JAMIA Open, 4(4), 2021, 1–5
https://doi.org/10.1093/jamiaopen/ooab099
Research and Applications

Research and Applications

AIM against survey fraud

Daniel Habib [1] and Nishant Jha [2]

[1]Department of Sociology, Johns Hopkins University, Baltimore, Maryland, USA, and [2]Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA

Daniel Habib and Nishant Jha contributed equally to this work.
Corresponding Author: Daniel Habib, Department of Sociology, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA; dhabib1@jhu.edu

Received 27 May 2021; Editorial Decision 22 October 2021; Accepted 30 October 2021

ABSTRACT

Objectives: Although there exists a variety of anonymous survey software, this study aimed to develop an improved system that incentivizes responses and proactively detects fraud attempts while maintaining anonymity.
Materials and Methods: The Anonymous Incentive Method (AIM) was designed to utilize a Secure Hash Algorithm, which deterministically assigned anonymous identifiers to respondents. An anonymous raffle system was established to randomly select participants for a reward. Since the system provided participants with their unique identifiers and passwords upon survey completion, participants were able to return to the survey website, input their passwords, and receive their rewards at a later date. As a case study, the validity of this novel approach was assessed in an ongoing study on vaping in high school friendship networks.
Results: AIM successfully assigned irreversible, deterministic identifiers to survey respondents. Additionally, the particular case study used to assess the efficacy of AIM verified the deterministic aspect of the identifiers.
Discussion: Potential limitations, such as scammers changing the entry used to create the identifier, are acknowledged and given practical mitigation protocols. Although AIM exhibits particular usefulness for network studies, it is compatible with a wide range of applications to help preempt survey fraud and expedite study approval.
Conclusion: The improvements introduced by AIM are 2-fold: (1) duplicate responses can be filtered out while maintaining anonymity and (2) the requirement for the participant to keep their identifier and password for some time before returning to the survey website to claim a reward ensures that rewards only go to actual respondents.

Key words: Anonymous Incentive Method (AIM), fraudulent data, internet research, privacy, Secure Hash Algorithm (SHA)

Lay Summary
Incentivizing a target population to complete a survey can be at odds with maintaining privacy. Anonymous survey software typically either collects some form of personal information (eg, emails) to send rewards to participants or upholds respondents' privacy by providing rewards without requiring personally identifiable information but risks being exploited by scammers. To address this gap, the Anonymous Incentive Method (AIM) was designed to only reimburse actual participants and allow for the detection of scammers without storing personal information of any kind. AIM can be adapted for a wide variety of commercial and research surveys. Despite not completely eradicating survey fraud, AIM can be bolstered by additional safeguards depending on the specific survey. AIM can not only increase the response rate by ensuring anonymity and making potential respondents more comfortable but also help qualify studies as exempt from full review to expedite approval by an Institutional Review Board.

INTRODUCTION

Internet research provides a unique opportunity for cost-effective, large samples from otherwise unreachable populations.[1] Online surveys hold considerable advantages compared to their physical in-person counterparts. Although physical survey drop-boxes allow for anonymous responses, participants may still be hesitant to respond, for example, in fear of being watched or judged. Online surveys replace the physical drop-box with a more neutral intermediary: one's own electronic device. Internet-based surveys distance survey administrators from participants, permitting complete anonymity and freedom from the stigma associated with certain responses such as those related to HIV/AIDS, substance use, and sexual activity.[2–5] However, the shift of research to an online setting raises new concerns for Institutional Review Boards (IRBs) and survey administrators alike.[2]

Fraudsters

The distance and anonymity afforded to survey respondents constitute a double-edged sword as these properties also allow respondents to submit multiple survey responses to increase their compensation.[2] Research has focused on reducing survey scamming that allows the same participant, coined a fraudster,[2] to potentially collect multiple payments and waste research funding.[6] Fraudsters were initially believed to pose a low-stakes threat, being infrequent and easily detected, but early studies on which this misconception was based failed to fully account for the effect of incentives.[6–8] After more assessment studies emerged to elucidate the unexpectedly high rate of participants submitting multiple responses to receive compensation,[9] common preventative measures were questioned for placing the responsibility on potential fraudsters to identify themselves, a mistake Bowen et al.[6] call naïve.

Gap in anonymous survey fraud protection

Despite the substantial literature aimed at mitigating fraud and hence data invalidity,[7,9–14] IRBs and researchers often lack systematic guidelines for assessing protocols that involve an online component.[2] Moreover, most efforts either rely on personally identifiable information, such as Internet Protocol (IP) addresses, names, phone numbers, and emails,[6] which render the survey confidential but not anonymous, or exhibit a level of uncertainty during the fraud detection process that occurs after survey administration (eg, analyzing responses after changing the question order) instead of designing a system that proactively and unambiguously checks for fraudulent data.[2] The development of the Anonymous Incentive Method (AIM) addresses this gap in the literature.

