

**IN THE UNITED STATES DISTRICT COURT
FOR THE MIDDLE DISTRICT OF ALABAMA
EASTERN DIVISION**

STATE OF ALABAMA, *et al.*,

Plaintiffs,

v.

UNITED STATES DEPARTMENT OF
COMMERCE, *et al.*,

Defendants.

Case No. 3:21-CV-211-RAH-ECM-KCN

DECLARATION OF JOHN M. ABOWD

I, John M. Abowd, make the following Declaration pursuant to 28 U.S.C. § 1746, and declare that under penalty of perjury the following is true and correct to the best of my knowledge:

BACKGROUND

1. I am the Chief Scientist and Associate Director for Research and Methodology at the United States Census Bureau. I have served in this capacity since June 2016. My statements in this declaration are based on my personal knowledge or on information supplied to me in the course of my professional responsibilities.
2. I received my Ph.D. in economics from the University of Chicago with specializations in econometrics and labor economics in 1977 (M.A. 1976). My B.A. in economics is from the University of Notre Dame.
3. I have been a university professor since 1976 when I was appointed assistant professor of economics at Princeton University. I was also assistant and associate professor of econometrics and industrial relations at the University of Chicago Graduate School of Business. In 1987, I was appointed associate professor of industrial and labor relations with indefinite tenure at Cornell University where I am currently the Edmund Ezra Day Professor. I am on unpaid leave from Cornell University to work in my current position at the Census Bureau as part of the Career Senior Executive Service.
4. I am a member and fellow of the American Association for the Advancement of Science, American Statistical Association, Econometric Society, and Society of Labor Economists (president 2014). I am an elected member of the International Statistical Institute. I am also a member of the American Economic Association, International Association for Official Statistics, National Association for Business Economists, American Association for Public Opinion Research, Association for Computing Machinery, and American Association of Wine Economists. I regularly attend and present papers at the meetings of these organizations.

5. I have served on the American Economic Association Committee on Economic Statistics. I have also served on the National Academy of Sciences Committee on National Statistics, the Conference on Research in Income and Wealth Executive Committee, and the Bureau of Labor Statistics Technical Advisory Board for the National Longitudinal Surveys (chair: 1999-2001).
6. I have worked with the Census Bureau since 1998, when the Census Bureau and Cornell University entered into the first of a sequence of Intergovernmental Personnel Act agreements and other contracts. Under those agreements, I served continuously as Distinguished Senior Research Fellow at the Census Bureau until I assumed my current position as Chief Scientist in 2016, under a new Intergovernmental Personnel Act contract. Since March 29, 2020, I have been in the Associate Director position at the Census Bureau as a Career Senior Executive Service employee.
7. From 2011 until I assumed my position as Chief Scientist at the Census Bureau in 2016, I was the lead Principal Investigator of the Cornell University node of the NSF-Census Research Network, one of eight such nodes that worked collaboratively with the Census Bureau and other federal statistical agencies to identify important theoretical and applied research projects of direct programmatic importance to the agencies. The Cornell node produced the fundamental science explaining the distinct roles of statistical policymakers and computer scientists in the design and implementation of differential privacy systems at statistical agencies.
8. I have published more than 100 scholarly books, monographs, and articles in the disciplines of economics, econometrics, statistics, computer science, and information science. I have been the principal investigator or co-principal investigator on 35 sponsored research projects. I was a founding editor of the [Journal of Privacy and Confidentiality](#) – an interdisciplinary journal, and I continue to serve as an editor and on the governance board. My full professional resume is attached to this report as Appendix A.

9. I have worked on and managed Census Bureau projects that were precursors to the Census Bureau's current program to implement differential privacy for the 2020 Census of Population and Housing. I was one of three senior researchers who founded the Longitudinal Employer-Household Dynamics (LEHD) program at the Census Bureau, which is generally acknowledged as the Census Bureau's first 21st Century data product: built to the specifications of local labor market specialists without additional survey burden, and published beginning in 2001 using state-of-the-art confidentiality protection via noise infusion. This program produces detailed public-use statistical data on the characteristics of workers and employers in local labor markets using large-scale linked administrative, census, and survey data from many different sources. In 2008, my work with LEHD led to the first production implementation worldwide of differential privacy as part of a product of the LEHD program called OnTheMap. The LEHD program also implemented other prototype systems to protect confidential information, including allowing the public to access synthetic micro-data confirmed via direct analysis of the confidential data on validation servers. A differentially private version of this system is under development at the Census Bureau but not for use with the 2020 Census.

IMPORTANCE OF CONFIDENTIALITY

10. Though participation in the census is mandatory under 13 U.S. Code § 221, in practice, the Census Bureau must rely on the voluntary participation of each household in order to conduct a complete enumeration.
11. One of the most significant barriers to conducting a complete and accurate enumeration are individuals' concerns about the confidentiality of census data. The Census Bureau's pre-2020 Census research showed that 28% of respondents were "extremely concerned" or "very concerned" and a further 25% were "somewhat concerned"

about the confidentiality of their census responses.¹ These concerns are even more pronounced in minority populations and represent a major operational challenge to enumerating traditionally hard-to-count populations.²

12. To secure voluntary participation, Congress first established confidentiality protections for individual census responses in the Census Act of 1879. These confidentiality protections were later expanded and codified in 13 U.S. Code §§ 8(b) & 9, which prohibits the Census Bureau from releasing “any publication whereby the data furnished by any particular establishment or individual under this title can be identified[,]” and allows the Secretary to provide aggregate statistics so long as those data “do not disclose the information reported by, or on behalf of, any particular respondent[.]” Title III of the Foundations for Evidence Based Policymaking Act of 2018 also requires statistical agencies to “protect the trust of information providers by ensuring the confidentiality and exclusive statistical use of their responses.”³
13. The broader scientific community generally concurs about the importance of rigorous protection of confidentiality by statistical agencies. For example, the National Academy of Sciences’ definitive guidebook for federal statistical agencies states “Because virtually every person, household, business, state or local government, and organization is the subject of some federal statistics, public trust is essential for the continued effectiveness of federal statistical agencies. Individuals and entities providing data di-

¹ U.S. Census Bureau (2019) “2020 Census Barriers, Attitudes, and Motivators Study Survey Report” <https://www2.census.gov/programs-surveys/decennial/2020/program-management/final-analysis-reports/2020-report-cbams-study-survey.pdf>, p.38-39.

² Ibid, p.39-42.

