Measuring the Effectiveness of Spinner-based Randomized-Response Differential Privacy Communication for Sensitive Data Sharing

Seo Young Ko, Sriram Viswanathan, Alan Esquenazi, Swadhin Routray, Jatan Loya, Tianshi Li, and Lorrie Cranor

Carnegie Mellon University

Abstract

While Differential Privacy (DP) is widely adopted to enable privacy-preserving data sharing, its usability by non-expert users is still a challenge. In this work, we measure the usability of a spinner-based interface to guide a user-led Randomized Response Technique (RRT). We conduct an exploratory between-group online survey that uses the spinner-based RRT interface and collects sensitive information in a hypothetical scenario. We found that many participants did not follow instructions on the spinner interface to add noise correctly, and did not understand the purpose of the spinner.

1 Introduction

Differential Privacy (DP) is one of the commonly adopted Privacy-Enhancing Technologies to enable privacy-preserving data sharing. Given the technical complexity of DP, the usability of such a technique by non-expert users is a challenge. Existing research has examined ways to enhance the description and interface to communicate the algorithmic concepts of DP to users [2, 3, 7, 8], while little studied the effectiveness of the process of applying DP empirically.

This work aims to improve the transparency of the process of applying DP and improve user agency in the process by designing interfaces that guide users in adding noise to their data. We focused on the Randomized Response Technique (RRT), a local DP mechanism, in which random noise is added at the individual level before sending the data to the server or administrator. Inspired by prior work that proposed spinner to explain the RRT mechanism [2], we designed a spinner interface to provide users with instructions for adding noise.

We conducted a between-group survey with one control group and two experimental groups to test the effectiveness of user-led RRT communication grounded in a survey-taking scenario. The study uses both simple definitions and the visual spinner tool to explain the mechanism of adding noise when a sensitive question is asked in the hypothetical online survey. Based on the data analysis, this study identifies some limitations of the spinner RRT interface proposed by existing literature [2] in terms of the understandability and the ability to guide users to add noise correctly. Also, we highlight suggestions to improve the RRT interface as future work.

2 Methodology

We recruited 60 Prolific participants to complete our online Qualtrics survey. The survey took about 10 minutes to complete and each participant was compensated $3. All 60 users were given a hypothetical scenario regarding the use of recreational drugs. Then, each user was randomly assigned to one of three groups of 20 participants. The first group is the control group, whose survey does not include DP explanation and mechanism. The second and third experimental groups use DP and receive explanations about it. The second one uses RRT with machine-added noise, so users will only be informed that their answer was anonymized through the use of DP without user interaction. The third group uses RRT with user-added noise, as shown in Figure 1. To implement the user noise addition, we have used a spinner, as has been previously done in studies such as Bullek et al. [2], Karegar et al. [5] and Blair et al. [1]. The code to implement our spinner was modified from publicly available code on Github [6]. The spinner's position is also recorded for use in analysis.

- `Group1`: No DP mechanism
- `Group2`: Textual RRT mechanism with automated noise
- `Group3`: Spinner RRT mechanism with user-led noise

Our study was approved by CMU's IRB. The usage of a fictitious scenario helps mitigate any concerns about collecting real sensitive data from the participants.

From a data analysis perspective, our independent variables are our three condition groups, whereas our dependent variables include the user's comprehension, trust level with the noise addition method, honesty level, and the comfort level regarding the sensitive questions asked of them. These are measured from quantitative data, organized on a 5-point Likert scale. We also collected qualitative data, where we asked

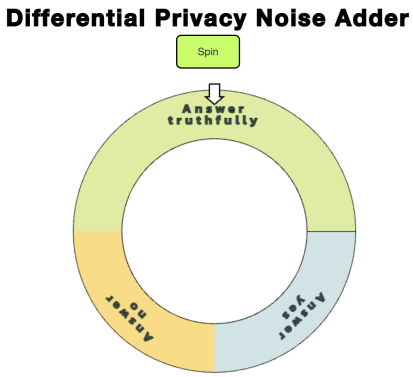

Figure 1: RRT Spinner Interface. To answer the sensitive question, an user first clicks on the "spin" button. If the arrow of the spinner points towards only "answer truthfully" (with 50% probability), the user should answer truthfully. In the other cases, when the arrow points towards "answer yes" or "answer no" (with 25% probability for each), the user should answer "Yes" or "No" for plausible deniability.

users to elaborate on their concerns and opinions in the form of free-text responses. For the analysis, two researchers coded the open-ended questions independently using inductive coding and reached a consensus by resolving conflicts.

