Second-Order Consequences of Differential Privacy

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

R.X. Schwartz

Spring, 2022

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Joshua Earle, Department of Engineering and Society

Introduction

The UN High Commissioner for Human Rights notes that privacy is now ever more threatened by the automated processing of data (2021). At the same time, differential privacy (DP) is an increasingly used data protection technique that supports the privacy of dataset participants. DP consists of adding specific levels of random noise to dataset outputs so that the presence or absence of an individual in the dataset is reasonably unlikely to be discerned (Dwork and Roth, 2014).

A form of DP is used in the publicly-released 2020 Census dataset, as well as in datasets in public health and other fields (US Census Bureau, 2019; Google and Apple, 2021; Ficek et al., 2021; Dyda et al., 2021; Kim et al., 2018). Although DP can increase privacy protections, the technique weakens data analysis, at times making valid data interpretation impossible (Santos-Lozada et al., 2020; Ficek et al., 2021; Swanson et al., n.d.; Swanson and Cossman, n.d.). In other words, DP is an innovation that can affect different stakeholders (social groups) in conflicting ways. This STS research paper highlights four social groups in terms of their relationship to the DP algorithm (namely, DP researchers, DP implementers, dataset analysts, and dataset participants), then examines second-order consequences that arise between these social groups based on their relationship to DP. This paper also considers how some of these consequences may be ameliorated.

Background: Differential Privacy

Statistical research has demonstrated that large, anonymized datasets can unintentionally reveal personal information, either by analyzing a single database on its own (*reconstruction attacks*) or by combining several databases at once (*re-identification attacks*) (Dwork et al., 2017). For example, an adversary may purchase a marketing dataset that contains a participant's

name, address, age, gender, and ethnicity. Another publicly-available dataset may provide *anonymized* health data that is aggregated by county, age, gender, and ethnicity. Participants in the marketing dataset—particularly those who have particularly distinctive combinations of age, gender, and ethnicity in a single county—are now able to have their health data predicted with increased accuracy when these two datasets are combined.

Identification of anonymized dataset participants poses risks both to participant rights and to the interactions that participants have with society. Identification of participants in anonymized datasets may be a violation of participant rights to privacy or confidentiality, and may also violate the trust the participants have placed in a data collector or processor. Identification of participants in anonymized health datasets—or datasets about controversial behavior—may also have consequences for the participants' insurance costs, employability, or social standing (Hansson et al., 2016). Additionally, identification of participants in anonymized datasets may lead to reduced participation in future datasets, causing data quality to decrease or data collection costs to increase. All of the above risks have become more relevant in the face of the increased ability to combine and analyze datasets from different sources, which in turn increases the ability to identify anonymized dataset participants, as mentioned above.

In response to these increasingly relevant risks, data publishers, including governments, businesses, and researchers, have turned to differential privacy, a technique that can provide a mathematical guarantee of privacy protection. In its simplest form, DP requires only a single parameter, epsilon, or ε (Dwork and Roth, 2014). ε controls the amount that DP-enabled (noisy) queries on the original dataset *would change* based on the inclusion of an additional individual in the dataset. A small ε value means that the noisy DP queries would not significantly change based on the inclusion of an additional individual to the dataset. In other words, the noise in the

DP query output nearly overwhelms the effect of adding an individual to the dataset (protecting the individual's privacy relatively well.) A DP query output becomes *less noisy* (and less privacy-protective) when we increase ε: eventually, ε can reach a value where the DP implementation provides negligible privacy protections, because the noise in the ostensibly noisy DP query output is so low that it exposes rather accurately when an individual is added to the dataset. Conversely, if ε is decreased all the way to 0, then a noisy DP query output on the original dataset will be identical to a noisy DP query output from the dataset with an additional individual. However, there is no useful information that is revealed from noisy DP dataset queries when ε=0, as the query outputs consist of pure noise. The challenge faced by implementers of DP is to pick a value of ε that balances the two goals of query utility (higher ε, less noise) and participant privacy protection (lower ε, more noise).