Study objective

The goal of this study was to develop an anonymous survey web application to provide additional protection against fraud while upholding anonymity during reward distribution. A Secure Hash Algorithm (SHA) was implemented to store deterministic, irreversible identifiers to detect multiple responses without personally identifiable information. An incentive system was developed where respondents were required to return to the web app to claim their reward but were still informed about reward distribution without providing emails or phone numbers.

MATERIALS AND METHODS

Web application

A web application was created that allowed respondents to input their names and survey responses. The Google Sheets Application Programming Interface (API) was used to populate a spreadsheet, which served as the backend data store.[15] An additional requirement involved collecting responses while maintaining the anonymity of the respondent. The following method was devised to ensure that the respondent's name was never stored. To this end, the application assigned an anonymous identifier to each name collected by the survey. The assigning process was designed to be both functionally irreversible and stable/deterministic. The irreversibility of the identifier was necessary to preserve anonymity, while the stability of the identifier allowed researchers to filter out repeat attempts.

Encryption function

The identity assignment process used the hashlib library from python's standard library—specifically, the implementation of the SHA-512 algorithm designed by the National Security Agency.[16] This algorithm remains 1 of 7 message digest algorithms approved as a U.S. Federal Information Processing Standard.[17] The web application appended a private secret key to the survey respondent's name and then hashed the resulting string to yield an identifier. This secret key, known as a "salt" in the cryptography field, was used to safeguard against brute force attacks where an attacker generates identifiers en masse and compromises anonymity. In addition, a unique password was generated by hashing the identifier with a second salt. Since the identifier was deterministically formed from the name, and the password was formed from passing the identifier into the same deterministic encryption algorithm, the password was deterministic as well (Table 1). In the interest of open science and to build confidence in our solution, the code for the web app was made publicly available on GitHub.

The name-to-identifier transformation exhibited two properties. First, it was infeasible to reverse. If the original name could be easily determined from the identifier, anonymity would be lost. Second, it was deterministic. Although the identifiers were anonymous, determinism requires each respondent's name to always return the same random identifier. For example, the name "Daniel Habib" was hashed to the string "A8DJS482." This hashed string served as Daniel's identifier, but no one could use the identifier to trace survey responses back to Daniel. Additionally, "Daniel Habib" was transformed to the same identifier (A8DJS482) every time "Daniel Habib" was entered into the survey.

Table 1. Pseudocode and explanation for encryption function implementation

Pseudocode	Explanation
identifier = hashlib.sha512(name + salt1)	The respondent's name is passed into the encryption function along with the salt and returns a deterministic, irreversible identifier.
password = hashlib.sha512(identifier + salt2)	The identifier is passed with a different salt into the same function to return a second identifier that serves as a password.

Anonymous Incentive Method (AIM)

On survey submission, the unique public identifier was shown to the survey respondent along with a private unique password. The respondents were asked to save the identifier and password as well as the data of survey closure. After the conclusion of the survey process, a list of winning identifiers was randomly selected and published on the survey website. Winning users proved their identities using the password to get access to survey rewards. In summary, we used a publicly vetted algorithm created by the US government, opened our source code to examination, and never stored any real names of respondents in the development of a novel incentive system. We are deeply committed to preserving the anonymity of study participants now and in the future.

CASE STUDY

AIM was assessed in an ongoing research study: Peer Influence on Vaping in High Schools. A link to a website survey, which asked for the participant's name and list of friends' names, was sent by school administrators to their 18-year-old high school seniors. After a response was submitted, the names of the respondent, school, and friends were transformed by immediately passing each name into the hashlib.sha512 function. Similar to network analysis of smoking behavior,[18] this study linked each respondent to members of their social network, many of whom were respondents themselves. For example, if survey taker Nishant Jha with hashed identifier "C2BXZ216" typed in "Daniel Habib" as a friend in the high school, Daniel's identifier (A8DJS482) would be returned and used to form a connection between the nodes with identifiers A8DJS482 and C2BXZ216. Upon survey completion, the respondents were provided with each of their identifiers and passwords. Since no emails were collected to maintain anonymity, administrators were asked to email the general student body upon survey completion, instructing anyone who participated to return to the website, check for their identifiers in the list of winners, and enter their passwords, which remained unknown to other respondents, to reveal the code for an electronic Amazon gift card. Thus, no personal identifiers (names, emails, etc.) were ever stored, but AIM provided the possibility to incentivize actual participants.