3 Results

Overall, our result demonstrates that the RRT spinner interface for DP communication is not as effective as expected to guide users to add noise by themselves.

3.1 User Capability To Add Noise

Based on the recorded spinner responses, we compare their self-reported answers with the spinner responses to verify if they are able to add noise correctly. Since we are using the hypothetical scenarios, we can know the ground truth of the answer and measure whether they follow the guide successfully or not.

We found that only 13 out of 20 in `Group3` participants were able to correctly follow the instructions and add noise. This could be attributed to the ineffectiveness of the interface in communicating the essential information required by the participant to add noise correctly, as per the spinner response.

3.2 Misunderstanding, Discomfort, Dishonesty, and Distrust

Our quantitative data analysis did not lead to any significant differences in comprehension, honesty, discomfort and trust levels of users between all three conditional groups. We also discovered that the qualitative data supports the quantitative findings. For example, the question on understanding the DP mechanism shows similar amounts of confusion, and similar concepts brought up between `Group2` and `Group3`. Some reasons for discomfort for `Group2`, which uses the automated noise mechanism, were the lack of transparency and lack of trust in the mechanism. This is also related to the reasons for distrust that users complains about the lack of transparency:

> "Just being told that it exists doesn't mean it will work"

For the reasons for untrustworthiness, we also saw some lack of understanding of the spinner mechanism in `Group3`:

> "I don't fully understand how a spinner creates noise over my data"

We also found significant bias among users, relating to their previous conceptions surrounding the internet and privacy-preserving mechanisms. Some users believed they were always being tracked; Others thought that the mechanism would not be able to protect them from data leaks. We also found one user mentioning unfamiliarity as a reason for distrust.

4 Limitations, Discussion and Future Work

Our study was limited by a small sample size of participants. We believe a larger sample size might reveal clearer patterns that we were not able to observe in this small dataset. In addition, our study only asks about one kind of sensitive information which is the usage of recreational drugs, and the finding may not apply to other questions with different perceived sensitivity. Lastly, we tried to mimic authentic survey-taking experiences by designing the survey question around a realistic situation, but the use of a hypothetical scenario may still cause confusion and result in a lack of ecological validity. Future work might also benefit from the addition of interviews for more in-depth insights into users' thought processes.

In our pilot tests, we observed that users cannot understand how the spinner interface works until they get a sense of the aggregation and possibility of data analysis with a noise-added database. We believe that instead of focusing on a narrow view of adding noise, we should provide a holistic view of the data life cycle to enhance understandability.

Another option is to guide users to follow an offline coin-flip mechanism in place of a spinner to add noise, as previously suggested in the foundational literature on DP [4]. A few participants expressed concerns over tracking and data leakage, which can be related to the fact that the noise-addition mechanisms are all virtual. We hypothesize that giving participants even more control over the noise-addition mechanism through the use of a coin flip, controlled by the participants themselves, might help ease these concerns.

More research is required in understanding the reasons for the lack of effectiveness of such an interface and ways in which we can improve the effectiveness.

References

[1] Graeme Blair, Kosuke Imai, and Yang-Yang Zhou. Design and analysis of the randomized response technique. *Journal of the American Statistical Association*, 110(511):1304–1319, 2015. doi: 10.1080/01621459.2015.1050028. URL https://doi.org/10.1080/01621459.2015.1050028.

[2] Brooke Bullek, Stephanie Garboski, Darakhshan J. Mir, and Evan M. Peck. Towards Understanding Differential Privacy: When Do People Trust Randomized Response Technique? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3833–3837, Denver Colorado USA, May 2017. ACM. ISBN 9781450346559. doi: 10.1145/3025453.3025698. URL https://dl.acm.org/doi/10.1145/3025453.3025698.

