NBER WORKING PAPER SERIES

PRIVACY PROTECTION AND ACCURACY: WHAT DO WE KNOW? DO WE KNOW THINGS?? LET'S FIND OUT!

Evan S. Totty
Thor Watson

Working Paper 32989
http://www.nber.org/papers/w32989

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

We have no funding and/or financial relationships to disclose. The views expressed herein are those of the authors and do not necessarily reflect the views of the U.S Census Bureau or the National Bureau of Economic Research.

Privacy Protection and Accuracy: What Do We Know? Do We Know Things?? Let's Find Out!
Evan S. Totty and Thor Watson
NBER Working Paper No. 32989
September 2024
JEL No. C42, C81, C83, J10

ABSTRACT

Statistical agencies have a dual mandate to provide accurate data and protect the privacy and confidentiality of data subjects. These mandates are fundamentally at odds and therefore must be balanced: more accurate data reduces privacy, while privacy protections introduce error that reduces accuracy. Balancing accuracy and privacy requires, among other things, that we can quantify accuracy and privacy. Quantifying privacy has become easier thanks to differential privacy. Quantifying accuracy may sound easy by comparison, but there are many challenges to doing this effectively. In this paper, we first discuss some challenges associated with quantifying data accuracy. We then focus on an often-ignored challenge, which is the existence of survey error in the data being protected. We provide an overview of how privacy protection error relates to total survey error. We also summarize recent work that uses validation data to quantify the impact of privacy protection error relative to and conditional on other sources of survey error. Finally, we discuss opportunities and challenges for future work on data privacy and survey error.

Evan S. Totty
US Census Bureau
4600 Silver Hill Road
Washington, DC 20233
evan.scott.totty@gmail.com

Thor Watson
US Census Bureau
4600 Silver Hill Road
Washington, DC 20233
thorwatson.econ@gmail.com

1 Introduction

Statistical agencies collect and disseminate survey data on virtually every aspect of the economy and society of the United States. The purpose of collecting the survey information is to provide accurate data to the public. At the same time, these agencies are bound by laws requiring them to protect respondent privacy. Therefore, every survey undergoes some type of privacy protection between data collection and data dissemination to prevent the public from re-identifying respondents in the data, a process known as statistical disclosure limitation (SDL). SDL introduces intentional error into the data for the sake of privacy.[1] There is an inherent trade-off between protecting respondent privacy and releasing accurate data to the public. Releasing more data and/or more accurate data reduces privacy (Dinur and Nissim, 2003). Managing this trade-off is a choice that attempts to balance demand for data accuracy and availability against demand for privacy (Abowd and Schmutte, 2019).

A principled evaluation of the privacy-accuracy trade-off is required in order to make choices regarding the appropriate balance of privacy and accuracy. Ideally, this includes quantifying privacy risk and data accuracy so that the trade-off can be clearly evaluated and communicated. In reality, this is a challenging task. In the past, assessments of privacy risk were limited in scope and relied on assumptions regarding an attacker's capabilities (Reiter, 2019). Privacy protection methods were ad hoc and provided no formal guarantees for preventing privacy loss (Abowd et al., 2020, 2023). Furthermore, details regarding even *what* was done to the data for the sake of privacy were often kept secret to reduce the likelihood that the protections could be undone.[2,3] More recently, "differential privacy" solved many of these challenges. Protection methods that satisfy differential privacy quantify bounded worst-case privacy loss risk without assuming an attacker's capabilities and provide a for-

[1] Examples include top- and bottom-coding, coarsening, rounding, suppression, sub-sampling, swapping, synthesis, and noise injection.

[2] Protection methods typically impacted only a small fraction of observations. Because of this, details were often kept secret to add uncertainty to an attacker's confidence.