Implications of connectedness for AIM

This survey was a good fit for AIM because its focus on network effects allowed us to cross-validate the anonymous identifier assignment system. Once the social graph was constructed, the connectedness of a node (indicated by an identifier) became a good proxy for the reliability of a survey response. Fraudulent respondents were not connected to other nodes in the network for three possible reasons. First, no other respondent listed them in their list of friends. This was acceptable and posed no threat to internal validity. Second, respondents made typing errors. This was somewhat precluded by using Regular Expression to check for first and last name, but a respondent could have typed "Daniwl Habib" instead of "Daniel Habib," which would go undetected after running our assignment procedure. Errors due to typos could have been avoided altogether by school administrators providing a list from which high school seniors could select friends, but administrators were not comfortable sharing this information. Third, the encryption algorithm could have somehow incorporated randomness and could have consequently returned a different identifier for the same name but different response. However, since the vast majority of respondents (83%) were linked to other nodes in the network, we concluded that the encryption algorithm proved deterministic, and respondents consistently typed their own names as well as their friends' names.

AIM reliability

Although additional details of the case study are tangential, this study showcased and verified the capabilities of AIM. Since the same name returned the same identifier, respondents could not simply scam the survey to increase the chance of getting a gift card: the identifier system was integrated so that if the same name was entered, the same corresponding initial and reward identifiers were assigned (ie, the system was deterministic). There was also no possibility of another person erroneously receiving an eGift card code from the survey administrator since the password was deterministically formed from the corresponding identifier. However, there existed the possibility of a student disclosing their password or someone else obtaining the password by accessing the respondent's personal property (laptop, phone, paper note on which the password was written, etc.).

DISCUSSION

AIM is a safeguard against survey fraud that can be adapted to the specific needs of the survey. For example, identifiers need not derive from respondents' names. Any piece of information that is unlikely to be the same for different respondents can be used. For example, birthdates can just as easily be passed into the encryption function to return deterministic, irreversible, identifiers. If the survey is large enough or if the target population is expected to share certain attributes (eg, some respondents of a survey for 18-year olds nationwide are likely to share birthdates), multiple pieces of information can be passed into the encryption function altogether to avoid duplicates that arise erroneously due to two different respondents sharing information rather than the same respondent submitting the same responses. Moreover, it may be obvious in this case that responses to other survey questions are nearly identical in the case of the scammer but markedly different in the case of another respondent with the same name, birthdate, etc.

Determinism is not only important for the particular case study where, without consistent identifiers, friends could not be linked to each other to assess network dynamics. In the case of survey scamming, a deterministic, irreversible identifier is more fruitful than a random identifier. For example, it is feasible to design one Qualtrics survey that provides the participant with two random identifiers (one "ID" and one "password") upon survey completion. Then, the survey administrator could potentially inform all respondents through a centralized system without personally identifiable information to enter their password into a second survey. For instance, the administrator could send a general email to the same general population from which participants were recruited but not to participants' emails in particular. Alternatively, the survey website itself could clearly display the end date of the survey and instruct respondents to return to the website on that date to claim their reward without the need for a general announcement. Although rewards would only go to actual respondents, this system is not robust against all scammers since each submission would simply return a different random identifier. Workarounds such as creating a survey link that could be used only once per device, email, etc. would also prove futile because scammers could simply take the survey using a different device, email, etc. Moreover, creating a unique link for each participant would prove challenging while maintaining anonymity.

Potential supplements to AIM

More complicated algorithms can be applied to further decrease mistakes in anonymous survey data. For example, heuristics can be used to calculate the likelihood that a name is mistyped: before transformation to an anonymous identifier, each name can be assigned a distrust score that scales with how far off the entry aligns with its most similar counterpart in an online repository of names. A certain threshold can be established beforehand beyond which names are discarded as fraudulent. For instance, "Daniwl" aligns closely with "Daniel" and is given a low distrust score. In contrast, "Asdfghjkl" is given a high score and is likely attributable to a potential scammer quickly typing in gibberish to maximize the rate of resubmission. As a note of caution, a name repository can reflect certain biases such as being populated predominantly by names from a specific region or culture.