[3] Rachel Cummings, Gabriel Kaptchuk, and Elissa M. Redmiles. "i need a better description": An investigation into user expectations for differential privacy. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, CCS '21, page 3037–3052, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384544. doi: 10.1145/3460120.3485252. URL https://doi.org/10.1145/3460120.3485252.

[4] Cynthia Dwork. Differential privacy. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming*, pages 1–12, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-35908-1.

[5] Farzaneh Karegar, Ala Sarah Alaqra, and Simone Fischer-Hübner. Exploring User-Suitable Metaphors for Differentially Private Data Analyses. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 175–193, Boston, MA, August 2022. USENIX Association. ISBN 978-1-939133-30-4. URL https://www.usenix.org/conference/soups2022/presentation/karegar.

[6] MikeyC0340. dinnerspinner, July 2022. URL https://github.com/MikeyC3040/dinnerSpinner.

[7] Jack Murtagh, Kathryn Taylor, George Kellaris, and Salil Vadhan. Usable Differential Privacy: A Case Study with PSI, September 2018. URL http://arxiv.org/abs/1809.04103. arXiv:1809.04103 [cs].

[8] Felix Wolter and Peter Preisendörfer. Asking sensitive questions: An evaluation of the randomized response technique versus direct questioning using individual validation data. *Sociological Methods & Research*, 42(3): 321–353, 2013.

5 Appendix

5.1 Appendix 1: Survey Questions

Here is our survey, including the flow, as seen in Qualtrics. Notice that users are divided into three condition groups. Also, in the online version, some of the questions (specifically, some of the open-response questions) were randomized in order to avoid biasing effects.

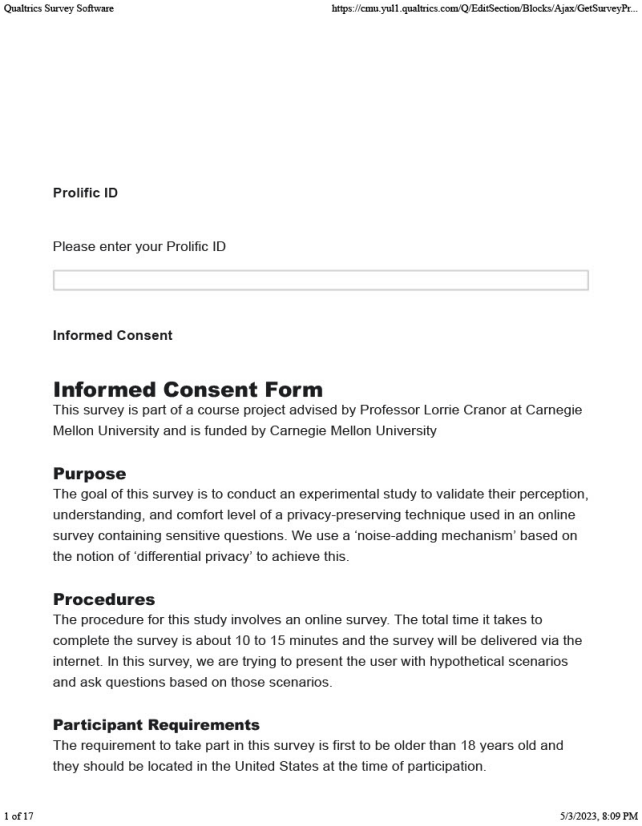

Compensation & Costs

The participants will be paid $3 after completion of the survey, and the compensation will be paid via the Prolific account. There will be no cost to you if you participate in this study.

Confidentiality

The data captured for the research does not include any personally identifiable information about you. Your IP address will not be captured. By participating in this research, you understand and agree that Carnegie Mellon may be required to disclose your consent form, data, and other personally identifiable information as required by law, regulation, subpoena, or court order. Otherwise, your confidentiality will be maintained in the following manner: Your consent form will be stored in a secure location on Carnegie Mellon property and will not be disclosed to third parties. By participating, you understand and agree that the data and information gathered during this study may be used by Carnegie Mellon and published and/or disclosed by Carnegie Mellon to others outside of Carnegie Mellon. However, your name, address, contact information, and other direct personal identifiers will not be mentioned in any such publication or dissemination of the research data and/or results by Carnegie Mellon. Note that per regulation all research data must be kept for a minimum of 3 years.