³ Title III of the Foundations for Evidence Based Policymaking Act of 2018, § 3563.

rectly or indirectly to federal statistical agencies must trust that the agencies will appropriately handle and protect their information.”⁴ The report also notes that respondents expect statistical agencies not to “release or publish their information in identifiable form.”⁵ The National Academies also broadly exhort statistical agencies to “continually seek to improve and innovate their processes, methods, and statistical products to better measure an ever-changing world.”⁶

14. The Census Bureau enjoys higher self-response rates than private survey companies in large part because the public generally trusts the Census Bureau to keep its data safe. The Census Bureau makes extensive outreach efforts to assure respondents and other data providers about the Bureau’s commitment to protection of confidential data. The criminal fines and imprisonment penalties that Census Bureau employees would face by unlawfully disclosing respondent information are frequently cited by the Census Bureau in these outreach efforts.⁷
15. This trust in the Census Bureau is particularly important for the decennial census, given the “civic ceremony” aspect of the census, akin to the civic ceremony aspect of elections and voting. The decennial census is an exercise where the nation comes together every ten years, under a strict promise of confidentiality, to provide information to help govern our nation. Were the Census Bureau to expose confidential information, there is no doubt that self-response rates would drop, increasing survey

⁴ National Academies of Sciences, Engineering, and Medicine 2021. Principles and Practices for a Federal Statistical Agency: Seventh Edition. Washington, DC: The National Academies Press. <https://doi.org/10.17226/25885>, p. 37-38.

⁵ Ibid., p.38.

⁶ Ibid., p.4.

⁷ <https://www.census.gov/content/dam/Census/library/factsheets/2019/comm/2020-confidentiality-factsheet.pdf>.

cost across programs by increasing in-person follow up, and decreasing the quality of the census overall.

PRIVACY PROTECTION AT THE CENSUS BUREAU

16. Protecting privacy is at the core of the Census Bureau’s mission. Our privacy promise to respondents is key to promoting response to our censuses and surveys. The Census Bureau – at the crux of its dual mandate to publish only statistical summaries and to protect the confidentiality of respondent data – is balancing the preferences of data users and data providers. An optimal choice must account for the preferences of data users and protect the data the American people entrust the Census Bureau with keeping safe.⁸
17. Data collected from the decennial census support a wide array of critical government and societal functions at the federal, state, tribal, and local levels. In addition to apportioning seats in the U.S. House of Representatives and supporting the redistricting of those seats, census data also support the allocation of over \$675 billion in federal

⁸ “Official Statistics at the Crossroads: Data Quality and Access in an Era of Heightened Privacy Risk,” *The Survey Statistician*, 2021, Vol. 83, 23-26 (available at [Survey Statistician 2021 January N83_03.pdf \(isi-iass.org\)](https://www.isi-iass.org/Survey-Statistician-2021-January-N83-03.pdf)). The paper is based on talks that I gave in 2019 to the Committee on National Statistics and the Joint Statistical Meetings. It summarizes the research in Abowd, J.M. and I. Schmutte “An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices,” *American Economic Review*, Vol. 109, No. 1 (January 2019):171-202, DOI:[10.1257/aer.20170627](https://doi.org/10.1257/aer.20170627).

funding each year based on population counts, geography, and demographic characteristics.⁹ Census data also support important public and private sector decision-making at the federal, state, tribal, and local levels, and serve as benchmark statistics for other important surveys and data collections throughout the decade.¹⁰

18. The Census Bureau publishes an enormous number of statistics calculated from its collected data. Following the 2010 Census, for example, the Census Bureau published over 150 billion independent statistics about the characteristics of the 308,745,538 persons in the resident population that were enumerated in the census. To serve their intended governmental and societal uses, the majority of these statistics needed to be published at very fine levels of detail and with geographic precision often down to the individual census tract or block.

19. While it would be quite difficult from any single one of those published statistics to ascertain the identity of any individual census respondent or the contents of that respondent's census response, the volume and detail of information published by the Census Bureau, taken together, pose a serious challenge for protecting the privacy and confidentiality of census data. Combining information from multiple published statistics or tables can make it easy to pick out those individuals in a particular geographic area whose characteristics differ from those of the rest of their neighbors. These individuals, who have unique combinations of the demographic characteristics

⁹ Hotchkiss, M., & Phelan, J. (2017). Uses of Census Bureau data in federal funds distribution: A new design for the 21st century. United States Census Bureau. <https://www2.census.gov/programs-surveys/decennial/2020/program-management/working-papers/Uses-of-Census-Bureau-Data-in-Federal-Funds-Distribution.pdf>.

¹⁰ Sullivan, T. A. (2020). Coming to Our Census: How Social Statistics Underpin Our Democracy (and Republic). *Harvard Data Science Review*, 2(1). <https://doi.org/10.1162/99608f92.c871f9e0>.

reported in statistical summaries, are known as “population uniques” and their records have traditionally been the target of the mechanisms that the Census Bureau uses to protect confidentiality in its data publications.

20. Traditional statistical disclosure limitation methods,¹¹ like those used in 2010 census, cannot defend against modern challenges posed by enormous cloud computing capacity and sophisticated software libraries. That does not mean traditional statistical disclosure limitation methods usually fail – they usually do not fail. But as computer scientists bring their expertise from the field of cryptography to the field of safe data publication, they have exposed significant vulnerabilities in traditional privacy methods. The Census Bureau’s own internal analysis, for example, confirmed that a modern database reconstruction-abetted re-identification attack can reliably match a large number of 2010 census responses to the names of those respondents – a vulnerability that exposed information of *at least* 52 million Americans and potentially up to 179 million Americans.¹² To defend against this known vulnerability, the Census Bureau explored different confidentiality methods that explicitly defend against database reconstruction attacks and concluded that the best tool to protect against this modern attack while also preserving the accuracy and usability of data products comes from the body of scientific work called “differential privacy.”

THE HISTORY OF INNOVATION IN THE DECENNIAL CENSUS

21. The decennial census, known officially as the *Decennial Census of Population and Housing*, is the flagship statistical product of the U.S. Census Bureau. Though the Census

¹¹ The technical field that addresses confidentiality is known as “statistical disclosure limitation.” At the Census Bureau, it is known as “disclosure avoidance.” It is also called “statistical disclosure control” by some statisticians and “privacy-preserving data analysis” by some computer scientists.

¹² See Appendix B for a summary of the Census Bureau’s simulated reconstruction and re-identification attacks.

Bureau conducts hundreds of surveys every year, the once-every-decade enumeration of the population of the United States, mandated by Article I, Section 2 of the U.S. Constitution, is the single largest and most complex data collection regularly conducted by the United States government. Since the very first U.S. census in 1790, the collection, processing, and dissemination of census data have posed unique challenges and have required the Census Bureau to improve its operations every decade.

22. The challenges faced by the Census Bureau have led to remarkable innovations. Herman Hollerith's electric tabulation machine, developed for the 1890 Census, revolutionized the field of data processing and led Hollerith to form the company that eventually became IBM.¹³ To conduct the 1950 Census, the Census Bureau commissioned the development of the first successful civilian digital computer, UNIVAC I.¹⁴ With each passing decade, the Census Bureau develops, tests, and deploys innovations to its statistical methods, field data collection methods, and data processing operations.

23. That spirit of innovation includes the Census Bureau's more recent implementation of cutting-edge privacy protections. Prior to the 1990 Census, the primary mechanism that the Census Bureau employed to protect the confidentiality of individual census responses was to withhold publication of (or "suppress") any table that did not meet certain household, population, or demographic characteristic thresholds. The 1970 Census, for example, suppressed tables reflecting fewer than five households, and would only publish tables of demographic characteristics cross-tabulated by race if

¹³ https://www.census.gov/history/www/census_then_now/notable_alumni/herman_hollerith.html.