Another way to decrease the resubmission rate involves extracting additional data (ie, IP addresses) from the requests made to the web app. Collecting IP addresses, a piece of personally identifiable information, might not be possible while maintaining anonymity, but IP addresses could be transformed by SHA-512 to a deterministic, irreversible identifier. This would at least limit the number of respondents to the number of their electronic devices. However, this might not prove effective for surveying underprivileged populations since many potential respondents might share devices. Additionally, fraudsters could change their IP address, requiring survey administrators to screen for small variations in the IP address.[6]

Limitations

There are potential limitations for AIM implementation. Changing one's name for each submission results in a new identifier. All of the information that is needed for a scammer to deduce this weakness is available in the consent form as it must be for ethical reasons. Although removing information or needlessly convoluting the language in the consent form is ill-advised and unethical to say the least, survey administrators need not highlight this potential weakness for the respondents. For instance, stating, "Do not scam the survey: we can detect if you submit multiple responses with the same name," would blatantly inform scammers that would have otherwise been detected that they must be more vigilant in their scamming technique.

If the time required for resubmission is too low or if the reward is sufficiently high to justify taking the time to submit a different response, respondents might be more inclined to exploit the incentive system. Additionally, a more educated or conversely underprivileged target population might be more prone to exploiting the incentive system by quickly understanding the basics of hashing or having a greater need for the reward respectively. Nevertheless, additional precautions, such as applying an appropriate incentive and using the Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA) to slow down the turnover rate, could decrease the feasibility of survey scamming. In addition, only a small percentage of respondents in the case study was not linked to other respondents. Together, these considerations suggest that scammers did not want to waste time inputting different data for multiple responses. Even though the data was still anonymous, determinism allowed the researcher to detect if the same anonymous identifier was produced multiple times, unambiguously indicating a scammer. Hence, the case study suggests that AIM was robust to scammers, allowing easy detection and deletion of copies of the same identifier (Table 2).

Although AIM does not completely protect against survey fraud, there is nothing about AIM that would exclude supplementation with additional safeguards. AIM decreases the risk of multiple responders without jeopardizing anonymity to any degree and should therefore serve as a core component of incentivized anonymous surveys. Future research is warranted to develop safeguards that could be implemented in conjunction with AIM.

Table 2. Detection of multiple entries by the same respondent

Single identifier	Multiple entries
538c6beb	41f39156
182f5f1e	41f39156
f2d28fc3	41f39156
f12c90e8	41f39156
277c4b2c	41f39156
815b102f	41f39156
cab5da31	41f39156
b4c1d510	41f39156
3123edd5	41f39156
ee415311	41f39156
47623d60	38cc332b
eb2cd702	38cc332b
cc75ad2f	38cc332b
44e74686	38cc332b
650ab161	49cd3f58
cc062fd0	49cd3f58
	49cd3f58
	49cd3f58
	49cd3f58
	49cd3f58

Implications for IRB exemption

The Code of Federal Regulations requires IRBs to conduct a limited or full review of a study if survey information is obtained in a way that the identity of the research subjects can "readily be ascertained, directly or through identifiers linked to the subjects."[19] Typical incentive systems that require emails would fall under this category. By maintaining anonymity, AIM fulfills this criterion and can thus qualify more studies for IRB exemption, provided that the other criteria for exemption are fulfilled. Especially for smaller studies and student projects, this process for obtaining IRB approval is more efficient for moving projects forward and allows surveys that are highly dependent on changing circumstances to be implemented within the appropriate timeframe. Therefore, AIM helps expedite study initiation by the researchers and alleviate the workload of IRBs, which increases the rate of study approval, increases the time IRBs can spend on full reviews of other proposals, and decreases the financial overhead of IRBs by decreasing the cost per study approval.

CONCLUSION

Python's hashlib function (and equivalent algorithms in other programming languages) can be applied to return deterministic, irreversible identifiers for anonymous surveys, allowing duplicate responses to be filtered out while maintaining anonymity. A potential limitation is a scammer delving deeply into how the hashing algorithm functions to know that one must change their name for the algorithm to return a different identifier, which would preclude the researcher from filtering out multiple responses from the same person. However, additional safeguards (CAPTCHA, appropriate incentives, IP addresses, identifier similarity heuristics, etc.) can be implemented depending on the level of security that the particular

study requires. Despite limitations, the simple addition of a deterministic hash secure algorithm effectively assigned unique identifiers and detected fraudsters, as evidenced in the case study where identifiers were verified by other participants listing them in their network. While ethically required to disclose methodology, it would be wise to not emphasize the fact that scamming with the same name would not work because identifiers are deterministic: the scammer may consequently turn to the longer process of inputting a different name for each submission. The amount of time that passes between submissions does not affect the deterministic conversion of the same name to the same identifier, thus stopping scammers at any point in time and allowing for data from the same participants in follow-up studies to be linked without the researchers accessing personal information. AIM is applicable to a wide range of anonymous surveys, can be modified for a particular study, and can help qualify the study as exempt from full IRB review. Hence, AIM should be incorporated into the recommended standard for anonymous survey protocols and become a reference in institutional guidance that is under development about the detection and mitigation of fraud in conducting online surveys.