Future Use of Information

Information collected will be de-identified and archived after the study and will not be used for any future work.

Right to Ask Questions & Contact Information

If you have any questions about this study, you should feel free to ask them by contacting the Principal Investigator now at Alan Esquenazi (Undergrad Student, ECE Department, Carnegie Mellon University, Pittsburgh, PA) (email: aesquena@andrew.cmu.edu, phone: 206-276-6169) . If you have questions later, desire additional information, or wish to withdraw your participation please contact the Principal Investigator by phone or e-mail in accordance with the contact information listed above. If you have questions pertaining to your rights as a research participant;

or to report concerns to this study, you should contact the Office of Research integrity and Compliance at Carnegie Mellon University. Email: irb-review@andrew.cmu.edu . Phone: 412-268-4721.

Voluntary Participation

Your participation in this research is voluntary. You may discontinue participation at any time during the research activity. You may print a copy of this consent form for your records.

Answer yes to the following questions to continue the survey.

I am 18 years or older in age.

Yes
No

I am located in the United Sates at the time of participation.

Yes
No

I have read and understood the information above.

Yes
No

I want to participate in this research and continue with the survey.

Yes
No

Condition Group 3

Imagine you have used recreational drugs in the past year. One day, you saw a paid survey on an online forum for research purposes. You were interested in earning some money and wanted to fill it out. Then you noticed that there were some sensitive questions related to your experience of using recreational drugs.

Consider the following definition of **LDP:**

Local Differential Privacy or LDP protects users' privacy by adding random noise to each response that users give and provides a statistical guarantee of privacy. So, no one can know, just by looking at the overall data, what a specific user's response is, or even if a user was part of the data collection. LDP has been used by companies such as Google and Apple to collect information from users in a privacy-preserving manner.

We are going to use the LDP technique here – you will be adding noise to the answer by using a spinner. To correctly add the noise to your answer, spin the wheel and follow the instructions.

Differential Privacy Noise Adder

Spin

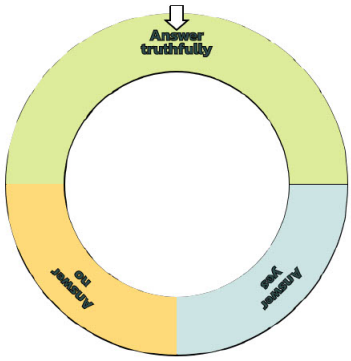

Please answer the question below, imagining what you would do if you had actually used recreational drugs in the past year and were presented with a paid survey online for research purposes that uses the LDP spinner technique described above.

Have you used recreational drugs in the past 1 year?

Spin the spinner and:
If the spinner lands on "Answer Yes", then answer "Yes".
If the spinner lands on "Answer No", then answer "No"
If the spinner lands on "Answer Truthfully", then answer truthfully.

Yes No

To what extent do you agree or disagree with the below statement about the noise-adding mechanism used in the previous question:

The noise addition guarantees the secrecy of my sensitive data (whether I used recreational drugs in the past year or not).

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

To what extent do you agree or disagree with the below statement about the noise-adding mechanism used in the previous question:

The noise addition guarantees that my sensitive data (whether I used recreational drugs in the past year or not) will not be revealed, even if the server's data is disclosed because of an attack.

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

How would you rate your level of understanding of the noise addition mechanism?

Very High
High
Medium
Low
Very Low

To what extent do you agree or disagree with this statement:

The prior description of the privacy protection technique was easy for me to understand.

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

To what extent do you agree or disagree with this statement:
I expect the noise addition technique to provide benefits.

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

What benefit do you expect from adding noise based on the aforementioned explanation?

What was your comfort level when taking this survey and being asked sensitive questions in a hypothetical scenario?

Very High
High
Medium
Low
Very Low

What was the reason for you to feel discomfort?

If this had been a real situation and you were asked to disclose your recreational drug use (or similar sensitive information) within the same online survey interface, how much do you agree that you would be honest in answering this question based on the noise-addition technique?

Strongly agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

What was the reason that you couldn't be honest?

What is your level of trust in the noise-adding mechanism presented before in terms of its privacy preserving capability?

Very High
High
Medium
Low
Very Low

What are the reasons for you to feel less trustworthy?