Differential privacy is a valuable privacy-protection technique because it provides a formal guarantee of privacy: the expected economic cost of an individual's participation in the dataset can be expressed mathematically in terms of (1) the ε value and (2) the individual's self-estimated economic loss if their participation is revealed (Hsu et al., 2014). Ideally, this expected economic cost should be low. Another benefit of DP is that the privacy guarantee provided by DP-protected datasets is not affected by the existence or non-existence of other datasets. However, it is important to note that a DP-protected dataset may reveal information about an individual, even if the individual is not a participant in the dataset: it is possible to draw general conclusions from a dataset that are also usually relevant on an individual basis (Dwork and Roth, 2014).

Several variations on differential privacy exist, including (ε, δ)-differential privacy, Rényi differential privacy, and Gaussian differential privacy (Dwork and Roth, 2014; Mironov,

2017; Dong et al., 2020). These versions of differential privacy relax the privacy guarantees of standard differential privacy in order to compensate for real-world cybersecurity, hypothesis testing, or data analysis concerns. The simplest form of DP (discussed above, which only uses ε) can also be implemented in two versions: a local or a global implementation. These are significant tradeoffs between each of these implementations: the local DP implementation has users apply DP before sending data to the central data processor, at the cost of increased noise (above the value targeted by ε) in the final noisy DP-enabled queries (Wood et al., 2020). In contrast, the global DP implementation allows the central data processor to access the raw data and then create the final DP data query output with significantly less noise than the local DP implementation (while maintaining the same ε value.) The local DP implementation should be used when the central data processor should not have a copy of the raw dataset, and the global DP implementation is acceptable when the central data processor can access a raw version of the dataset and subsequently apply the DP algorithm.

Background: Social Construction of Technology and Second-Order Consequences

Wiebe Bijker's social construction of technology, or SCOT (developed with Trevor Pinch and Thomas Hughes), is a sociotechnical theory which was formed and defined in opposition to technological determinism. Technological determinism is the concept that "(1) technology develops autonomously and (2) technology determines societal development to an important degree" (Bijker, 2015, p. 136). Bijker instead argues under the SCOT theory that "technological development should be viewed as a social process… relevant social groups will be the carriers of that process." (Bijker, 1997, p. 42). Further, Pinch and Bijker posit that analyses of technological development should be driven by problems, which are held by social groups (Pinch and Bijker, 1984, p. 414).

The sociotechnical concept of “second-order consequences” describes the social outcomes of a technical innovation outside of the direct technical task at hand (Bauer et al., 1969, p. 14). The purpose of an analysis of second-order consequences is "....to make as many second-order consequences as possible intended, anticipated, and desirable." (Bauer et al., 1969, p. 18). Factors needed to achieve this goal include “(1) an increased understanding of our society… and (2) an ability to detect unanticipated consequences as rapidly as possible." (Bauer et al., 1969, p. 19). This STS paper aims to understand the problems that various social groups have with respect to differential privacy and to leverage this understanding to best shape the second-order consequences of differential privacy.

Understanding Social Groups and their Problems in Differential Privacy

This paper considers DP researchers, DP implementers, dataset analysts, and dataset participants as relevant social groups. These social groups are the carriers of “problems” with respect to DP, as proposed in the SCOT method.

Differential Privacy Researchers

The research lineage of DP can be traced to a 1965 paper which proposed the randomized response survey method (Warner, 1965). This method was referenced almost forty years later in Dwork et al.’s 2006 seminal DP paper. Since 2006, the number of DP papers has grown in both size and scope, with subtopics including DP algorithmic improvements and applied evaluations of DP methods. The “problem” most relevant to DP researchers is the improvement of DP techniques.

Differential Privacy Implementers

Local and global DP are implemented by many large technology companies (including Uber, Apple, and LinkedIn) in internal and public-facing contexts (Tezapsidis, 2017; Apple,

2017; Kenthapadi et al., 2019). Additionally, DP is used in its global form in the 2020 US Census data output (Abowd and Velkoff, 2020). DP has also been proposed and used (in both local and global forms) in health contexts, such as public health and genomics (Dyda et al., 2021; Kim et al., 2018; Google & Apple, 2021). DP libraries such as PipelineDP and OpenDP have been developed, supporting the ability of non-expert users to adopt DP privacy protections (PipelineDP, 2022; OpenDP, 2022).