[3] Famous examples of undoing privacy protections include the identification of Massachusetts governor William Weld's medical records in a dataset released by the state and the identification of Netflix users in a dataset released by the company as part of the Netflix Prize competition (Heffetz and Ligett, 2014).

mal guarantee for the bounds (Dwork et al., 2006; Dwork, 2006; Dwork and Roth, 2014). This bounded privacy loss risk is represented by the parameter known as "epsilon," which measures how much a given database output can change based on the inclusion of a single record in the database. Differential privacy is typically achieved by adding random noise to queries on a database. Because of the quantified and guaranteed bounds on privacy loss risk, agencies can be transparent about the protection method, implementation details, and resulting privacy risk.[4]

Evaluating the impact of privacy protections on data accuracy may sound easy by comparison, but there are several challenges. Many of these challenges arise due to the fact that there is no single metric for accuracy like what differential privacy provides for privacy loss via the epsilon parameter (Brenner and Nissim, 2014). Data collected by federal agencies cover many topics and are used for numerous purposes. Therefore, an accuracy assessment requires the data provider to select a set of priority use cases and a set of particular metrics for evaluation.[5] A given application of privacy protection may have different impacts on different use cases and/or metrics, which means that accuracy results depend on these selections. Another challenge is that increased transparency regarding the application of privacy protection may seem to imply the choice of more privacy and less accuracy (e.g., Ruggles et al., 2019), which is not necessarily correct.[6] These challenges are understood and documented (Abowd and Hawes, 2023; Drechsler, 2019), although they continue to be a challenge for balancing privacy versus accuracy and obtaining buy-in and trust from stakeholders.

However, there is yet another challenge to quantifying accuracy: while accuracy metrics usually involve comparing the data before and after privacy protection, survey data already

[4]See, e.g., the Census Bureau's release of the production code base for the disclosure avoidance system (DAS) used for the 2020 Decennial Census redistricting data (U.S. Census Bureau, 2021*b*) and the Census Bureau's choice of epsilon for the production run of the DAS (U.S. Census Bureau, 2021*a*).

[5]Accuracy assessments sometimes evaluate specific utility for given use cases (e.g., changes to a statistic between the original and protected data). Other times, assessments evaluate general utility based on multidimensional distributional similarity (e.g., classifying whether a given record belongs to the original data or protected data). See Snoke et al. (2018) for additional discussion.

[6]It is important to distinguish between characteristics of a disclosure avoidance system (i.e., *how* data will be protected) and choices regarding its implementation (e.g., parameter choices and other details that will impact the accuracy and privacy of the resulting data). See Hawes et al. (2024) for additional discussion.

contain error before applying privacy protection. This fact is fundamental to survey data and infects all accuracy metrics that are based on simply comparing the original data to the protected data.[7] Such metrics are inherently flawed in the sense that they implicitly assume the original data are without error and thus attribute any differences to reductions in accuracy.[8] When survey data already contain error, differences due to the privacy protection no longer directly translate to reductions in accuracy (nor do they directly translate to reductions in privacy loss risk). Ignoring the reality of survey error before privacy protection contradicts a large literature documenting the existence of survey error and its impact on bias and uncertainty for official statistics and econometric analyses.[9] This challenge is pervasive and difficult to overcome, but there are existing tools that may provide insights.

In this chapter, we first review the taxonomy of survey error and discuss a framework for connecting privacy protection and total survey error. Next, we discuss the results from Totty and Watson (2024), which quantify the contribution of privacy protection to total survey error in a synthetic version of the American Community Survey (ACS). Synthetic data can be differentially private, but even when it is not it still allows for transparency and correct inferences, unlike some other SDL methods (Abowd and Schmutte, 2015). Finally, we wrap up with a discussion of future challenges and opportunities for the evaluation of data accuracy.

2 The Sources of Survey Error

Broadly speaking, there are three traditional sources of survey error: measurement error, item non-response error, and generalized coverage error. Measurement error arises when there is a difference between an individual's survey response and the truth. For example,

[7] Rising survey error has led to conclusions that household surveys are in "crisis" (Meyer, Mok and Sullivan, 2015) and interest in transforming statistical agencies to rely less on survey data (Jarmin, 2019).

[8] Alternatively, such metrics would be ideal if one assumed the goal of a data user was to infer about the original data rather than the population from which the original data were generated, but this contradicts survey methodology and the fact that agencies produce margins of error alongside published statistics.