How confident are you that you correctly added noise to your answers?

Very High
High
Medium
Low
Very Low

Condition Group 2

Imagine you have used recreational drugs in the past year. One day, you saw a paid survey on an online forum for research purposes. You were interested in earning some money and wanted to fill it out. Then you noticed that there were some sensitive questions related to your experience of using recreational drugs.

Consider the following definition of **LDP:**

Local Differential Privacy or LDP protects users' privacy by adding random noise to each response that users give and provides a statistical guarantee of privacy. So, no one can know, just by looking at the overall data, what a specific user's response is, or even if a user was part of the data collection. LDP has been used by companies such as Google and Apple to collect information from users in a privacy-preserving manner.

We are going to use the LDP technique here – random noise will be added automatically once you choose an option and the noise-added answer will be stored in the server.

Please answer the question below, imagining what you would do if you had actually used recreational drugs in the past year and were presented with a paid survey online for research purposes that uses the LDP technique described above.

Have you used recreational drugs in the past 1 year?

Yes No

To what extent do you agree or disagree with the below statements about the noise-adding mechanism used in the above question:

The noise addition guarantees the secrecy of my sensitive data (whether I used recreational drugs in the past year or not).

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

The noise addition guarantees that my sensitive data (whether I used recreational drugs in the past year or not) will not be revealed, even if the server's data is disclosed because of an attack.

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

How would you rate your level of understanding of the noise addition mechanism?

Very High
High
Medium
Low
Very Low

To what extent do you agree or disagree with this statement:

The prior description of the privacy protection technique was easy for me to understand.

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

To what extent do you agree or disagree with this statement:

I expect the noise addition technique to provide benefits.

Strongly Agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

What benefit(s) do you expect/see from adding noise based on the explanation?

What was your comfort level when taking this survey and being asked sensitive questions in a hypothetical scenario?

Very High
High
Medium
Low
Very Low

What was the reason for you to feel discomfort?

If this had been a real situation and you were asked to disclose your recreational drug use (or similar sensitive information) within the same online survey interface, how much do you agree that you would be honest in answering this question based on the noise-addition technique?

Strongly agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

What were the reason(s) that you couldn't be honest?

What is your level of trust in the noise-adding mechanism presented before in terms of its privacy preserving capability?

Very High
High
Medium
Low
Very Low

What are the reason(s) you felt less trustworthy of the mechanism?

Condition Group 1

Imagine you have used recreational drugs in the past year. One day, you saw a paid survey on an online forum for research purposes. You were interested in earning some money and wanted to fill it out. Then you noticed that there were some sensitive questions related to your experience of using recreational drugs.

Please answer the question below, imagining what you would do if you had actually used recreational drugs in the past year and were presented with a paid survey online for research purposes.

Have you used recreational drugs in the past 1 year?

Yes No

What was your comfort level when taking a survey and being asked sensitive questions in a hypothetical scenario?

Very High
High
Medium
Low
Very Low

What were the reason(s) for you to feel discomfort?

If this had been a real situation and you were asked to disclose your recreational drug use (or similar sensitive information) within the same online survey interface, how much do you agree that you would be honest in answering this question?

Strongly agree
Somewhat agree
Neither agree nor disagree
Somewhat disagree
Strongly disagree

What were the reason(s) that you couldn't be honest?

Demographics

What is your age range?

18-22
23-27
28-32
33-37
38-42
43-47
48-52
53 or older

What is your gender?

Male
Female
Prefer not to disclose
Prefer to self-describe

What is your highest education level?

Middle School or lower
Partial completion of high school
High School
Associate's Degree
Bachelor's Degree
Graduate Degree

General Attitudes about Technology

In this section, we'd like to understand your general understanding of technology in your daily life.

How much do you agree or disagree that I can usually figure out new high-tech products and services without help from others?

Strongly disagree
Somewhat disagree
Neither agree nor disagree
Somewhat agree
Strongly agree

General Attitudes about Privacy

In this section, we'd like to understand your general awareness and attitude toward privacy in your daily life.

