided by the system (Jacko and Sears, 2003). It could therefore be that participants are being persuaded to be more private than they desire to be.

Table IV details the exit-survey results. Participants who selected either strongly or somewhat agree are said to have agreed with the statement in questions and likewise for the disagree range of responses.

Group	% of “Green” questions answered	% of “Yellow” questions answered	% of “Red” questions answered
Control	83	82	81
PC1	77	73	73
PC2	65	41	37

Table II. Spread of disclosure across suggested sensitivities

Group	Test	Green	Yellow	Red
PC1	Mann–Whitney	0.242	0.192	0.175
PC2	Mann–Whitney	0.027	< 0.0001	< 0.0001

Table III. Statistical tests comparing each sensitivity category to the control group

Statement	Agreed (%)	Neutral (%)	Disagreed (%)
I found the privacy information helpful	58	32	10
The privacy information helped me answer	63	32	5
I believe the privacy information would be beneficial in the long run	42	47	12
I acted differently due to the privacy information	42	37	21

Table IV. Exit-survey results’ summary

5. Discussion

The literature review identified that privacy can suffer from a secondary goal problem and hence is not considered during the implementation of other pre-defined goals. In this experiment, the goal was to create an account on a new social network and this is achieved through the completion of the smaller sub-goals of answering a series of questions to complete the profile. Certainly, disclosure within the control group seems to be fairly high with a fairly even spread across the sensitivity categories defined by the treatment. This would suggest that their privacy is not being considered during the completion of the task set before them. Indeed, when asked post experiment about why they behaved in such a way the response was: *I do not really know, I just answered the questions* and *I did not think, now I would have left some questions out*. These responses are similar to wider work (Strater and Lipford, 2008) and are indicative of a lack of privacy thought during the interaction. Some participants described themselves as "completionists" and wanted to answer each question they could; indeed, wider research suggests that disclosing information about the self is intrinsically rewarding and potentially addictive (Tamir and Mitchell, 2012).

The decreases in the more sensitive information categories in the treatment group would suggest that participants were making a selection of what to disclose based on the potential privacy invasiveness of the information asked of them. However, the question remains as to whether or not participants are enacting *their* privacy preferences or are responding the potentially persuasive goal of the treatment (that places the interaction squarely within a privacy context). That is, the system now wants them to not answer questions and improve the "P-Score", evidence for which could be found in the reduction in the green category of questions (with statistical significance). However, some questions in the green may indeed be sensitive to some people due to the personal and contextual nature of privacy (Masiello, 2009).

Interestingly, post-experiment results explore this further, as participant's responses suggested that the dynamic "Privacy Score" was the most influential in convincing them to remove submitted information: *I wanted to get a low score, it was like a game*. Also, when asked if they changed response: *yes, it seemed to want me to*. The score would therefore seem to provide participants with a real-time reaction to their interaction that gave them something to aim for with a real and tangible goal. However, the main reason for removing information may be attributed to gaining a low score and not as a result of thinking about their own privacy needs; although, disclosure was lessened in the more sensitive categories. Privacy would therefore seem to be a more prevalent goal of the interaction where the treatment is present. However, the treatment may have made privacy the *primary* goal of interaction and therefore was persuasive in the same vein as a social network may be with gathering participant information.

From the exit survey, the majority of participants agreed with the statements that the treatment was useful in aiding their selection. However, only 42 per cent believed that they acted differently due to its presence; this is despite the change in disclosure when the treatment is introduced to the group (i.e. there is strong evidence that they were effected). Participants may be unwilling to admit the extent to which they were influenced by the treatment, should it be persuasive in convincing them to act in accordance with strong privacy recommendations. Indeed, there is evidence in the literature that users tend to downplay the effect of perceived counterintuitive behaviour on themselves but do perceive it to be persuasive to others (Debatin *et al.*, 2009).

6. Conclusions and further work

This paper has presented the results of an experiment based on the perceived control aspect of the TPB and explored a means of including the idea of privacy as part of the interaction with social networks. This treatment aimed to introduce privacy into the goal of interaction in attempt to provide a solution to privacy as a secondary goal problem by making the identification of sensitive information more salient and the impact of control over that information more obvious. Participants in the treatment group did disclose less than the control, and this disclosure was specifically reduced in the more sensitive categories of questions. Such a dynamic score could be added to real social networks or other software through browser extensions or the use of APIs to aid users in the use of such systems. It would appear that seeing the immediate consequence of behaviour through a changing score clearly defines a goal and provides direction to users; this may be very useful in systems with a wide user base with variable levels of technical skill. However, the extent to which participants enacted their privacy preferences is unclear due to the potential persuasion the treatment may have introduced; that is, privacy may have become the *primary* goal of interaction and other goals may have suffered.

This experiment does show that one form of salience can be particularly effective in persuading users at the point of interaction through dynamic UI elements that instantly show the tangible effect of an interaction. However, the experiments took place in a controlled context and as such do not model the real-world setting of privacy and social network behaviour. Hence, participants may have acted according to the perceived aim of the experiment explaining the extensive reduction in the more sensitive data categories. To verify the results here, the UI elements should be placed in a real-world setting to examine the potential effects in an actual context. For example, in Facebook, participants may disclose due to their own pre-defined goals rather than goals defined by the system (account creation, reduce privacy score, etc.). Would such UI elements influence these more personal goals?