[9] See, e.g., Bound, Brown and Mathiowetz (2001); Celhay, Meyer and Mittag (2024); Jarmin (2019); Manski (2015); Meyer, Mittag and George (2022); Meyer, Mok and Sullivan (2015); Schennach (2016).

an individual may inflate their income when responding to a survey to save face or they might provide a rough estimate off the top of their head. Measurement error can arise for a multitude of reasons such as the wording of a survey question, miscommunication with the interviewer, or during the data cleaning process. Meanwhile, item non-response error is created when an individual refuses to respond to a specific survey question. For example, an individual may fill out basic demographic questions on a survey but leave an income question blank because they consider it sensitive information. Lastly, generalized coverage error encapsulates the remaining sources of survey error such as sampling error, frame error, unit non-response error, and any survey weight adjustments.[10] These forms of errors arise when the sample is not representative of the target population for various reasons. For example, a surveyor interested in estimating average income by zip code may draw a random sample of households to contact from a public directory of homeowners. As a result, the sample would exclude everyone who is not a homeowner and likely draw from the upper half of the population income distribution (frame error). Even a perfectly designed sampling process can produce a non-representative sample as some individuals may refuse to participate in the survey (unit non-response error), a subgroup may be under-represented in the random sample by pure chance (sampling error), or the application of survey weights to fix other issues can induce its own form of error.

One way to measure survey error and its components is to utilize the Total Survey Error (TSE) framework proposed in Groves and Lyberg (2010) alongside population-level data from administrative records, e.g., tax records, as a source of the truth. By linking survey data to administrative records, it is possible to quantify survey error in key statistics and provide a measure of accuracy. Consider, for instance, that means of survey variables are often used as key indicators of the economy (e.g., unemployment rate, average salary, eviction rate). Using the TSE framework, it can be shown that that total survey error for a mean

[10] Although sampling error and coverage error are usually considered separate sources of survey error, the Total Survey Error framework often combines the two into generalized coverage error for the sake of simplicity. See Meyer and Mittag (2021) for more details.

statistic can be reduced to the following simple formula:

$$\hat{\varepsilon}_{TSE} = \hat{\varepsilon}_{ME} + \hat{\varepsilon}_{INRE} + \hat{\varepsilon}_{GCE} \tag{1}$$

The first term on the right-hand side is measurement error ($\hat{\varepsilon}_{ME}$), which is calculated as the weighted average difference between the survey response for individuals who responded to the relevant survey question (e.g., "What was your total wage and salary income last year?") and their administrative data value (e.g., wage and salary income reported to the IRS). The second term is item non-response error ($\hat{\varepsilon}_{INRE}$), which is calculated as the weighted average difference between the imputed survey response for individuals who did not respond to the survey question and their administrative data value. The third term is generalized coverage error ($\hat{\varepsilon}_{GCE}$), which is calculated by taking the difference between the population average from the administrative data and the weighted survey average when the survey responses are replaced with administrative data values. Note that each component is calculated based on the original survey values and can be positive or negative in value.

Considering the fact that applying privacy protection to a survey introduces some error, a natural extension is to model that error as another component within the TSE framework. Doing so provides a way to evaluate the impact of privacy protection relative to the three traditional sources of error and gain a better understanding of its impact on accuracy. As demonstrated in Totty and Watson (2024), the inclusion of privacy protection error within the TSE framework for a mean statistic can be formulated as the following:

$$\hat{\varepsilon}_{TSE} = \hat{\varepsilon}_{ME} + \hat{\varepsilon}_{INRE} + \hat{\varepsilon}_{GCE} + \hat{\varepsilon}_{SDLE} \tag{2}$$

The new fourth term on the right-hand side is statistical disclosure limitation error ($\hat{\varepsilon}_{SDLE}$), which is calculated as the weighted average difference between an individual's response in the protected survey data and the original survey data.[11] Note that the additive

[11]See Totty and Watson (2024) for additional derivations. Other papers also discuss the need to connect SDL and TSE (Eltinge, 2022; Gong, 2022; Hotz et al., 2022; Karr, 2017), but Totty and Watson (2024) is the first to extend the statistical framework from Meyer and Mittag (2021) to include SDL.

nature of the individual components implies that applying privacy protection to a survey may paradoxically reduce the total survey error for a particular statistic if $\hat{\varepsilon}_{SDLE}$ is opposite in sign of the sum of the three other components ($\hat{\varepsilon}_{ME} + \hat{\varepsilon}_{INRE} + \hat{\varepsilon}_{GCE}$).