Problems faced by DP implementers include commitments to *privacy by design*, the pursuit of business advantages by providing more private data to external or internal stakeholders, justifying the ε value chosen in a DP implementation, and the support of the rights of data participants (Tezapsidis, 2017; Apple, 2017; Garfinkel et al., 2020; Kenthapadi et al., 2019). For example, some app usage data (such as typing suggestions or health data) should only be provided to stakeholders if it has significant privacy protections (potentially including DP) (Apple, 2017). Government and corporate implementers are also subject to particular pressure from outside groups, such as the public or regulatory agencies (Abowd and Velkoff, 2020). This pressure influences the question of if or how DP is adopted.

Dataset Analysts

Dataset analysts can be the public, employees, specialists in a particular field, marketers, or insurers, among others. This group encounters the "problem" of being able to use a DP-protected dataset to conduct their desired analysis.

Dataset Participants

Participants in DP datasets are the individuals whose data makes up the dataset. These individuals have the "problem" with DP as to if DP affords them sufficient privacy protections. DP may also raise a "problem" for these users if the perception of DP query randomness causes

these users to be less honest about their responses (John et al., 2018). Additionally, DP may be irrelevant or opaque to participants if they do not have a good understanding of the risk posed by a certain ε value or DP implementation.

Second-Order Consequences of Differential Privacy

Differential privacy is increasingly shifting from a theoretical or specialized technology to a general-purpose technology, leading to a variety of real-world implementations. This section considers the second-order (social) consequences which stem from the use of DP by the above social groups (DP researchers, DP implementers, dataset analysts, and dataset participants). These consequences are supported by the mathematical characteristics of the local and global differential privacy implementations.

Concretization of Conflict

Differential privacy can create conflicts between stakeholder groups where no conflict (or a latent conflict) previously existed. Perhaps the most straightforward example of the concretization of conflict is between *dataset analysts* and *dataset participants*. This concretization occurs because an increase in ε causes an increase in the usability of the dataset at the expense of participant privacy (i.e., the privacy-utility tradeoff). Before differential privacy is implemented on a dataset, there is arguably a latent privacy conflict between dataset analysts and dataset participants that falls to the benefit of dataset analysts; by implementing DP, this conflict is concretized through the epsilon value ε. (However, it is important to note that the privacy-utility tradeoff curve may also be modified by choosing other DP implementations, such as local vs. global, or different privacy preserving methods, such as data enclaves.)

A second case of the concretization of conflict can be seen in the fact that the same ε value is applied to all dataset participants: in the standard implementations of both the local and

global DP methods, a single ε value represents the level of privacy-protection for all dataset participants. Some dataset participants may feel more strongly about privacy than others, yet a dataset-wide ε value concretizes this difference into a conflict mediated by a single value. This ε value may also fail to represent the actual privacy desires of any user in the dataset (i.e., it could be the average of two groups with extreme privacy desires).

A third case of the concretization of conflict occurs with respect to the composition of participants in a dataset. In a non-DP environment, the study of x+n individuals (x and n are positive, nonzero integers) is more desirable than the study of x individuals, *ceteris paribus*, due to both a reduction in sampling error and better claims to comprehensive analysis. Yet in local DP, the inclusion of new individuals in a DP-processed dataset who do not have the same behavior as the main study group can weaken the confidence of the analysis of the main study group. This effect occurs in local DP because as the measured subgroup stays the same size while the sample size scales up by a factor of x, the standard deviation of the DP statistic describing the proportion of the sample in the subgroup is reduced by a factor of sqrt(x), while the proportion itself is reduced by a factor of x (own experimentation). In this way, the proportion becomes lower at a faster rate than the standard deviation of the proportion, leading the proportion measurement to eventually become meaningless due to high amounts of noise. The same effect can be seen in the count measurement of a subgroup in local DP when the sample size scales up by a factor of x and the target subgroup remains the same size. The standard deviation of the count measurement increases by a rate of sqrt(x), forcing the count measurement to become less and less accurate (own experimentation).

This scaling property of local DP creates a new conflict between a desire for dataset comprehensiveness and a desire for accuracy among a measured subgroup; before the

implementation of local DP, the inclusion of an additional participant in a dataset had no influence on the accuracy of an analysis of a different dataset subgroup. Yet with local DP implemented, there now exists a risk of a negative influence on pre-existing subgroup analysis by adding participants of a different subgroup.