¹⁴ https://www.census.gov/history/www/innovations/technology/univac_i.html.

there were at least five individuals in each reported race category.¹⁵ These suppression routines helped to protect privacy by reducing the detail of data published about individuals who were relatively unique within their communities. By the 1990 Census, however, the Census Bureau transitioned away from suppression methodologies for two reasons: first, data users were dissatisfied with missing details caused by suppression and second, the Bureau realized that the suppression routines it had been using were insufficient to fully protect against re-identification.¹⁶

24. For the 1990 Census, the Bureau began using a technique known as noise infusion to safeguard respondent confidentiality. Noise infusion helps to protect the confidentiality of published data by introducing controlled amounts of error or “noise” into the data. The goal of noise infusion is to preserve the overall statistical validity of the resulting data while introducing enough uncertainty that attackers would not have any reasonable degree of certainty that they had isolated data for any particular respondent. The noise infusion used in 1990 was a very simple form of data swapping between paired households in a geographic area with similar attributes, and for small

¹⁵ Zeisset, P. (1978), “Suppression vs. Random Rounding: Disclosure Avoidance Alternatives for the 1980 Census,” <https://www.census.gov/content/dam/Census/library/working-papers/1978/adrm/Suppression%20vs.%20Random%20Rounding%20Disclosure-Avoidance%20Alternatives%20for%20the%201980%20Census.pdf>.

¹⁶ McKenna, L. (2018), “Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing,” <https://www.census.gov/content/dam/Census/library/working-papers/2018/adrm/Disclosure%20Avoidance%20for%20the%201970-2010%20Censuses.pdf>, p.6.

block groups the Census Bureau replaced the collected characteristics of households with imputed characteristics.¹⁷

25. For the 2000 and 2010 censuses, the Census Bureau began to infuse noise using a more advanced “data swapping” method. The Census Bureau first identified households most vulnerable to re-identification—especially households on smaller-population blocks whose residents had differing demographic characteristics from the remainder of their block. While every non-imputed¹⁸ household record in the Census Edited File (CEF) had a chance of being selected for data swapping, records for more vulnerable households (typically those on low-population blocks) were selected with greater probability. Then, the records for all members of those selected households were exchanged with the records of households in nearby geographic areas that matched on key characteristics. For the 2000 and 2010 censuses, those key matching characteristics were (1) the whole number of persons in the household, and (2) the whole number of persons aged 18 or older in the household. These swapping criteria resulted in the total population and total voting age population for each block being held “invariant” —that is, while noise was added to all remaining characteristics, no noise was added to the block-level total population or block-level voting age population

¹⁷ Ibid., p. 6-7. An “imputed characteristic” is the prediction of a statistical model used in place of a missing characteristic, when used in standard editing procedures, or in place of a collected characteristic, when used for confidentiality protection.

¹⁸ When a respondent household provides only a count of the number of persons living at that address or when the housing unit population count is itself imputed, the Census Bureau imputes all characteristics: sex, age, race, ethnicity, and relationship to others in the household. Such persons are called “whole-person census imputations” in technical documentation. When a household consists entirely of whole-person census imputation records, it is called an “imputed” household. A “non-imputed” household contains at least one person whose characteristics were collected on the census form for the household.

counts.¹⁹ *The selection and application of these particular invariants is not an innate feature of data swapping; invariants are implementation parameters that can be applied to (or removed from) any counted characteristic under any noise infusion methodology.*

THE PRIVACY PROTECTIONS USED FOR THE 2010 CENSUS ARE NO LONGER SUFFICIENT

26. While the Census Bureau's confidentiality methodologies for the 2000 and 2010 censuses were considered sufficient at the time, advances in technology in the years since have reduced the confidentiality protection provided by data swapping.
27. Disclosure avoidance has been a recognized branch of statistics since the 1970s, but it has only been since the late 1990s that it has evolved into a distinct scientific field of study in both statistics and computer science. Prof. Latanya Sweeney's 1997 revelation that she had re-identified then Massachusetts Governor William Weld's medical records in a purportedly "deidentified" public database²⁰ prompted the Census Bureau and many other statistical agencies to re-examine the efficacy of their disclosure avoidance techniques.
28. *Re-identification attacks.* Prior to 2016, disclosure risk assessments usually focused on assessing the vulnerability of microdata releases (data products that contain individual records for all or some of the data subjects deidentified by removing names and addresses), rather than the rules used for aggregated data releases (data compiled and aggregated into tables). Simulated "re-identification attacks" analyze the risk that an external attacker could use individuals' characteristics that are included on a published microdata file (e.g., location, age, and sex) and link those records to a third-

¹⁹ Ibid. p. 8-10.

²⁰ Sweeney, L. (2002). "k-anonymity: a model for protecting privacy." *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10 (5); 557-570, also recounted in Ohm, P. (2009) "Broken promises of privacy: Responding to the surprising failure of anonymization." *UCLA L. Rev.* 57: 1701.

party data source (e.g., commercial data or voter registration lists) that contains those characteristics along with the individuals' names and addresses. The resulting rates of "putative" (suspected) and confirmed linkages show the overall degree of vulnerability of the data. If those linkage rates are deemed too large, then additional disclosure avoidance is necessary to mitigate the disclosure risk.

29. The general problem with relying exclusively on re-identification studies to assess disclosure risk is that they can only provide a "best-case" approximation of the underlying disclosure risk of the data. If a real attacker has access to more sophisticated tools (e.g., optimization algorithms or computing power) or to higher quality external data (e.g., with better age and address information) than the tools or data used in the simulated attack, then the real disclosure risk will be substantially higher than what is estimated via the study. This limitation is particularly vexing for statistical agencies that must rely on a "release and forget" approach to data publication, where disclosure avoidance safeguards must be selected without foreknowledge of the better tools and external data that attackers may have at their disposal after the data are published.

30. Re-identification studies also underestimate the risk from releasing aggregated data. The Census Bureau has long relied on re-identification studies to assess the disclosure risk of its microdata releases, but the majority of Census Bureau data products are aggregated data releases. Over the past decade, aggregated data releases have become increasingly vulnerable to sophisticated "reconstruction attacks" that have emerged as computing power has improved and gotten substantially cheaper.

31. *Reconstruction attacks.* The theory behind a “reconstruction attack” is that the release of *any* statistic calculated from a confidential data source will reveal a potentially trivial, but non-zero, amount of confidential information.²¹ As a consequence, if an attacker has access to enough aggregated data with sufficient detail and precision, then the attacker may be able to leverage information from each statistic in the aggregated data to reconstruct the individual-level records that were used to generate the published tables. This process is known as a “reconstruction attack,” and it adds a new degree of disclosure vulnerability against which statistical agencies must defend. While the statistical and computer science communities have been aware of this vulnerability since 2003, only over the last few years have computing power and the sophisticated numerical optimization software necessary to perform these types of reconstructions advanced enough to permit reconstruction attacks at any significant scale.

