Trust, Transparency & Control in Inferred User Interest Models

Sebastian Schnorf

Google

Zurich, Switzerland

sebschnorf@google.com

Nikhil Sharma

Google

Mountain View, CA

nikhilsh@google.com

Martin Ortlieb

Google

Zurich, Switzerland

mortlieb@google.com

Abstract

This paper explores the importance of transparency and control to users in the context of inferred user interests. More specifically, we illustrate the association between various levels of control the users have on their inferred interests and users' trust in organizations that provide corresponding content. Our results indicate that users value transparency and control very differently. We segment users in two groups, one who states to not care about their personal interest model and another group that desires some level of control. We found substantial differences in trust impact between segments, depending on actual control option provided.

Author Keywords

Trust; Transparency; Control; Personalization; Privacy

Introduction

In the context of privacy enhancing technologies (PETs), offering transparency and control is often mentioned as a strategy to mitigate between users' privacy concerns and data practices of Internet companies. Previous research (e.g. [1], [2]) has shown, that people are more likely to share information if they feel they have overview knowledge of personal data and are able to act on data controls. If trust is established through such means, users are ready to share more online [3] and vice versa [4].

Offering transparency and control to users involves several challenges for companies. The simple display of all personal data and eventual behavioral traces available can be an overwhelming and anxiety producing experience for users. When making this information progressively transparent without the right level of control, for example in the context of an account profile, users can feel disempowered - they might react with withdrawal. Offering control to users without experiencing the benefit of providing personal data, such as more relevant content recommendations, can also lead to unwanted outcomes. Essentially, no one likes to be out of control, so as soon as this becomes apparent or is perceived, users will either execute control or go to places where they have these options. Recent research also shows that technologies that make individuals feel more in control over the release of personal information may have the unintended consequence of eliciting greater disclosure of sensitive information [1]. Giving users control may be an essential step, but, to some degree, it may not be enough on its own. Instead of maximizing control, the solution to mitigate between users' privacy concerns and data practices of Internet companies might be more about providing the right form of transparency and control, in the right situation.

The concept of trust is an extensively studied concept. As denoted by Luhmann [5], trust is a social mechanism for reducing complexity. Transposing this to the world of products we could argue that the cumulative experience with a product or brand leads to confidence [6]. In the realm of online services this could mean confidence in a company's practices such as never selling personal data to any third party. Trust online is a widely studied area of research and of

particular interest due to the absence of real life signals (e.g. [7], [8], [9]): trust needs to be established through mediated interaction. The challenge for companies dealing with personal data is how to setup and maintain the right level and form of data controls for users, if they want to establish user trust.

The major contribution of this research is that it shows how trust, transparency, and control are intertwined. Furthermore, we show how one can explore such aspects with a novel survey approach. First, we wanted to see how important, in general, control is to users in the context of inferred user interests (IUIs). IUIs contain information about a user's preference for certain content based on explicitly stated or behavioral inferred traces. Second, we wanted to investigate the association of control and users' trust in Internet companies. Finally, we wanted to correlate different levels of control with overall notions of trust.

Method
Above stated research questions were examined employing online surveys. In the following section we will briefly describe our tool and outline some limitations of our approach.

Online Surveys
Google Consumer Surveys (GCS) [10] provide a new method for performing probability based internet surveying which produces timely and cost-effective results while still maintaining much of the accuracy of pre-existing surveying techniques. GCS presents one or two questions to users, providing them access to content that is behind a pay-wall and might not be otherwise available to them for free. A comparison found GCS to be more accurate than both the

probability and non-probability based Internet panels in three separate measures: average absolute error (distance from the benchmark), largest absolute error, and percent of responses within 3.5 percentage points of the benchmark [11].