3 The Relative Impact of Privacy Protection

After extending the Total Survey Error (TSE) framework to include error from privacy protection, an application is demonstrated in Totty and Watson (2024) using American Community Survey (ACS) data, with and without privacy protection, linked to several administrative or proprietary datasets.[12] The administrative and proprietary datasets provide population-level proxies for the true information that the ACS variables are trying to measure. Using these linked datasets, the paper analyzes the effect of generalized coverage error, non-response error, measurement error, and SDL error on simple but important statistics such as variable means (of wage and salary income, retirement income, home value, property taxes, and birth year) and population sizes of demographic groups (based on race categories, Hispanic status, and citizenship status). The paper also details how the various sources of error differentially impact demographic sub-groups and thereby influence important estimates of inequality.

Comparing the relative size and impact of different types of error is useful for communicating aspects of data quality and uncertainty to data users. For instance, privacy protection is often seen as a detriment to survey quality since it intentionally introduces some error; meanwhile, the original survey data are often viewed through rose-tinted glasses, as if the unaltered survey responses are a close approximation to the truth. However, instances that contradict this narrative are found in Totty and Watson (2024). Consider the wage and salary results which find that measurement error is $4,237; that is, people who respond to

[12]Privacy protection for the data used in the paper is applied in the form of partial synthesis. More specifically, ACS variables of interest are synthesized using classification and regression tree (CART) synthesis methods. For more details on the synthesis model and synthetic data in general, please refer to Abowd and Schmutte (2015); Benedetto, Stanley and Totty (2018); Totty and Watson (2024).

the survey over-report their income on average by $4,237 relative to the amount reported to the IRS. Meanwhile, the amount of error due to privacy protection (SDL error) is $207.10; that is, individuals in the survey with privacy protection have $207.10 more income on average than the version of the survey without privacy protection. In other words, SDL error is only about 5% the size of measurement error. Moreover, SDL error increases the estimated average income in the survey by a mere 0.72% relative to the survey target, i.e., the average income in the administrative IRS data, while the three other sources of error combined increase it by 15.52%. In plain English, the original survey without privacy protection over-estimates average income by 15.52% and applying privacy protection to the survey only increases that bias by an additional 0.72 percentage points.

An even more interesting finding is brought to light when the wage and salary results are stratified by gender in Totty and Watson (2024). More specifically, measurement error overstates the gender wage gap in the survey by $2,091; that is, women do earn less than men on average but measurement error in the ACS is over-estimating that wage gap by $2,091. Meanwhile, SDL error *reduces* the estimated gender wage gap in the survey by $149.80. In a vacuum, one could argue an application of privacy protection that shrinks the estimated gender wage gap is not ideal; however, the net effect of privacy protection must be considered in a broader context that includes the other sources of survey error. In this particular case, for instance, the SDL error is offsetting some of the measurement error, and thereby bringing the estimated gender wage gap closer to the target value in the administrative IRS data. This finding illustrates the uncertain impact of privacy protection on data quality and the difficulty in quantifying its overall impact on accuracy. Although the net effect of privacy protection reduced the accuracy of the average income estimate in the aggregate, it also increased the accuracy of the estimated gender wage gap.

Lastly, for the sake of brevity, two metrics to convey some high-level takeaways are provided in Totty and Watson (2024). First, SDL error on average induces a bias of -0.72% in the survey estimates relative to the target values in the administrative data. This metric

is referred to as the average percentage error (APE), which allows for standardizing the error amounts across variables to then aggregate into a single value. In a similar manner, the average *absolute* percentage error (AAPE) of SDL error is 3.17%. The AAPE provides a useful complement to APE when a source of error is large (in absolute value) but the direction of the error differs across variables such that the APE is deceptively small. When comparing the four sources of error at an aggregate level, SDL error (-0.72% APE, 3.17% AAPE) is smaller on average than coverage error (-1.00% APE, 12.14% AAPE) and measurement error (5.14% APE, 7.64% AAPE). It is also smaller than non-response error (1.31% APE, 1.67% AAPE) in terms of APE, but non-response error is smaller in terms of AAPE. Moreover, even when comparing the sources of error on a variable-by-variable basis, SDL error is generally one of the smallest error components across all variables.

4 Discussion

The framework and results described in this chapter are only a starting point. The results only speak to the accuracy of a single survey, a limited set of variables, and a particular application of privacy protection. Future work should assess to what extent the results generalize by applying the framework to more surveys, variables, administrative data sources, synthesis configurations, and privacy protection methods. This requires access to a variety of administrative and commercial data sources for validation, which will be a challenge in some cases, but data sharing agreements between agencies and the availability of commercial data should make it feasible to expand the use of this framework.