32. The risk of reconstruction and re-identification attacks is real and substantiated. The Census Bureau has been approached by Prof. Sweeney and others who claim that they have identified specific vulnerabilities in our standard disclosure avoidance methodologies.²² The vulnerabilities in the disclosure avoidance protections for the Census Bureau’s Survey of Income and Program Participation (SIPP) identified by Prof. Sweeney led the Census Bureau to immediately implement permanent changes to the

²¹ Dinur, I. and Nissim, K. (2003) “Revealing Information while Preserving Privacy” PODS, June 9-12, San Diego, CA. <https://doi.org/10.1145/773153.773173>.

²² McKenna, L. (2019b). “U.S. Census Bureau Reidentification Studies,” available at <https://www.census.gov/library/working-papers/2019/adrm/2019-04-ReidentificationStudies.html>.

disclosure avoidance rules used for SIPP data, including increased noise infusion and delayed reporting of survey participants' major life events.²³

33. Statistical releases do not all need to be of the same type, or contain the same data fields, to enable re-identification by reconstruction. For example, a 2015 interagency report published by the National Institute of Standards and Technology (NIST) written by my colleague Simson Garfinkel provided examples of using disparate data sets to reconstruct hidden underlying data.²⁴ Some of these examples are quoted here:

34. "*The Netflix Prize*: Narayanan and Shmatikov showed in 2008 that in many cases the set of movies that a person had watched could be used as an identifier.²⁵ Netflix had released a dataset of movies that some of its customers had watched and ranked as part of its "Netflix Prize" competition. Although there was [sic] no direct identifiers in the dataset, the researchers showed that a set of movies watched (especially less popular films, such as cult classics and foreign films) could frequently be used to match a user profile from the Netflix dataset to a single user profile in the Internet Movie Data Base (IMDB), which had not been de-identified and included user names, many of which were real names. The threat scenario is that by rating a few movies on IMDB, a person might inadvertently reveal *all* of the movies that they had watched, since the person's IMDB profile could be linked with the Netflix Prize data."²⁶ (emphasis in original)

²³ McKenna, L. (2019b). p. 2-3.

²⁴ Garfinkel, S. (2015) "De-Identification of Personal Information," National Institute of Standards and Technology, available at <http://dx.doi.org/10.6028/NIST.IR.8053> at 26-27.

²⁵ Narayanan, A. and Shmatikov V. "Robust De-anonymization of Large Sparse Datasets," *IEEE Symposium on Security and Privacy* (2008): 111-125.

²⁶ Garfinkel, S. (2015), p. 26-27.

35. “*Credit Card Transactions*: Working with a collection of de-identified credit card transactions from a sample of 1.1 million people from an unnamed country, Montjoye *et al.* showed that four distinct points in space and time were sufficient to specify uniquely 90% of the individuals in their sample.²⁷ Lowering the geographical resolution and binning transaction values (*e.g.*, reporting a purchase of \$14.86 as between \$10.00 and \$19.99) increased the number of points required.”²⁸
36. “*Mobility Traces*: Montjoye *et al.* showed that people and vehicles could be identified by their “mobility traces” (a record of locations and times that the person or vehicle visited). In their study, trace data from a sample of 1.5 million individuals was processed, with time values being generalized to the hour and spatial data generalized to the resolution provided by a cell phone system (typically 10-20 city blocks).²⁹ The researchers found that four randomly chosen observations of an individual putting them at a specific place and time was sufficient to uniquely identify 95% of the data subjects.³⁰ Space/time points for individuals can be collected from a variety of sources, including purchases with a credit card, a photograph, or Internet usage. A similar study performed by Ma *et al.* found that 30%-50% of individuals could be identified with 10 pieces of side information.³¹ The threat scenario is that a person who

²⁷ Montjoye, Y-A. et al. “Unique in the shopping mall: On the reidentifiability of credit card metadata,” *Science*, 30 (January 2015) Vol 347, Issue 6221.

²⁸ Garfinkel, S. (2015), p. 27.

²⁹ De Montjoye, Y. A., Hidalgo, C. A., Verleysen, M., & Blondel, V. D. (2013). Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 3(1).

³⁰ *Ibid.*, p. 1-5.

³¹ C. Y. T. Ma, D. K. Y. Yau, N. K. Yip and N. S. V. Rao (2013) "Privacy Vulnerability of Published Anonymous Mobility Traces," in *IEEE/ACM Transactions on Networking*, vol. 21, no. 3, pp. 720-733, June 2013, doi: 10.1109/TNET.2012.2208983.

revealed five place/time pairs (perhaps by sending email from work and home at four times over the course of a month) would make it possible for an attacker to identify his or her entire mobility trace in a publicly released dataset. As above, the attacker would need to know that the target was in the data.”³²

37. The same general principles apply to census data. The difference between census data and the examples above is that census data can be combined in vastly more ways with other information because all the tables published from census data share basic standardized identifiers including location, age, sex, race, ethnicity, and marital status. Even if each of these identifiers is not included in every table, their use and combinations across many different tables creates the disclosure risk. The Census Bureau understood this emerging risk even before the 2010 Census. As field collection for the 2010 Census was first beginning, the Census Bureau had already flagged the heightened disclosure risk of releasing detailed block level population data, even with the 2010 Census swapping mechanism in place.³³ After tracking this growing risk of reconstruction and re-identification attacks for several years, the Census Bureau decided in 2015 to establish a new team to comprehensively evaluate the Census Bureau’s disclosure avoidance methods to determine if they were sufficient to protect against these disclosure risks.³⁴

³² Garfinkel, S. (2015), p. 27-28.

³³ During a January 2010 meeting of the Census Bureau’s Data Stewardship Executive Policy (DSEP) Committee, the chair of the Disclosure Review Board voiced her concerns about the 2010 Census swapping mechanism’s ability to adequately protect future censuses, noting specifically the challenge posed by “continuing to release data at the block level, as block populations continue to decrease (e.g., 40% of blocks in North Dakota have only 1 household in them)” Based on this warning, DSEP decided that “the problem of block population size and disclosure avoidance is real, and that it deserves attention in the context of 2020 planning.” DSEP Meeting Record, January 14, 2010. See Appendix C.

³⁴ DSEP Meeting Record, February 5, 2015. See Appendix D.

2010 CENSUS SIMULATED RECONSTRUCTION-ABETTED RE-IDENTIFICATION ATTACK

38. The results from the Census Bureau's 2016-2019 research program on simulated reconstruction-abetted re-identification attack were conclusive, indisputable, and alarming. Appendix B, attached to this declaration, provides an overview of that simulation and the results. The bottom line is that our simulated attack showed that a conservative attack scenario using just 6 billion of the over 150 billion statistics released in 2010 would allow an attacker to accurately re-identify *at least* 52 million 2010 Census respondents (17% of the population) and the attacker would have a high degree of confidence in their results with minimal additional verification or field work. In a more pessimistic scenario, an attacker with access to higher quality commercial name and address data than those used in our simulated attack could accurately re-identify around 179 million Americans or around 58% of the population.
39. Emerging attack scenarios and our own internal simulated attacks show that were the Census Bureau to use the disclosure avoidance mechanism implemented for the 2010 Census again for the 2020 Census, the results would be vulnerable to reconstruction and re-identification attacks because of the parameters of the swapping mechanism's 2010 implementation: an overall insufficient level of noise, the invariants preserved without noise, and the geographic and demographic detail of the published summary data. The Census Bureau can no longer rely on the swapping implementation used in 2010 if it is to meet its obligations to protect respondent confidentiality under 13 U.S. Code §§ 8(b) & 9. Protecting against new technology-enabled re-identification attacks, while maintaining the high quality of the decennial census data products, requires the implementation of a disclosure avoidance mechanism that is better able to protect against these new, sophisticated vectors of attack.