We ran our surveys in December 2012 with a sample approximately representative of the US Internet population (see [11] for more in-depth discussion on representation of GCSs). We fielded our questions until 200 people responded to any of the subsequent options to control IUIs (see Fig. 1, right hand side). We ran one survey (N=1982) with conditioning on answer 1.1 and a second survey (N= 2038) piping options 1.2-1.4.
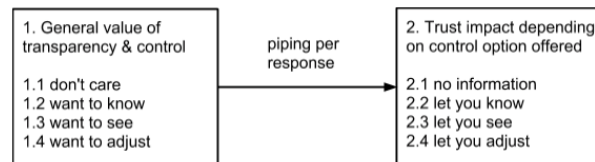
Figure 1. Screener and follow-up conditions in our survey

In order to keep this survey short, we did not prompt an extensive trust scale. We directly asked participants how each of the given options in Fig. 1 would affect their trust level ("How will this change your trust in the company?").

Limitations of our approach
There are several limitations to this setup that we want to mention before we turn to results. As with every survey, we measure reported and not actual behavior. This is particularly relevant for this study as in the realm of privacy; discrepancies between attitudinal and actual behavior can be quite significant (e.g. [12]).

Further, respondents in our study could not factor in the cost of managing transparency and control over their IUIs in that users need to imagine an actual situation that they may well have never been in. It is also worth to note that, in this study, we gauge the impact of controlling IUIs in a centralized form, such as on an account's profile page. This form of control differs from others such as with an app at installation time or more product-embedded experiences when seeing recommended content. Finally, language in surveys has the potential of impacting findings. For instance, the term "interest profile" as well as the various control options might be difficult to understand and, depending on the individual user, have different connotations.

Results
We present our findings along our three goals of this study. We start by demonstrating the overall importance of transparency and control to users, and then we show the association of control and trust. Finally, we illustrate the relationships between different levels of control with overall notions of trust.
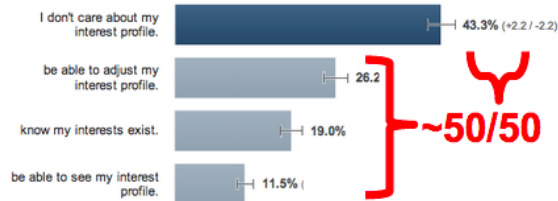
Importance of transparency and control
The study's first research goal was to substantiate evidence for the general value of transparency and control to users. We asked users about the importance of various levels of control over an interest model about them.

As seen in Fig. 2 below, there is more than one dimension in response to this question. We could split the users in two groups of around 50% each: the don't care users, who state not to care about an interest model about them at all and the care users, who want some level of control. This general response pattern

(N=1982) was validated through rerunning the same question for any of the subsequent conditions (see Fig. 1). The significance evaluation of the GCS tool is calculated using the Wilson score interval and the confidence level is at 95%.

Figure 2. Users value transparency and control very different - we could split between groups of users who care and users who don't care (n=1982).

As for interpretation, Fig. 2 illustrates that transparency over interest models is not a very salient concern to almost one in two users. We have anecdotal evidence from users interviews, that some people actually expect interest modeling already be the case, but without much detailed knowledge about the actual practices employed.

On the other side of the spectrum, users who want some level of transparency and control, state a strong tendency to desire full control (be able to adjust their interest model). Overall, we expected much more users would want to have full control if this option was offered. We are going to explore these various attitudinal dimensions further from the angle of this broad user segmentation.

Association of control and trust
The second research goal investigated the association of various behavioral controls and users' trust in companies. We asked users: "If an internet company… decides to [options: not provide info about / inform you about the existence of / allow you to view / allow you

to adjust] your interest profile. How would this change your trust in the company?"

Let us first focus on the scenario, when a company would not provide information about their interest modeling practices. In other words, we explored the situation, when people might be aware about interest modeling, but the company would decide to not inform about its practices, hence, not giving transparency to its users. Under this condition, users who don't care (see Fig. 3) reacted balanced: In total, 68 users out of 200 stated their trust would decrease, versus 60 respondents stated their trust level would increase. Overall, this response pattern makes sense as people get what they want.