Future work should also attempt to extend the TSE framework in new directions. One direction is to apply the framework to more statistics, such as variance or mean squared error. Additionally, while the TSE framework has traditionally focused on descriptive statistics, it would be valuable to extend the framework to model-based statistics. Another direction for future work is to address the fact that this framework treats the administrative data as if it

were the truth. Even when there is error in administrative data, the framework is still useful compared to treating the original data as the truth and using that as the only evaluation benchmark. Nonetheless, relaxing this assumption would be valuable. Finally, methods for evaluating data accuracy within a holistic approach to survey error without relying on validation/auxiliary data would help grow the feasibility of this work. Use of validation data will always be limited by availability and subject to questions about their own accuracy.

These extensions may require adopting measurement error models from other disciplines. For example, recent work on measurement error in economics has developed methods for validation studies that relax the assumption that administrative data are without error (Bingley and Martinello, 2017; Crossley, Fisher and Hussein, 2023; Jenkins and Rios-Avila, 2023), methods for evaluating measurement error in model-based statistics rather than descriptive statistics (Bingley and Martinello, 2017; Nguimkeu, Denteh and Tchernis, 2019), and methods for bounds and bias-corrections that avoid the use of validation data altogether (Nguimkeu, Denteh and Tchernis, 2019; Tommasi and Zhang, 2024). These methods have not been extended to study errors in protected data relative to errors in the original data, but there should be opportunities to connect these methods with data privacy.

As discussed in the introduction, assessing data accuracy in a way that facilitates a principled evaluation of the trade-off with data privacy is challenging even when ignoring error that already exists in the original data. The adoption of a Total Survey Error (or related) framework could help address these challenges as well. Accuracy metrics based on differences between the original and protected data can be difficult to communicate to stakeholders because they are implicitly framed as a change from zero error to non-zero error. Stakeholders who interpret the accuracy trade-off in this way may be unlikely to buy into the usability of the protected data. Alternatively, comparing the impact of privacy protection to the impact of other error sources re-frames the trade-off in terms of the *amount* of error rather than the *existence* of error and emphasizes that privacy protection is only one source of survey error. It also provides alternative metrics and practical benchmarks that may be

useful for determining acceptable levels of error from privacy protection.[13]

5 Conclusion

Rising demand for data coupled with rising reconstruction and re-identification risk presents a challenge for statistical agencies such as the Census Bureau. Agencies have long used SDL to protect respondent information, but legacy methods are now seen as insufficient given the increased privacy risks. Regardless of the SDL method, careful attention must be paid to data accuracy. This requires a holistic approach to survey error that moves beyond simply comparing the original and protected data.

In this chapter, we motivated the need for this holistic approach, reviewed a framework for evaluating privacy protection error within the scope of total survey error, and discussed existing research comparing the impact of privacy protection error to measurement error, non-response error, and coverage error. Results in Totty and Watson (2024) based on synthetic data demonstrate that error from privacy protection can appear small when compared to other sources of error and can sometimes offset other sources thereby *reducing* total survey error. This represents a crucial departure from only comparing statistics generated before versus after applying privacy protection, in which case any deviation is often interpreted as an increase from zero error to non-zero error. The results also demonstrate that effects of privacy protection on demographic sub-groups can differ from the effects of other sources of error, thereby having important (and unexpected) impacts on inequality statistics.

We hope this discussion will spur future work on the challenges associated with evaluating data accuracy and the need to place privacy protection error within the larger context of survey error. The literature that addresses this challenge is small (e.g., Agarwal and Singh, 2024; McKinney et al., 2021), but we have described a few areas in which future research could contribute to this work.

[13] E.g., an accuracy metric could be based on how privacy protection changes total survey error in a statistic rather than how it changes the original statistic itself. Other benchmarks could compare the average error from privacy protection to the average error from other components.

References

Abowd, J. M., and M. B Hawes. 2023. "21st Century Statistical Disclosure Limitation: Motivations and Challenges." *arXiv preprint arXiv: 2303.00845*.

Abowd, John M., and Ian M. Schmutte. 2015. "Economic Analysis and Statistical Disclosure Limitation." Brookings Papers on Economic Activity.