How much do you agree or disagree with the following statements:

	Strongly Agree	Somewhat agree	Neither agree nor disagree	Somewhat disagree	Strongly disagree
It usually bothers me when online companies ask me for personal information.	○	○	○	○	○
When online companies ask me for personal information, I sometimes think twice before providing it.	○	○	○	○	○
It bothers me to give personal information to so many online companies.	○	○	○	○	○
I'm concerned that online companies are collecting too much personal information about me.	○	○	○	○	○

Hypothetical Scenario: Imagine, you are visiting a website of a discount club. The club offers discounts on consumer products (e.g., electronics, CDs, books) to its members. Generally, an annual membership fee is $50. To obtain free membership, you are required to fill out your personal financial information (e.g., annual income, current debt, annual mortgage payment, checking and saving balances, and any other investments).

Based on this scenario, how much do you agree or disagree with the following statements:

	Strongly Agree	Somewhat agree	Neither agree nor disagree	Somewhat disagree	Strongly disagree
In general, it would be risky to give the information to online companies.	○	○	○	○	○
There would be high potential for loss associated with giving the information to online firms.	○	○	○	○	○
There would be too much uncertainty associated with giving the information to online firms.	○	○	○	○	○
Providing online firms with the information would involve many unexpected problems.	○	○	○	○	○
I would feel safe giving the information to online companies.	○	○	○	○	○

5.2 Appendix 2: Codebook and results

Here is the codebook we used for the open-response questions. Each answer may be coded in more than one category.
Understanding of the benefits of Differential Privacy:
Definitions:

1. Privacy - the participant named privacy as one of the main benefits they expect from the mechanism.

2. Anonymity - the participant named anonymity as one of the main benefits they expect from the mechanism.

3. Anonymity- Obfuscation/Cloaking - the participant named anonymity as one of the main benefits they expect from the mechanism, and specifically mentioned obfuscation or cloaking as the reason for anonymity.

4. Security - the participant named security as one of the main benefits they expect from the mechanism.

5. Increased Honesty - the participant named increased honesty when answering sensitive questions as one of the main benefits they expect from the mechanism.

Count:

Privacy	10
Anonymity	9
Anonymity - Obfuscation/cloaking	3
Security	4
Increased Honesty	2

Reason for discomfort:
Definitions:

1. Disclosure of Sensitive Data - The user named the possible disclosure of sensitive data as a reason for their discomfort.

2. Law enforcement - The participant named the possible involvement of law enforcement with the survey (due to questioning about drug usage) as a reason for their discomfort.

3. Lack of trust - The participant named their lack of trust in the mechanism as a reason for their discomfort.

4. Lack of transparency - The participant named the lack of transparency of the mechanism as a reason for their discomfort.

5. Other - Response not fitting in any of the aforementioned categories.

Count:

Disclosure of Sensitive data	1
Law enforcement	1
Lack of trust	1
Lack of transparency	1
Other	1

Reason for dishonesty:
Definitions:

1. No trust: The participant named their lack of trust in the mechanism as a reason for being dishonest.

2. Fully trust the platform without noise mechanism: The participant thought that the differential privacy mechanism was unnecessary in the face of other mechanisms used by the survey platform to maintain anonymity.

3. Law Enforcement: The participant was afraid of potential disclosure of data to law enforcement

4. Self-esteem: The participant did not want to admit to taking drugs for the sake of their self-esteem.

Count:

No trust	2
Fully trust the platform without noise mechanism	1
Law enforcement	2
Self-esteem	1

Reason for lack of trustworthiness:

1. Lack of understanding: The participants admitted to not fully understanding the workings of the mechanism.

2. Lack of trust in privacy mechanisms in general: The participants do not trust privacy-enhancing mechanisms in general.

3. Unfamiliarity: The participant was unfamiliar with the technique used, and therefore did not trust it.

4. Lack of trust in this specific mechanism: The participant did not trust this mechanism's capacity of providing privacy.

5. Lack of transparency: The participant did not trust this mechanism due to a lack of transparency of its inner workings.

6. Concerned with tracking: The participant was concerned with being tracked while using the internet.

7. Concerned with data leakage: The participant was concerned with their data being leaked.

Count:

Lack of understanding	3
Lack of trust of privacy mechanisms in general	2
Unfamiliarity	1
Lack of trust in this specific mechanism	3
Lack of transparency	1
Concerned with tracking	2
Concerned with data leakage	1