If provided the opposite option, i.e. the company does not provide transparency to care users (all three options lumped together, see Fig. 2), the group would, as expected, lose trust in the internet company: In total, 103 out of 200 users stated a decrease versus 50 an increase in trust (significant difference with chi-square test). In this scenario care users assume and speculate about the data practices around interest models and this tends to be evaluated negatively by users. One result that is interesting to note is that trust increases, when a company acts entirely secretive, regardless if a user is a care or non-care user. Maybe the explanation for this kind of response behavior is along the line of "what you don't know won't hurt you" [13].

Now we turn to the scenario, when a company would provide some level of control to users. As can be seen in Fig. 3 and 4, there is a different pattern between the groups (the Chi-squared test is significant: $p < 0.05$):

The don't care users in Fig. 3 tend to remain neutral, as can be seen in the peak. One of our hypotheses in this study was that if a company provides higher level of control, users could be scared. As illustrated, users who don't care do not significantly reduce their level of trust, so based on this survey this hypothesis does not hold true.
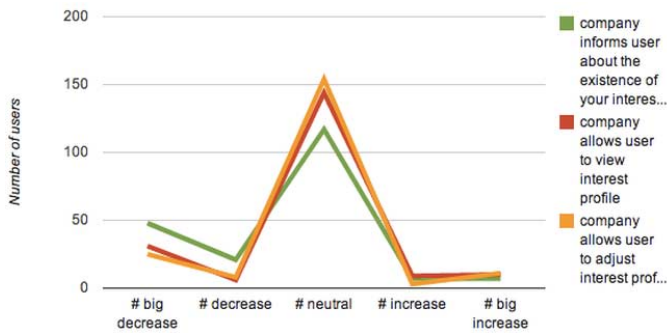
users who don't care. Interestingly, with regard to overall trust impact, ratings for viewing and adjusting interest models were about equal.

Figure 3. Trust change if control provided: "Don't care" users remain neutral (n=200 for each condition).
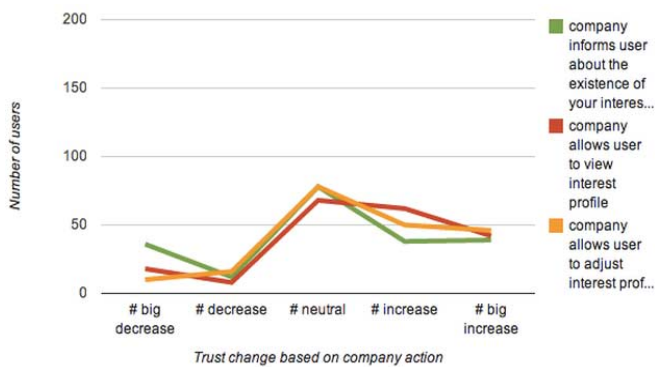
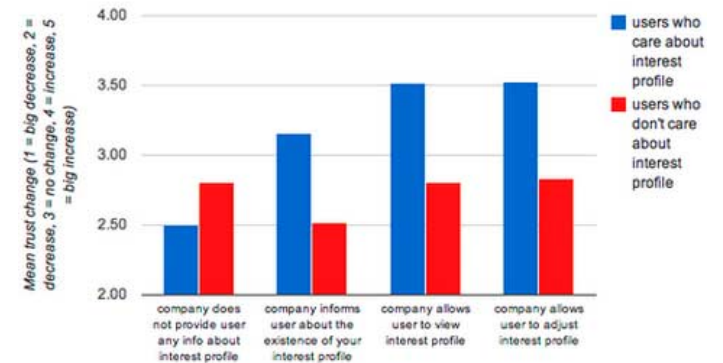In Fig. 4, we see how care users react if control is offered. These users tend to remain neutral or increase their trust levels: from 100 respondents with declined trust in total to 277 respondents with increased trust (not significant). This result is not particularly surprising. However, across both charts in Fig. 3 and 4, responses to various control options of letting know, see and adjust overlap.

Overall trust and levels of control
The third research goal examined overall notions of trust with different levels of control. We present our results along the user segmentation we have constructed earlier: users who don't care and users who care about inferred user interests.