Abowd, John M., and Ian M. Schmutte. 2019. "An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices." *American Economic Review*, 109(1): 171–202.

Abowd, John M., Gary L. Benedetto, Simson L. Garfinkel, Scot A. Dahl, Aref N. Dajani, Matthew Graham, Michael B. Hawes, Vishesh Karwa, Daniel Kifer, Hang Kim, Philip Leclerc, Ashwin Machanavajjhala, Jerome P. Reiter, A. Rodríguez, Ian M. Schmutte, William N. Sexton, Phyllis E. Singer, and Lars Vilhuber. 2020. "The modernization of statistical disclosure limitation at the U.S. Census Bureau." U.S. Census Bureau Working Paper.

Abowd, John M., Tamara Adams, Robert Ashmead, David Darais, Sourya Dey, Simson L. Garfinkel, Nathan Goldschlag, Daniel Kifer, Philip Leclerc, Ethan Lew, Scott Moore, Ramy N. Tadros Rolando A. Rodríguez, and Lars Vilhuber. 2023. "The 2010 Census Confidentiality Protections Failed, Here's How and Why." National Bureau of Economic Research working paper 31995.

Agarwal, Anish, and Rahul Singh. 2024. "Causal Inference with Corrupted Data: Measurement Error, Missing Values, Discretization, and Differential Privacy." *arXiv preprint arXiv: 2107.02780*.

Benedetto, Gary, Jordan Stanley, and Evan Totty. 2018. "The Creation and Use of SIPP Synthetic Beta v7.0." Center for Economic Studies, U.S. Census Bureau CES Technical Notes Series 18-03.

Bingley, Paul, and Alessandro Martinello. 2017. "Measurement error in income and schooling and the bias of linear estimators." *Journal of Labor Economics*, 35(4): 1117–1148.

Bound, John, Charles Brown, and Nancy Mathiowetz. 2001. "Measurement error in survey data." In *Handbook of Econometrics Vol. 5.* , ed. James J. Heckman and Edward Leamer, Chapter 59, 3705–3843. North-Holland.

Brenner, H., and K. Nissim. 2014. "Impossibility of differentially private universally optimal mechanisms." *SIAM Journal on Computing*, 43(5): 1513–1540.

Celhay, P., B. D. Meyer, and N. Mittag. 2024. "What leads to measurement errors? Evidence from reports of program participation in three surveys." *Journal of Econometrics*, 238(2): 105581.

Crossley, Thomas F, Paul Fisher, and Omar Hussein. 2023. "Assessing data from summary questions about earnings and income." *Labour Economics*, 81: 102331.

Dinur, I., and K. Nissim. 2003. "Revealing information while preserving privacy." In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 202–210.

Drechsler, Jörg. 2019. "Differential Privacy for Government Agencies – Are We There Yet?" *Journal of the American Statistical Association*, 118(541): 761–773.

Dwork, Cynthia. 2006. "Differential privacy." In *International colloquium on automata, languages, and programming*. 1–12.

Dwork, Cynthia, and Aaron Roth. 2014. "The algorithmic foundations of differential privacy." *Foundations and Trends in Theoretical Computer Science*, 9(3–4): 211–407.

Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. "Calibrating noise to sensitivity in private data analysis." In *Theory of Cryptography: Third Theory of Cryptography Conference*. 265–284.

Eltinge, John L. 2022. "Disclosure protection in the context of statistical agency operations: Data quality and related constraints." *Harvard Data Science Review*, Special Issue 2.

Gong, Ruobin. 2022. "Transparent privacy is principled privacy." *Harvard Data Science Review*, Special Issue 2.

Groves, Robert M., and Lars Lyberg. 2010. "Total survey error: Past, present, and future." *Public Opinion Quarterly*, 74(5): 849–879.

Hawes, Michael B., Evan M. Brassell, Anthony Caruso, Ryan Cumings-Menon, Jason Devine, Cassandra Dorius, David Evans, Kenneth Haase, Michele C. Hedrick, Scott H. Holan, Cynthia D. Hollingsworth, Eric B. Jensen, Dan Kifer, Alexandra Krause, Philip Leclerc, James Livsey, Roberto Ramirez, Rolando A Rodríguez, Luke T. Rogers, Matthew Spence, Victoria Velkoff, Michael Walsh, James Whitehorne, and Sallie Ann Keller. 2024. "Towards a Principled Discussion of a Disclosure Avoidance Framework: Identifying the Characteristics of an Ideal, Applied Disclosure Avoidance System." Working paper. https://conference.nber.org/conf_papers/f193633.pdf.