DISCLOSURE AVOIDANCE OPTIONS CONSIDERED FOR THE 2020 CENSUS

40. Faced with this compelling mathematical and empirical evidence of the inherent vulnerability of the 2010 Census swapping mechanism to protect against reconstruction-abetted re-identification attacks, the Census Bureau began exploring the available data protection strategies that it could employ for the 2020 Census. The three methods the Census considered were *Enhanced Data Swapping*, *Suppression*, and *Differential Privacy*.
41. The Census Bureau decided that differential privacy was the best tool after analyzing the various options through the lens of economics. Efficiently protecting privacy can be viewed as an economic problem because it involves the allocation of a scarce resource—confidential information—between two competing uses: public data products and privacy protection. If we produce more accuracy, we will have less privacy, and vice versa. And just like in the classic economic example of the trade-off between producing guns and butter, the tradeoff between privacy and accuracy can be analyzed with a production possibility curve. Our empirical analysis showed that differential privacy offered the most efficient trade-off between privacy and accuracy—our calculations showed that the efficiency of differential privacy dominated traditional methods.³⁵ In other words, regardless of the level of desired confidentiality, differential privacy will always produce more accurate data than the alternative traditional methods considered by the Census Bureau.
42. *Enhanced Data Swapping*. Enhancing the data swapping mechanism used for the 2010 Census in a manner sufficient to protect against emerging threats like reconstruction

³⁵ See Abowd, J. M., & Schmutte, I. M. (2019). An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, 109(1), 171-202.

attacks would have a significant, detrimental impact on data quality. With an estimated 57% of the population³⁶ known to be unique at the block level, a swapping mechanism that targets vulnerable households for swapping would require significantly higher rates of swapping than were used in 2010 to protect against a reconstruction attack. Implementing swapping in 2020 would also require abandoning the total population and voting-age population invariants that were used in 2010. There are two technical reasons for this. First, at swap rates sufficient to counter the reconstruction of microdata accurate enough to enable large-scale reidentification, it is impossible to find enough paired households with the same number of persons and adults without searching well outside the neighborhood of the original household. Finding swap pairs was a challenge for some states even at the 2010 swap rate. Second, holding the total and adult populations invariant gives the attacker a huge reconstruction advantage—exact record counts in each block for persons and adults. This advantage vastly improves the accuracy of the reconstructed data. Even a small amount of uncertainty about the block location of an individual greatly expands the variability in the reconstructed microdata effectively reducing the chances of a correct linkage in a re-identification attack. If a block is known to contain exactly seven persons in the confidential data, then every feasible reconstructed version of those data will have exactly seven records in that block, meaning that the block identifier will be correct on every record of every feasible reconstructed database. But if the block population is reported with some random fluctuation around seven, then only by chance will the

³⁶ Fifty-seven percent of the 308,745,538 person records in the confidential 2010 Census Edited File, the definitive source for all 2010 Census tabulations, were unique on their block location, sex, age (in years), race (any combination of the 6 OMB-approved race categories, 63 possibilities in all) and Hispanic/Latino ethnicity. This previously confidential statistic was approved for publication with DRB clearance number CBDRB-FY21-DSEP-003.

block identifier be correct in the reconstructed data. Compound this effect over 8,000,000 blocks and the number of feasible reconstructions explodes exponentially. This is what provides the protection against re-identification from the reconstructed data.³⁷ Internal experiments also confirmed that increasing the swap rate from the level used in 2010 and removing the invariants on block-level population counts (to permit the increased level of swapping and protect against reconstruction attacks) would render the resulting data unusable for most data users.

43. *Suppression*. While the Census Bureau could use suppression to protect from a reconstruction attack, the resulting data would be only available at a very high level of generality. Today's data users, including redistricters, rely on detailed block and tract-level data, which would not be available for many areas if the Census were to return to suppression to protect against modern attacks.
44. *Differential Privacy*. Differential privacy, first developed in 2006, is a framework for quantifying the precise disclosure risk associated with each incremental release from a confidential data source.³⁸ In turn, this allows an agency like the Census Bureau to quantify the precise amount of statistical noise required to protect privacy. This precision allows the Census to calibrate and allocate precise amounts of statistical noise in a way that protects privacy while maintaining the overall statistical validity of the data.

³⁷ Garfinkel, S., Abowd, J. M., & Martindale, C. (2018). Understanding Database Reconstruction Attacks on Public Data: These attacks on statistical databases are no longer a theoretical danger. *Queue*, 16(5), 28-53.

³⁸ Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006, March). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference* (pp. 265-284). Springer, Berlin, Heidelberg.

45. The Census Bureau first began using differential privacy to protect its statistical data products in 2008, with the launch of its [OnTheMap](#) tool for employee commuting statistics and its heavily used extension [OnTheMap for Emergency Management](#). In the years since, the Census Bureau has also successfully used differential privacy in a number of other innovative statistical products, such as the Post-Secondary Employment Outcomes and Veteran Employment Outcomes products. Differential privacy is also being used by many of the major technology firms, including Apple³⁹, Google,⁴⁰ Microsoft,⁴¹ and Uber.⁴² Other statistical agencies, such as the Statistics of Income Division of the Internal Revenue Service, have also begun implementing differential privacy.⁴³ Internationally, the Australian Bureau of Statistics,⁴⁴ the Office of National

³⁹Differential Privacy Team. (2017). "Learning with Privacy at Scale." *Apple Machine Learning Journal*, 1(8).

⁴⁰Erlingsson, U., V. Pihur, and A. Korolova. (2014). "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response." *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security - CCS '14*, 1054–1067.

⁴¹ Ding, B., J. Kulkarni, and S. Yekhanin. (2017). "Collecting Telemetry Data Privately." *Advances in Neural Information Processing Systems* 30.

⁴² Near, J. (2018) "Differential Privacy at Scale: Uber and Berkeley Collaboration," *Enigma 2018* (January) USENIX Assoc. <https://www.usenix.org/node/208168>.

⁴³ Bowen, C. et al. (2020) "A Synthetic Supplemental Public-Use File of Low-Income Information Return Data: Methodology, Utility, and Privacy Implications," (July) Tax Policy Center, The Brookings and Urban Institutes. https://www.urban.org/sites/default/files/publication/102547/a-synthetic-supplemental-public-use-file-of-low-income-information-return-data_2.pdf.