Ultimately, this experiment has examined if privacy can be made a salient goal of interaction in an attempt to tackle the potential secondary goal problem it has. The treatment designed here placed privacy squarely into the interactive experience of the participant and made it a clear, tangible goal of system use. UI elements described demonstrate how privacy goals can be made to be a more persuasive part of end-user needs and directly influence their behaviour within the system.

References

Acquisti, A. and Gross, R. (2006), "Imagined communities: awareness, information sharing, and privacy on the facebook", *Proceedings of the 6th Workshop on Privacy Enhancing Technologies*.

Ajzen, I. (1991), "The theory of planned behaviour", *Organizational Behaviour and Human Decision Processes*, Vol. 50, pp. 179-211.

Amichai-Hamburger, Y. and Gideon, V. (2010), "Social network use and personality", *Computers in Human Behavior*, Vol. 26 No. 6, pp. 1289-1295.

Bonneau, J., Anderson, J. and Church, L. (2009), "Privacy suites: shared privacy for social networks", *5th Symposium on Usable Privacy and Security*, New York, NY.

Brandimarte, M., Acquisti, A. and Loewenstien, G. (2012), *Misplaced Confidences: Privacy and the Control Paradox*, Workshop on the Economics of Information Security, Harvard.

Breakwell, G.M. (2006), *Research Methods in Psychology*, Sage Publications, Oxford.

Debatin, B., Lovejoy, J.P., Horn, A.K. and Hughes, B.N. (2009), "Facebook and online privacy: attitudes, behaviors, and unintended consequences", *Journal of Computer-Mediated Communication*, Vol. 15 No. 1, pp. 83-108.

Fang, L. and LeFevre, K. (2010), "Privacy wizards for social networking sites", *Proceedings of the 19th International Conference on World wide Web*, *New York, NY*.

Fogg, B.J. (2003), *Persuasive Technology: Using Computers to Change what We Think and Do*, Morgan Kaufmann, San Francisco, CA.

Fogg, B.J. (2009), *The Behaviour Grid: 35 Ways Behaviour Can Change*, PERSUASIVE, Clairemont, CA.

Fogg, B.J. and Iizawa, D. (2008), in Oinas-Kukkonen, H. *et al.* (Eds), *Online Persuasion in Facebook and Mixi: A Cross-Cultural Comparison*, PERSUASIVE, Berlin, pp. 35-46.

Houghton, D.J. and Joinson, A. (2010), "Privacy, social network sites, and social relations", *Journal of Technology in Human Services*, Vol. 28 Nos 1/2, pp. 74-94.

Jacko, J.A. and A. Sears (2003), *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies, and Emerging Applications*, Lawrence Erlbaum and Associates, Mahwah, NJ.

John, L., Alessandro, A. and George, L. (2009), "The best of strangers: context-dependent willingness to divulge personal information", Social Science Research Network.

Knijnenburg, B.P., Kobsa, A. and Jin, H. (2013), "Dimensionality of information disclosure behavior", *International Journal of Human-Computer Studies*, Vol. 71 No. 12, pp. 1144-1162.

Kolter, J. and Pernul, G. (2009), "Generating user-understandable privacy preferences", *International Conference on Availability, Reliability and Security*, *Fukuoka*, pp. 299-306.

Lew, L., Nguyen, T., Messing, S. and Westwood, S. (2011), "Of course I Wouldnt do that in real life: advancing the arguments for increasing realism in HCI Experiments", *Computer Human Interaction*, Association for Computing Machinery, Vancouver, BC.

Livingstone, S. (2008), "Taking risky opportunities in youthful content creation: teenagers' use of social networking sites for intimacy, privacy and self-expression", *New Media and Society*, Vol. 10 No. 3, pp. 393-411.

Masiello, B. (2009), "Deconstructing the privacy experience", *IEEE Security and Privacy*, Vol. 7 No. 4, pp. 68-70.

Miller, R.E., Salmona, M. and Melton, J. (2011), "Students and social networking site: a model of inappropriate posting", *Proceedings of the Southern Association for Information Systems Conference*, *Atlanta*.

Portes, A. (1998), "Social capital: its origins and applications in modern sociology", *Annual Review of Sociology*, Vol. 24, pp. 1-24.

Rosenblum, D. (2007), "What anyone can know: the privacy risks of social networking", *IEEE Security and Privacy*, Vol. 5 No. 3, pp. 40-49.

Strater, K. and Lipford, H.R. (2008), "Strategies and struggles with privacy in an online social networking community", *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction*, British Computing Society, Swinton.

Tamir, D.I. and Mitchell, J.P. (2012), "Disclosing information about the self is intrinsically rewarding", *Proceedings of the National Academy of Sciences*, Vol. 109 No. 21, pp. 8038-8043.

Ur, B., Kelley, P.G., Komanduri, S., Lee, J., Maass, M., Mazurek, M.L., Passaro, T., Shay, R., Vidas, T., Bauer, L., Christin, N. and Cranor, L.F. (2014), "How does your password measure up? The effect of strength meters on password creation", *Proceeding Security 12 Proceedings of the 21st USENIX Conference on Security Symposium*, *Berkeley, CA*.

Corresponding author
Thomas Hughes-Roberts can be contacted at: thomas.hughesroberts@ntu.ac.uk