Heffetz, Ori, and Katrina Ligett. 2014. "Privacy and data-based research." *Journal of Economic Perspectives*, 28(2): 75–98.

Hotz, V. Joseph, Christopher R. Bollinger, Tatiana Komarova, Charles F. Manski, Robert A. Moffit, Denis Nekipelov, Aaron Sojourner, and Bruce D. Spencer. 2022. "Balancing data privacy and usability in the federal statistical system." *Proceedings of the National Academy of Sciences*, 119(31): e2104906119.

Jarmin, Ron S. 2019. "Evolving measurement for an evolving economy: thoughts on 21st century US economic statistics." *Journal of Economic Perspectives*, 33(1): 165–184.

Jenkins, Stephen P, and Fernando Rios-Avila. 2023. "Reconciling reports: modelling employment earnings and measurement errors using linked survey and administrative data." *Journal of the Royal Statistical Society Series A: Statistics in Society*, 186(1): 110–136.

Karr, Alan F. 2017. "The role of statistical disclosure limitation in total survey error." In *Total Survey Error in Practice*. , ed. Paul P. Biemer, Edith de Leeuw, Stephanie Eckman, Brad Edwards, Frauke Kreuter, Lars E. Lyberg, N. Clyde Tucker and Brady T. West, Chapter 4, 71–94. John Wiley & Sons.

Manski, Charles F. 2015. "Communicating uncertainty in official economic statistics: An appraisal fifty years after Morgenstern." *Journal of Economic Literature*, 53(3): 631–653.

McKinney, Kevin L., Andrew S. Green, Lars Vilhuber, and John M. Abowd. 2021. "Total error and variability measures for the quarterly workforce indicators and LEHD origin-destination employment statistics in OnTheMap." *Journal of Survey Statistics and Methodology*, 9(5): 1146–1182.

Meyer, Bruce D., and Nikolas Mittag. 2021. "An empirical total survey error decomposition using data combination." *Journal of Econometrics*, 224(2): 286–305.

Meyer, Bruce D., Nikolas Mittag, and Robert M. George. 2022. "Errors in Survey Reporting and Imputation and their Effects on Estimates of Food Stamp Program Participation." *The Journal of Human Resources*, 57(5): 1605–1644.

Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan. 2015. "Household Surveys in Crisis." *Journal of Economic Perspectives*, 29(4): 199–226.

Nguimkeu, Pierre, Augustine Denteh, and Rusty Tchernis. 2019. "On the estimation of treatment effects with endogenous misreporting." *Journal of Econometrics*, 208(2): 487–506.

Reiter, Jerry P. 2019. "Differential privacy and federal data releases." *Annual review of statistics and its application*, 6(1): 85–101.

Ruggles, Steven, Catherine Fitch, Diana Magnuson, and Jonathan Schroeder. 2019. "Differential Privacy and Census Data: Implications for Social and Economic Research." *AEA Papers and Proceedings*, 109: 403–408.

Schennach, S. M. 2016. "Recent advances in the measurement error literature." *Annual Review of Economics*, 8(1): 341–377.

Snoke, J., G. M. Raab, B. Nowok, C. Dibben, and A. Slavkovic. 2018. "General and specific utility measures for synthetic data." *Journal of the Royal Statistical Society Series A: Statistics in Society*, 118(3): 663–688.

Tommasi, Denni, and Lina Zhang. 2024. "Bounding program benefits when participation is misreported." *Journal of Econometrics*, 238(1): 105556.

Totty, Evan, and Thor Watson. 2024. "Statistical Disclosure Limitation and Total Survey Error." U.S. Census Bureau Working Paper ced-wp-2024-001.

U.S. Census Bureau. 2021*a*. "Census Bureau Sets Key Parameters to Protect Privacy in 2020 Census Results." *https://www.census.gov/newsroom/press-releases/2021/2020-census-key-parameters.html*, Accessed: 2024-07-18.

U.S. Census Bureau. 2021*b*. "DAS 2020 Redistricting Production Code Release." *https://github.com/uscensusbureau/DAS_2020_Redistricting_Production_Code*, Accessed: 2024-07-25.