⁴⁴ Australian Bureau of Statistics, (2019) "Protecting the Confidentiality of Providers," January 2019, 1504.0 - *Methodological News*, <https://www.abs.gov.au/ausstats/abs@.nsf/Previousproducts/1504.0Main%20Features9999Jan%202019?opendocument&tabname=Summary&prodno=1504.0&issue=Jan%202019&num=&view=>, accessed on March 31, 2021.

Statistics in the United Kingdom,⁴⁵ and Statistics Canada⁴⁶ explicitly recognize the threat from combining multiple statistical tabulations to re-identify respondent information and recommend output noise infusion systems, including differential privacy.

46. Faced with the alarming results of the simulated reconstruction attack, which indicated that the established swapping mechanism resulted in far less disclosure protection than it was intended to provide, and considering the available alternatives, the Census Bureau's Data Stewardship Executive Policy Committee (DSEP)⁴⁷ determined that the Census Bureau should proceed with the deployment and testing of differential privacy for use in the 2020 Census given its obligations to produce high quality statistics from the decennial census while also protecting the confidentiality of respondents' census records under 13 U.S. Code §§ 8(b) & 9.⁴⁸

⁴⁵ United Kingdom Office for National Statistics, (2021) "Policy on Protecting Confidentiality in Tables of Birth and Death Statistics," <https://www.ons.gov.uk/methodology/methodologytopicsandstatisticalconcepts/disclosurecontrol/policyonprotectingconfidentialityintablesofbirthanddeathstatistics#annex-a-understanding-the-legal-and-policy-framework>, accessed on March 31, 2021.

⁴⁶ Statistics Canada, (2021) "A Brief Survey of Privacy Preserving Technologies," March 2021, *Data Science Network for the Federal Public Service*, <https://www.statcan.gc.ca/eng/data-science/network/privacy-preserving>, accessed on March 31, 2021.

⁴⁷ The Data Stewardship Executive Policy Committee (DSEP) is a committee chaired by the Deputy Director/Chief Operating Officer and composed of career senior executives with expertise in confidentiality practice, the uses of Census Bureau data, and policy. DSEP is the parent organization for the Disclosure Review Board (DRB), which reviews and approves individual data releases to ensure that no confidential data is released.

⁴⁸ On May 10-11, 2017 DSEP decided that "any request for disclosure avoidance of proposed publications for the 2020 Census be routed to the 2020 DAS team before going to the DRB" meaning that all 2020 Census publications would be subject to differential privacy. See Appendices E and F. On February 15, 2018 DSEP suspended publication of "all proposed tables in Summary File 1 and Summary File 2 for the 2020 Census at the block, block-group, tract, and county level except for the PL94-171 tables, as announced in Federal Register Notice 170824806-7806-01..." acknowledging that "...these data in many

47. The best disclosure avoidance option that offers a solution capable of addressing the new risks of reconstruction-abetted re-identification attacks, while preserving the fitness-for-use of the resulting data for the important governmental and societal uses of census data, is differential privacy. I have summarized here what I consider to be the most important reasons that the Census Bureau decided to adopt differential privacy.
48. **Disclosure avoidance must be proactive.** The fundamental objective of disclosure avoidance protections is to proactively prevent disclosures. Just like corporations are not expected to wait until they have suffered a major data breach before upgrading their IT security systems to protect against known threats, statistical agencies should not wait until they suffer a confirmed breach before improving their disclosure avoidance protections to account for known threats. The expectation, for both IT security and disclosure avoidance, is to remain vigilant about emerging threats and risks, and to take appropriate action *before* those risks lead to a breach.
49. **The privacy risk landscape has fundamentally changed since 2010.** Traditional methods of assessing disclosure risk rely on knowing what tools and resources an attacker might leverage to undermine confidentiality protections. These tools, however, are ever evolving. Over the last decade, technological advances have made powerful cloud computing environments, with sophisticated optimization algorithms

cases were accurate to a level that was not supported by the actual uses of those data, and such an approach is simply untenable in a formally private system.” DSEP further decided that “SF1 and SF2 will be rebuilt based on use cases.” See Appendix G. In parallel with these decisions by DSEP, the disclosure risks identified by the preliminary results of the simulated reconstruction attack also led to this issue being added to the Census Bureau’s risk management portfolio. On April 17, 2017 the risk of reconstruction attacks was proposed for inclusion in the Research and Methodology Directorate’s risk registry. On September 12, 2017 it was escalated and included on the Enterprise-level Risk register. Finally, on January 30, 2018, it was further escalated to the Enterprise-level Issue register, with the development and use of the 2020 Census Disclosure Avoidance System as an identified resolution action to be taken. .

capable of performing large-scale attacks, cheap and easily available. While these tools were not yet a viable attack model in 2010, they certainly represent a credible threat today.⁴⁹

50. **Internal research has conclusively proven the fundamental vulnerabilities of the 2010 swapping methodology.** The Census Bureau has performed extensive empirical analysis of the disclosure risk inherent to the 2010 Census swapping methodology as detailed in Appendix B. No technique can produce usable data with absolutely zero risk of re-identification, but the re-identification rates from our internal experiments on the 2010 Census swapping methodology are orders of magnitude higher than what they were intended to be. The privacy threat landscape has evolved over the last decade and compels the Census Bureau to adapt its protections accordingly.

51. **The Census Bureau determined that differential privacy was the only method that could adequately protect the data while preserving the value of census data products.** When our internal research demonstrated the vulnerabilities of the swapping mechanism used for the 2010 Census, we considered a range of options for the 2020 Census. The three leading options were differential privacy, an enhanced version of data swapping, and a return to whole-table suppression. But to achieve the necessary level of privacy protection, both enhanced data swapping and suppression had severely deleterious effects on data quality and availability. With its enhanced privacy protections and precision control over the tuning of privacy/accuracy tradeoff, the Census Bureau determined that differential privacy was the only viable solution for the 2020 Census.

⁴⁹ DSEP drew this conclusion from the simulated reconstruction-abetted re-identification attack in Appendix B. The Office of National Statistics reached the same conclusion in its 2018 “Privacy and data confidentiality methods: a Data and Analysis Method Review (DAMR)” at [Privacy and data confidentiality methods: a Data and Analysis Method Review \(DAMR\) – GSS \(civilservice.gov.uk\)](#) (cited on April 10, 2021).

52. Differential privacy can be fine-tuned to strike a balance between privacy and accuracy. DSEP made the preliminary decision to pursue differential privacy on May 10-11, 2017. Since that decision was announced, the Census Bureau has worked extensively with our advisory committees, federal agency partners, American Indian and Alaska Native tribal leaders, the Committee on National Statistics, professional associations, data user groups, and many others at the national, state, and local levels to understand how they use decennial census data and to ensure that our implementation of differential privacy will preserve the value of the decennial census as a national resource. The Census also released sets of demonstrative data to allow the public and end-users to provide feedback that allowed us to fine-tune and tweak how we will ultimately implement differential privacy.⁵⁰

53. The need to modernize our privacy protections has been confirmed by external experts. The Census Bureau's ongoing partnerships with scientific and academic experts from around the country helped us conduct the internal evaluation of the disclosure risk of the 2010 Census swapping methodology and confirmed the need to modernize our privacy protections. To supplement this ongoing work and to get external expert confirmation of the conclusions that we have drawn from it, the Census Bureau also commissioned an independent expert review by JASON, an independent group of elite scientists that advise the federal government on science and technology. The JASON report confirmed our findings regarding the re-identification risk inherent to the 2010 Census swapping methodology.⁵¹

⁵⁰ U.S. Census Bureau "Developing the DAS: Demonstration Data and Progress Metrics" <https://www.census.gov/programs-surveys/decennial-census/2020-census/planning-management/2020-census-data-products/2020-das-development.html>.