An example can be considered among different local DP dataset analysts who are both responsible for a single local DP tool and DP output dataset. One DP analyst may not want the DP tool to be shared among a wider population that has a lower occurrence of a target behavior, because this expansion of the dataset will impair their capacity to accurately analyze the occurrence of the target behavior among current participants. Another example can be seen when adversarial participants want to prevent data from being analyzed among active participants in a different subgroup under local DP analysis. These adversarial participants can join the dataset in high numbers, thereby increasing the sample size in database queries while reducing the statistical accuracy of the other subgroup's DP data query.

Centralization

DP puts pressure on centralizing services and technical capacity due to the complexity and novelty of the technique. The result is that social groups may benefit unequally from the implementation of DP: implementers which have better access to statistical and technical resources will be more apt to use DP, as can be seen by the initial development and implementation of DP by large technology corporations. Over time, DP has become more accessible to users via the creation of several software libraries (OpenDP, 2022; PipelineDP, 2022). However, the statistical and technical knowledge needed to successfully implement a DP protocol, analyze DP data, and communicate the benefits of DP is often outside of the technical expertise of data analysts. In some cases, DP will also increase the cost of studies, due to the

need to increase the number of participants (Oberski and Kreuter, 2020). Further, the priorities held by researchers within the DP field may not accurately reflect real-world needs of DP implementers, leading to the creation of DP research designs that never make it into real-world use or testing (Bambauer et al., 2014).

Obscuring Certain Forms of Analysis

One more second-order consequence of DP implementation is that certain forms of analysis will be precluded. This effect mainly occurs in cases where the metric of interest is below the signal/noise ratio: technologies in this space will be in a "dead zone" wherein they cannot be evaluated (or data collection procedures will fail to generate useful results). Additionally, some metrics (such as those targeted towards uncommon subgroups) will not be able to be measured using DP, due to accuracy or cost limitations (Oberski and Kreuter, 2020).

The concern of a metric falling into a "dead zone" can be particularly troublesome when the measurement of interest is a part of an online (i.e., continuously operating) system. It may be that a measurement of interest is pushed into the "dead zone" and that the only way to get out of the "dead zone" is to obtain a more accurate measurement of interest, thereby creating a vicious cycle. For example, the total sample size of a DP dataset could decrease, worsening the accuracy of a measurement of interest—yet a more accurate measurement may be needed in order to reverse this effect and increase the sample size of the dataset. In cases such as this, the system at hand may be prone to an increased risk of massive failure due to an inability to obtain accurate measurements.

Recommendations and Conclusion

This section explores how different stakeholder groups may work to reduce the negative effects of second-order consequences. One of the most straightforward ways of reducing

unintended and negative second-order consequences that come from the use of DP is to continue to diffuse DP education, training, and discussion into non-technical society and among practitioners in applied fields, such as health and political science. Among practitioners, the development of DP libraries and the increase in practical, field-of-study-specific demonstrations are good avenues for better understanding and implementation of DP. Additionally, the use of simulations to demonstrate risk and data analysis capabilities may be instructive for specialists in different fields to be able to evaluate the effectiveness and desirability of DP implementations. In this way, optimal tradeoffs within DP solutions (or outside of DP-specific solutions) can be more easily found and adopted. Further, the publishing and justification of particular epsilon values is a valuable communication task that increases the legitimacy of DP (Dwork et al., 2019; Garfinkel et al., 2020).

Another way to reduce unintended and negative second-order consequences is to emphasize that any particular DP implementation should be considered as one of many potential solutions in a risk framework. Potential solutions can extend to DP characteristics (such as the distribution of the privacy budget, the level of epsilon, or the capability of a (ε, δ) DP implementation to degrade gracefully) or to the use of DP in combination with (or in place of) other privacy-protecting solutions, such as data enclaves. The 2020 US Census DP implementation provides a good case study of the ways in which a two-way conversation about DP decisions can result in a more successful DP implementation in terms of stakeholder satisfaction (Abowd and Velkoff, 2020). There also may be possible tradeoffs (either in post-processing, data collection, or DP methods) that can make the dataset more suitable for research purposes while still presenting a suitable privacy guarantee for users against data use for nefarious purposes (Abowd and Velkoff, 2020). However, some practitioners contend that the

privacy-utility tradeoff provided by DP has no point which satisfies the mutual goals of privacy and utility, thereby supporting the selection of an alternative privacy protection method (Bambauer et al., 2014).

This paper indicates that rather than being a purely privacy-oriented technology, differential privacy presents a variety of second-order consequences. By anticipating these second-order consequences, DP can be better managed and shaped to serve the needs of different social groups.