In the following Fig. 5, we have calculated an index to see the overall impact of control on trust. As one can see, overall levels of trust increase for users who care when more control is offered, whereas trust remains on the same level for

Figure 5. Users who care increase overall trust with more control (n=200 for each condition).

Discussion

In this paper we explored the interplay between trust, transparency and control. The major contributions of this study are: 1.) In contrast to previous research, the value of transparency and control differs widely among users. In our segmentation, half of the users responded as if they don't care and the other half stated that transparency and control is important to them, with a tendency to desire full control. 2.) Between these two user groups, we found substantial differences regarding trust in internet companies. For instance, when offered more control, such as seeing and adjusting their profile, users who don't care remained neutral regarding trust. 3.) Providing more control options to users who care about inferred user interests, increased their overall level of trust.

Figure 4. "Care" users increase trust if company does provide control (n=200 for each condition).

In the introduction of this paper we have illustrated several challenges for companies with regard to providing transparency and control to users. The hypothesis, that offering more control could generally scare users and lead to less trust, could not be supported through this particular study. From our research, we tentatively conclude that for one group of users, the complexity reduction mechanism of trust may operate at a higher level of abstraction. In other words, users who do not care about the detailed workings of personalization may simply want assurance that their data is protected, and desire experiences that continuously support this notion of trust. For the other group, the care users, transparency and control is an essential part to establish trust, hence, they want more detailed information about personalization and be able to easily manage their data. Both user groups may need individual experiences reaffirming their conceptualization of trust.

In this paper we studied transparency and control at an abstract level also exploring a new survey approach. The research could be extended substantially, for instance, by exploring transparency and control in a more product embedded experience and with other methods such as experiments.

References

[1] L. Brandimarte, A. Acquisti and G. Loewenstein, "Misplaced Confidences: Privacy and the Control Paradox." Social Psychological and Personality Science. in press.

[2] C. Porter and N. Donthu, "Cultivating Trust and Harvesting Value in Virtual Communities. Management Science 54(1), 113–128, 2008.

[3] J. Staddon et al., "Are privacy concerns a turn-off? Engagement and privacy in social networks". Symposium on Usable Privacy and Security (SOUPS), Washington, DC, USA. 2012.

[4] Y. Wang et al., "I regretted the minute I pressed share": a qualitative study of regrets on Facebook. In Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS '11). ACM, New York, NY, USA,

[5] N. Luhmann, "Familiarity, Confidence, Trust: Problems and Alternatives", in Gambetta (ed.) Trust: Making and Breaking Cooperative Relations, electronic edition, Department of Sociology, University of Oxford, chapter 6, 94-107, 2000.

[6] L.G. Zucker, "Production of Trust: Institutional Sources of Economic Structure, 1820-1920," In. Staw/Cummings (Eds.), Research in Organizational Behavior, Vol 8, 53-111, Greenwich, CT: JAI Press, 1986.

[7] Karahasanović et al., "Ensuring Etiquette, Trust, and Privacy when developing Web 2.0 Applications." Computer 42(6), 42-49, 2009.

[8] X. Luo, "Trust production and privacy concerns on the Internet. A framework based on relationship marketing and social exchange theory." Industrial Marketing Management 31, 111-118, 2012.

[9] J. Camp, H. Nissenbaum, and C. McGrath: "Trust: a Collision of Paradigms" Lecture Notes in Computer Science Volume 2339, 2002, pp 91-105, 2002.

[10] GCS, "Google Consumer Surveys", Google Inc. 2013.

[11] P. McDonald, M. Mohebbi, and B. Slatkin, "Comparing Google Consumer Surveys to Existing Probability and Non-Probability Based Internet Surveys", 2012.

[12] B. Debatin et al., "Facebook and Online Privacy: Attitudes, Behaviors, and Unintended Consequences. Journal of Computer-Mediated Communication 15, 83–108, 2009.

[13] V. Garg, V. and J. Camp, "End User Perception of Online Risk under Uncertainty" System Science (HICSS), 2012 45th Hawaii International Conference.