⁵¹ JASON (2020). "Formal Privacy Methods for the 2020 Census" JASON Report JSR-19-2F. <https://www2.census.gov/programs-surveys/decennial/2020/program-management/planning-docs/privacy-methods-2020-census.pdf>.

54. **Differential Privacy can produce highly accurate data.** One key benefit of differential privacy is the ability to fine-tune privacy and accuracy. The next iteration of demonstration data will establish that differential privacy protections can produce extremely accurate redistricting data. While the full April 2021 Demonstration Data Product⁵² and supporting metrics will be released by April 30, 2021, I can provide a high-level summary of key metrics:⁵³

- Total populations for counties have an average error of +/- 5 persons (reflecting a mean absolute percent error of 0.04% of the counties' population) as noise from differential privacy.⁵⁴ This is extremely accurate considering that if we simulate the errors in census counts as estimates of the true population, then the average county-level estimation uncertainty of the census is +/- 960 persons (averaging 1.6% of the county census counts).⁵⁵

⁵² The April 2021 demonstration data uses a global privacy-loss budget of 10.3 with a very substantial proportion allocated to detailed race and ethnicity statistics at the block and block group levels.

⁵³ Statistics for the April 2021 Demonstration Data Product are preliminary, based on the internal research version. The production version will be used for the detailed summary statistics when they are posted on census.gov.

⁵⁴ The statistics are the mean absolute error and the mean absolute percentage error in county population comparing the April 2021 Demonstration Data Product to the data released in the 2010 Summary File 1.

⁵⁵ The inherent error in the census counts as estimates of the true population can be simulated using data-defined person and correct-enumeration rates from coverage measurement estimates, in this case from the most recent decennial census in 2010. (See Mule, T. "2010 Census Coverage Measurement Estimation Report: Summary of Estimates of Coverage for Persons in the United States", Report G-10, g01.pdf (census.gov). Table 3, in particular.) An alternative modeling perspective simulates the natural variation of census population estimates using the natural variation in census estimates due to erroneous enumerations and other sources of error inherent in the Census. For county populations

- At the block level the differentially private data have an average population error of +/- 3 persons, which includes both housing unit and group quarters populations. Compare that with the simulated error inherent in the census which puts the average error uncertainty of block population counts at +/- 6 people.⁵⁶

55. **The April 2021 demonstration data show no meaningful bias in the statistics for racial and ethnic minorities** even in very small population geographies like Federal American Indian Reservations. The data permit assessment of the largest OMB-designated race and ethnicity group in each geography – the classification used by the Department of Justice for Voting Rights Act scrutiny – with a precision of 99.5% confidence in variations of +/- 5 percentage points for off-spine geographies as small as 500 persons, approximately the minimum voting district size in the redistricting plans that the Department of Justice provided as examples.

56. **The accuracy of differential privacy increases at higher levels of geography, even for arbitrary geographic areas like Congressional and legislative districts.** The Census Bureau designed its implementation of differential privacy to increase accuracy

this natural variation is about +/- 120 persons (0.3% of population), also based on coverage data from the 2010 Census. As with all simulation estimates, there is sensitivity to the assumptions. The reported statistics are the mean absolute error and the mean absolute percentage error. Differentially private statistics include both the housing unit and group quarters populations. Simulations exclude the group quarters population because there are no coverage estimates for that group.

⁵⁶ The simulation of the natural variation of census block-level populations is +/- 1.5 persons, which excludes the group quarters population because there are no coverage estimates for that group. As with all simulation estimates, there is sensitivity to the assumptions. The reported statistics are the mean absolute errors. Mean absolute percentage errors are not useful statistics for block populations because more than 2,000,000 blocks with positive housing units have populations between 0 and 9. Differentially private statistics include both the housing unit and group quarters populations. Simulations exclude the group quarters population because there are no coverage estimates for that group.

as blocks are aggregated into larger geographic areas like neighborhoods, voting districts, towns, and other places. Rather than infusing noise at the block level and aggregating upwards, which would cause error to compound at larger geographic levels, the Disclosure Avoidance System's TopDown Algorithm (TDA) takes the opposite approach. Starting at the national level, the algorithm establishes very precise (but still privacy-protected) tabulations for all characteristics at the national level, then works its way down the geographic hierarchy, ensuring that all of the geographic entities at each level (e.g., the Census tracts within a county) add up precisely to the established characteristics of the level above (e.g., the county). This approach limits the distortions that can arise from noise infusion and ensures the reliability of statistics as the underlying size of the population increases. Plaintiffs argue that "the November 2020 demonstration data also skewed the 2010 tabulations enough to create a population deviation in Alabama's Congressional districts on a level that courts have found in other contexts to violate voters' equal population rights," with districts losing up to 73 individuals or gaining 206 individuals over reported values. While this may have been true for the November 2020 Demonstration Data Product, this is not true for the Demonstration Data Product that will be produced by the end of April. In the April 2021 Demonstration Data Product, Congressional districts as drawn in 2010 have a mean absolute percentage error of 0.06%. If the Congressional districts had been drawn using the April 2021 Demonstration Data Product, their statistical composition for the purposes of Voting Rights Act scrutiny would not be affected. Even for state legislative districts, which had average sizes of 159,000 (upper chambers) and 64,000 (lower chamber), the mean absolute percentage errors are 0.09% (upper chambers) and 0.16% (lower chambers), respectively. Such errors are trivial and imply that the difference between districts drawn from the April 2021 Demonstration Data Product and those drawn from the original 2010 P.L. 94-171 Redistricting Data Summary File would be statistically and practically imperceptible. *Most importantly*

for the redistricting use case, the TDA, when properly tuned, ensures that redistricters can remain confident in the accuracy of the population counts and demographic characteristics of the voting districts they draw, despite the noise in the individual building blocks.