References

Abowd, J., & Velkoff, V. (2020, June 18). *Modernizing Disclosure Avoidance: A Multipass Solution to Postprocessing Error*. The United States Census Bureau. Retrieved February 21, 2022, from https://www.census.gov/newsroom/blogs/research-matters/2020/06/modernizing_disclosu.html

Apple. (2017, November). *Differential Privacy*. https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

Bambauer, J., Muralidhar, K., & Sarathy, R. (2014). Fool's Gold: An Illustrated Critique of Differential Privacy. *Vanderbilt Journal of Entertainment & Technology*, *16*(4), 701–755. https://scholarship.law.vanderbilt.edu/jetlaw/vol16/iss4/1/

Bauer, R. S., Rosenbloom, R. S., & Sharpe, L. (1969). *Second-Order Consequences: A Methodological Essay on the Impact of Technology*. The MIT Press.

Bijker, W. E. (1997). *Of Bicycles, Bakelites, and Bulbs*. Amsterdam University Press.

Bijker, W. E. (2015). Technology, Social Construction of. *International Encyclopedia of the Social & Behavioral Sciences*, 135–140. https://doi.org/10.1016/b978-0-08-097086-8.85038-2

Dong, J., Roth, A., & Su, W. (2020). *Gaussian Differential Privacy*. Royal Statistical Society. https://rss.org.uk/RSS/media/Training-and-events/Events/2020/Dong-et-al-jrssb-final.pdf

Dwork C., McSherry F., Nissim K., Smith A. (2006). *Calibrating Noise to Sensitivity in Private Data Analysis.* In: Halevi S., Rabin T. (eds) Theory of Cryptography. TCC 2006. Lecture Notes in Computer Science, vol 3876. Springer, Berlin, Heidelberg. https://doi.org/10.1007/11681878_14

Dwork, C., & Roth, A. (2014). *The Algorithmic Foundations of Differential Privacy*. Macmillan Publishers. https://doi.org/10.1561/0400000042

Dwork, C., Smith, A., Steinke, T., & Ullman, J. (2017). Exposed! A Survey of Attacks on Private Data. *Annual Review of Statistics and Its Application*, *4*(1), 61–84. https://doi.org/10.1146/annurev-statistics-060116-054123

Dwork, C., Kohli, N., & Mulligan, D. (2019). Differential Privacy in Practice: Expose your Epsilons! *Journal of Privacy and Confidentiality*, *9*(2). https://doi.org/10.29012/jpc.689

Dyda, A., Purcell, M., Curtis, S., Field, E., Pillai, P., Ricardo, K., Weng, H., Moore, J. C., Hewett, M., Williams, G., & Lau, C. L. (2021). Differential privacy for public health data: An innovative tool to optimize information sharing while protecting data confidentiality. *Patterns*, *2*(12), 100366. https://doi.org/10.1016/j.patter.2021.100366

Ficek, J., Wang, W., Chen, H., Dagne, G., & Daley, E. (2021). Differential privacy in health research: A scoping review. *Journal of the American Medical Informatics Association*, 28(10), 2269–2276. https://doi.org/10.1093/jamia/ocab135

Garfinkel, S., Rodriguez, R., & Leclerc, P. (2020, January). *Differential Privacy at the US Census Bureau: Status Report*. U.S. Census Bureau. https://csrc.nist.gov/CSRC/media/Projects/pec/documents/stppa-01-20200127-talk03-Garfinkel-diff-priv-census.pdf

Google & Apple. (2021, April). *Exposure Notification Privacy-preserving Analytics (ENPA) White Paper*. Apple. https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf

Hansson, M. G., Lochmüller, H., Riess, O., Schaefer, F., Orth, M., Rubinstein, Y., Molster, C., Dawkins, H., Taruscio, D., Posada, M., & Woods, S. (2016). The risk of re-identification versus the need to identify individuals in rare disease research. *European Journal of Human Genetics*, *24*(11), 1553–1558. https://doi.org/10.1038/ejhg.2016.52

Hsu, J., Gaboardi, M., Haeberlen, A., Khanna, S., Narayan, A., Pierce, B. C., & Roth, A. (2014). Differential Privacy: An Economic Method for Choosing Epsilon. *2014 IEEE 27th Computer Security Foundations Symposium*. https://doi.org/10.1109/csf.2014.35

