November 17, 2017

Federal Trade Commission
600 Pennsylvania Avenue, NW
Washington, DC 20580

RE: Paper Submission for Participation in FTC PrivacyCon 2018

Dear PrivacyCon Organizers,

We would like to submit the attached paper for consideration for presentation at PrivacyCon 2018. Our research examines the factors that influence consumers to engage in privacy-protective behaviors in interpersonal and commercial settings online. This study fits the "Quantifying Costs and Benefits of Privacy From a Consumer Perspective" topic of interest in the call for presentations in the following ways:

- We identify several factors that influence consumers' decisions to lie about themselves to protect their privacy online, including individual characteristics and preferences
- We determine how contextual factors play a role in such decisions, and how consumers' perceptions of privacy risks can vary based on whether they are interacting with people or systems (such as commercial websites) online

As per the call for participation, the details for the project are included below. Please feel free to contact us if you would like any additional information. Thank you for considering our application.

Best regards,
Shruti Sannon, Dr. Natalya N. Bazarova, Dr. Daniel Cosley

REQUEST SUMMARY

First and last name, email address, and phone number of researcher(s) making the Request:

Shruti Sannon

Dr. Natalya N. Bazarova

Dr. Daniel Cosley

Title and abstract of the research you propose to present, summarizing your methodology, findings, and how your research differs from prior research in this area:

Title: Understanding People's Decisions to Tell Privacy-Protecting Lies in Multiple Online Contexts

Abstract:

In this paper, we study "privacy lies", which are lies people tell about themselves in order to protect their privacy online. We examine this concept through a mixed-method study, which includes a survey and a qualitative analysis of consumers' engagement in privacy lies in commercial and interpersonal settings online. Going beyond commonly used privacy lenses that focus on privacy concerns or cost/benefit analyses, we explore how contextual factors, motivations, and individual-level characteristics affect privacy-protective lying behavior through a 356-person survey.

Building on previous work that predicts privacy lies based solely on self-reported privacy concerns (R^2=.10), we present a statistical model with significantly greater predictive power (R^2=.35) that includes new predictive variables, including people's perceived control over their data online. A thematic analysis of privacy lies further shows that people weigh the nature of the request and the requestor when deciding on whether to engage in privacy lies, and that their goals and concerns vary by interaction context.

Furthermore, a thematic analysis of responses from people who do not engage in privacy lies reveals multiple reasons to not tell privacy lies; for example, that such lies are unneeded because privacy is not a large concern, or that lying is always immoral. Together, our results point to the need for conceptualizations of privacy lies—and privacy-protective behaviors (PPBs) more broadly—that account for multiple goals, control, context, and attitudes about PPBs.

Methodology:

We collected the data through a survey of Amazon Mechanical Turk (mTurk) workers that invited them to participate in a study on online behaviors and attitudes. Participants were asked whether they had ever lied to protect their privacy online, and if so, whether they had told these lies to people in online interpersonal contexts or directly to websites or other devices where another human was not directly involved. These participants were then asked to recall and describe a detailed example of when they had told a privacy lie to a person and/or a system, as applicable, and their reasons for telling the lie. Participants who reported not having told privacy lies were also asked to explain their reasoning in an open-ended response. All participants completed several privacy-related measures and provided us with demographic information.

We conducted two main analyses to address our research questions. First, we developed a predictive model based on quantitative survey data to see how individual characteristics influence privacy lie behavior. Second, we coded and performed a thematic analysis of the examples of privacy lies to

understand the motivations underlying the decision to engage in this type of PPB as well as the contexts in which it occurs. We then coded and analyzed the reasons people gave for not telling privacy lies.

Findings:

- In this study, we put forth an improved model to predict privacy lies. We found that, controlling for age and gender, frequency of privacy lies was predicted by a) perceived control of data, b) general willingness to disclose online, c) the number of other PPBs people engage in online, and d) attitudes about privacy lies.
- We investigated the contextual factors that may influence privacy lies by coding the reasons our participants provided for telling privacy lies. We identified 11 different factors that influenced people to lie to protect their privacy online. These factors could be broadly categorized into four overarching themes: the first two themes involve people's perceptions of the information request and the requesting party, while the second two themes represent people's motivations to gain or avoid a particular outcome. Our cross-comparison of interpersonal and system contexts also highlights several differences in how people perceive privacy threats and engage in PPBs in these settings: for example, 24.6% of participants who told privacy lies to another person did so because they were interacting with someone they were relatively unfamiliar with; in contrast, unfamiliarity was rarely called out in privacy lies to systems (3.2%); instead, participants mentioned telling privacy lies to websites they used frequently.
- We also sought to understand the reasons for not engaging in privacy lies through a thematic content analysis of responses from 91 participants who reported never telling privacy lies. Participants' reasons to not tell privacy lies could be organized into four main categories: that they thought such lies were unneeded, that they preferred to use another PPB, that lying -- and privacy lies by extension -- were morally wrong, and that privacy lies could have unforeseen negative consequences.

Novelty:

- Prior research on consumers' privacy-protective behaviors has found that privacy concerns can predict participants' likelihood to falsify information, use privacy-enhancing tools, or withhold information. However, privacy concerns do not explain the whole story behind people's decision to tell privacy lies, and predictive models have not been able to explain more than 10% of the variance in this behavior. Our study puts forth a model that can explain 35% of the variance for telling such lies, and we identify four important predictors of this behavior: perceived control of data, general willingness to disclose online, use of other PPBs (such as VPNs and do-not-track plugins), and attitudes about privacy lies.
- Our study is one of the first to provide a detailed thematic analysis of the various types of contexts and motivations that underlie consumers' decisions to lie to protect their privacy in both interpersonal and commercial settings online. As such, we find that a person's decision to tell privacy lies is highly nuanced and goes beyond concern for their personal data. For example, we find that lying can be a way to maintain privacy to explore a stigmatized identity or to avoid a range of negative repercussions. Similarly, we find that people can have strong moral stances

against telling such lies, which can also be used to inform the design of privacy-protective solutions.

- Overall, this study contributes to discussions about privacy research more broadly by highlighting the need to examine a range of contextual factors and motivations that underlie people's decisions to engage in privacy-protective behaviors. Our study suggests that there is much value in measuring factors beyond general privacy concerns and cost-benefit analyses in understanding consumer behavior around privacy.

Publication details for any research that has been previously published or accepted for publication:

This paper is currently under blind review at the following conference:

Shruti Sannon, N. Bazarova, D. Cosley. 2018. How Goals, Attitudes, and Interaction Context Shape Privacy-Protecting Lies Online. *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI2018)*. Under Review.