HUMAN SUBJECTS PROTECTIONS

The case study was approved by the Johns Hopkins University Homewood Institutional Review Board under an exempt review type.

FUNDING

The case study was funded by the Johns Hopkins University.

AUTHOR CONTRIBUTIONS

DH contributed to the conception of the work; the acquisition, analysis, and interpretation of data; the drafting of the work; and critical revision. NJ contributed to the conception of the work, the drafting of the work, and critical revision.

ACKNOWLEDGMENTS

The authors gratefully acknowledge David Fearon and the Johns Hopkins IRB support staff for helpful feedback as well as Andrew Cherlin for his expertise on survey methodology and the case study in particular.

CONFLICT OF INTEREST STATEMENT

None declared.

DATA AVAILABILITY

The code underlying this article is open source at https://github.com/ninjha01/vape-survey.

REFERENCES

1. Pequegnat W, Rosser BRS, Bowen AM, *et al.* Conducting internet-based HIV/STD prevention survey research: considerations in design and evaluation. *AIDS Behav* 2007; 11 (4): 505–21.
2. Teitcher JEF, Bockting WO, Bauermeister JA, Hoefer CJ, Miner MH, Klitzman RL. Detecting, preventing, and responding to "fraudsters" in internet research: ethics and tradeoffs. *J Law Med Ethics* 2015; 43 (1): 116–33.
3. Bauermeister JA, Zimmerman MA, Johns MM, *et al.* Innovative recruitment using online networks: lessons learned from an online study of alcohol and other drug use utilizing a web-based, Respondent-Driven Sampling (webRDS) strategy. *J Stud Alcohol Drugs* 2012; 73 (5): 834–8.
4. Bauermeister JA, Pingel E, Zimmerman M, *et al.* Data quality in HIV/AIDS web-based surveys: handling invalid and suspicious data. *Field Methods* 2012; 24 (3): 272–91.
5. Catania JA. A comment on advancing the frontiers of sexological methods. *J Sex Res* 1999; 36 (1): 1–2.
6. Bowen AM, Daniel CM, Williams ML, *et al.* Identifying multiple submissions in internet research: preserving data integrity. *AIDS Behav* 2008; 12 (6): 964–73.
7. Birnbaum MH. Human research and data collection via the internet. *Annu Rev Psychol* 2004; 55: 803–32.
8. Reips U-D. Internet-based psychological experimenting: five dos and five don'ts. *Soc Sci Comput Rev* 2002; 20: 241–9.
9. Konstan JA, Simon Rosser BR, Ross MW, *et al.* The story of subject naught: a cautionary but optimistic tale of internet survey research. *J Comput Commun* 2006; 10 (2): JCMC1029.
10. Miner MH, Bockting WO, Romine RS, Raman S. Conducting Internet Research With the Transgender Population: Reaching Broad Samples and Collecting Valid Data. *Soc Sci Comput Rev* 2012; 30 (2): 202–11.
11. Mustanski BS. Getting wired: exploiting the Internet for the collection of valid sexuality data. *J Sex Res* 2001; 38: 292–301.
12. Nosek BA, Banaji MR, Greenwald AG. E-research: ethics, security, design, and control in psychological research on the internet. *J Soc Issues* 2002; 58: 161–76.
13. Reips U-D. Standards for Internet-based experimenting. *Exp Psychol* 2002; 49 (4): 243–56.
14. Riggle EDB, Rostosky SS, Reedy CS. Online surveys for BGLT research: issues and techniques. *J Homosex* 2005; 49 (2): 1.
15. Google. *Google Sheets for Developers: API v4*. 2021. https://developers.google.com/sheets/api/quickstart/python Accessed April 15, 2021.
16. Python Software Foundation. *hashlib—Secure Hashes and Message Digests*. 2021. https://docs.python.org/3/library/hashlib.html# Accessed April 15, 2021.
17. National Institute of Standards and Technology. Secure Hash Standard. *Fed Inf Process Stand Publ* 2015; 4: 180–4.
18. Christakis NA, Fowler JH. The collective dynamics of smoking in a large social network. *N Engl J Med* 2008; 358 (21): 2249–58.
19. Office of the Federal Register, National Archives and Records Administration. 45 CFR § 46.104 - Exempt research. 2019. https://www.govinfo.gov/app/details/CFR-2019-title45-vol1/CFR-2019-title45-vol1-sec46-104