John, L. K., Loewenstein, G., Acquisti, A., & Vosgerau, J. (2018). When and why randomized response techniques (fail to) elicit the truth. *Organizational Behavior and Human Decision Processes*, *148*, 101–123. https://doi.org/10.1016/j.obhdp.2018.07.004

Kenthapadi, K., Tran, T., Dietz, M., & Koeppe, I. (2019, April 10). *Privacy-preserving analytics and reporting at LinkedIn*. LinkedIn Engineering. Retrieved February 21, 2022, from https://engineering.linkedin.com/blog/2019/04/privacy-preserving-analytics-and-reporting-at-linkedin

Kim, J. W., Jang, B., & Yoo, H. (2018). Privacy-preserving aggregation of personal health data streams. *PLOS ONE*, *13*(11), e0207639. https://doi.org/10.1371/journal.pone.0207639

Mironov, I. (2017). Renyi Differential Privacy. *arXiv*. https://arxiv.org/abs/1702.07476

Oberski, D. L., & Kreuter, F. (2020). Differential Privacy and Social Science: An Urgent Puzzle. *2.1*, *2*(1). https://doi.org/10.1162/99608f92.63a22079

OpenDP. (2022). OpenDP. Retrieved May 12, 2022, from https://opendp.org/

Pinch, T. J., & Bijker, W. E. (1984). The Social Construction of Facts and Artefacts: or How the Sociology of Science and the Sociology of Technology might Benefit Each Other. *Social Studies of Science*, *14*(3), 399–441. https://doi.org/10.1177/030631284014003004

PipelineDP. (2022). PipelineDP. Retrieved May 12, 2022, from https://pipelinedp.io/

Santos-Lozada, A. R., Howard, J. T., & Verdery, A. M. (2020). How differential privacy will affect our understanding of health disparities in the United States. *Proceedings of the National Academy of Sciences*, 117(24), 13405–13412. https://doi.org/10.1073/pnas.2003714117

Swanson, D. A., Bryan, T. M., & Sewell, R. (n.d.). *The Effect of the Differential Privacy Disclosure Avoidance System Proposed by the Census Bureau on 2020 Census Products: Four Case Studies of Census Blocks in Alaska*. National Conference of State Legislatures. Retrieved October 3, 2021, from https://www.ncsl.org/Portals/1/Documents/Elections/Four_Case_Studies_of_Census_Blocks_in_Alaska.docx.pdf

Swanson, D. A., & Cossman, R. E. (n.d.). *The Effect of the Differential Privacy Disclosure Avoidance System Proposed by the Census Bureau on 2020 Census Products: Four Case Studies of Census Blocks in Mississippi*. National Conference of State Legislatures. Retrieved October 3, 2021, from https://www.ncsl.org/Portals/1/Documents/Elections/Four_Case_Studies_of_Census_Blocks_in_Mississippi.pdf

Tezapsidis, K. (2017, July 13). *Uber Releases Open Source Project for Differential Privacy*. Uber Privacy and Security. Retrieved February 21, 2022, from https://medium.com/uber-security-privacy/differential-privacy-open-source-7892c82c42b6

United Nations High Commissioner for Human Rights. (2021, September). *The right to privacy in the digital age: report (2021)* (A/HRC/48/31). Human Rights Council - United Nations High Commissioner for Human Rights.

https://www.ohchr.org/EN/HRBodies/HRC/RegularSessions/Session48/Documents/A_HRC_48_31_AdvanceEditedVersion.docx

US Census Bureau. (2019, July 1). *Census Bureau Adopts Cutting Edge Privacy Protections for 2020 Census*. The United States Census Bureau. https://www.census.gov/newsroom/blogs/random-samplings/2019/02/census_bureau_adopts.html

Warner, S. L. (1965). Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias. *Journal of the American Statistical Association*, *60*(309), 63–69. https://doi.org/10.1080/01621459.1965.10480775

Wood, Alexandra, Micah Altman, Kobbi Nissim, and Salil Vadhan. (2020). "Designing Access with Differential Privacy." In: Cole, Dhaliwal, Sautmann, and Vilhuber (eds), *Handbook on Using Administrative Data for Research and Evidence-based Policy*. Retrieved May 22, 2022, from https://admindatahandbook.mit.edu/book/latest/diffpriv.html