IMPLEMENTING DIFFERENTIAL PRIVACY FOR THE 2020 CENSUS

57. Census announced that it planned to use Differential Privacy for the 2020 Census in a few different venues: (1) August 3, 2018, 2020 Census Program Management Review; (2) December 6, [2018, Census Scientific Advisory Committee Meeting](#); and (3) [May 2, 2019, Census National Advisory Committee meeting](#).
58. The Bureau has engaged in a years-long campaign to educate the user community and solicit their views about how differential privacy should be implemented. Census Bureau staff have made hundreds of public presentations, held dozens of webinars, held formal consultations with American Indian and Alaska Native tribal leaders, created an extensive website with plain English blog posts, and conducted regular outreach with dozens of stakeholder groups. We have made presentations to our scientific advisory committees and provided substantial information to oversight entities such as the Government Accountability Office and the Office of the Inspector General.
59. Part of the Bureau's effort to inform the public and solicit feedback involved releasing a series of Demonstration Data Products. There are many different ways to implement differentially private disclosure avoidance mechanisms, and the design and parameters of these mechanisms can substantially impact the fitness-for-use of the resulting data. The Census Bureau's TopDown Algorithm (TDA) was specifically designed to address the reconstruction-abetted re-identification vulnerability risks, while allowing the Bureau to tune the accuracy of the statistics to ensure fitness-for-use.

60. To date, the Census Bureau has released four sets of Demonstration Data Products (in October 2019, May 2020, September 2020, and November 2020). The Census Bureau has received substantial, actionable feedback after each release that has contributed to the system’s design and optimization.
61. All four of these demonstration products used a lower privacy-loss budget than we anticipate using for the final 2020 Census data – that is, these demonstration data were purposefully “tuned” to privacy and not “tuned” for producing highly accurate re-districting data. We held the privacy-loss budget roughly the same across these four releases to allow us to compare effects of incremental improvements in the system. After each release, these demonstration files enabled data users to help the Census Bureau identify areas where the algorithm needed to be tuned to meet their specific use cases. While the Census Bureau has not yet set the final privacy-loss budget, we have been clear that all the demonstration data released to date have used a lower privacy-loss budget (more privacy, less accuracy) than will be selected for the final production run of the redistricting data.⁵⁷
62. This degree of transparency into the design and implementation of a disclosure avoidance methodology is unprecedented in the federal government. The Census Bureau has submitted its differential privacy mechanisms, programming code, and system architecture to thorough outside peer review. We have also committed to publicly releasing the entire production code base and full suite of implementation settings and parameters. Many traditional disclosure avoidance methods, most notably swapping techniques, must be implemented in a “black box.” Implementation parameters for these legacy disclosure avoidance methods, especially swapping rates, are often

⁵⁷ Most recently on February 23, 2021 in [The Road Ahead: Upcoming Disclosure Avoidance System Milestones \(govdelivery.com\)](https://www.govdelivery.com).

some of the most tightly guarded secrets that the Census Bureau protects. But differential privacy does not rely on the obfuscation of its implementation as a means of protecting the data. The Census Bureau's transparency will allow any interested party to review exactly how the algorithm was applied to the 2020 Census data, and to independently verify that there was no improper or partisan manipulation of the data.

INVARIANTS ARE NOT REQUIRED FOR ACCURACY.

63. Invariants – or data held constant when applying differential privacy – introduce privacy risks and are not necessary to ensure accuracy. Invariants were not well understood either theoretically or empirically in 2016 when the Census Bureau began its research on differential privacy for decennial census data, but we now understand that invariants defeat the privacy protections and must be limited in order to protect the integrity of the system as a whole. Unlike traditional approaches to disclosure avoidance, differentially private noise infusion offers quantifiable and provable privacy guarantees. These guarantees, reflected in the global privacy-loss budget and its allocation to each statistic, serve as a promise to data subjects that there is an inviolable upper bound to the risk that an attacker can learn or infer something about those data subjects through publicly released data products. While that upper bound is ultimately a policy decision, and may be low or high depending on the balancing of the countervailing obligations to produce accurate data and to protect respondent confidentiality, the level of the global privacy-loss budget is central to the ability of the approach to protect the data. Invariants are, by their very nature, the equivalent of assigning infinite privacy-loss budget to particular statistics, which fundamentally violates the central promise of differentially private solutions to controlling disclosure risk. By excluding the accuracy of invariant data elements from the control of the privacy-loss budget, invariants exclude the disclosure risk and potential inferences that can be drawn from those data elements from the formal privacy guarantees. Thus,

instead of being able to promise data subjects that the publication of data products will limit an attacker to being able to infer, at most, a certain amount about them (with that amount being determined by the size of the privacy-loss budget and its allocation to each characteristic), the inclusion of one or more invariants fundamentally excludes attacker inferences about the invariant characteristic(s) from the very nature of that promise. The qualifications and exclusions to the privacy guarantee weaken the strength of the approach and make communicating the resulting level of protection substantially more difficult. This is the reason that DSEP removed the block-level invariant on population and voting-age population. Below the state level, DSEP only authorized block-level invariants that were necessary to conduct the field operations of the 2020 Census: housing unit address counts, and occupied group quarters address counts and types. As noted above, if the block population is reported with some random fluctuation around the confidential value, then only by chance will the block identifier be correct in any potential reconstructed microdata. Compound this effect over 8,000,000 blocks and the number of feasible reconstructions explodes exponentially. This is what provides the protection against re-identification from the reconstructed data.

64. Invariants are not required to improve the accuracy of any statistic processed by differential privacy. Assigning sufficiently high (but not infinite) privacy-loss budget to any statistic can ensure perfect accuracy for that statistic while still allowing the resulting privacy-loss to be communicated in the privacy guarantee. For example, the state-level population of the American Indian and Alaska Native tribal areas has been given sufficient privacy-loss budget to ensure that those populations are presented accurate to the number of persons in the units column; the mean absolute error is 1 person, essentially invariant and the same precision as the state populations themselves. But this solution still requires balancing accuracy and privacy-loss overall. All characteristics cannot have large privacy-loss budget allocations at every geographic

level. If they did, the published tables would be exact images of the confidential data and subject to the same vulnerability as the 2010 Census.

65. The forthcoming April 2021 Demonstration Data Product illustrates this tradeoff. These new demonstration data use a global privacy-loss budget for persons of 10.3, which is much larger than the 4.0 budget used in the earlier releases but is still allocated in a manner that provides a level of protection for every census record and every published characteristic. The April 2021 demonstration data also fully satisfy a tightly specified set of accuracy criteria specialized to the redistricting use case. Specifically, populations, voting-age populations, and the proportion of the largest OMB-designated race and ethnicity groups are all reliable for redistricting and Voting Rights Act scrutiny in arbitrary contiguous block aggregates for both on-spine and off-spine political and legal entities. Because new districts cannot be drawn before the 2020 P.L. 94-171 Redistricting Data Summary File is released, counties, block groups, minor civil divisions, incorporated places, and Census-designated places were all used as on- and off-spine geographic entities for tuning purposes.
66. In the April 2021 Demonstration Data Product, all the targeted small population statistics for race and ethnic groups are far more accurate than in previous demonstration data products, even though no additional invariants were used. The gain in accuracy is entirely due to dedicating more of the privacy-loss budget to the block- and block group-level statistical tables and carefully specifying the differentially private measurements to target the OMB-designated race and ethnicity groups. Biases in the tribal areas' race and ethnicity data were also greatly reduced.
67. The Census Bureau has received substantial feedback from our data user community highlighting distortions that were present in the early versions of our demonstration data, particularly in the version released in October 2019. Based on that feedback, the Census Bureau has identified and corrected the algorithmic sources of those distortions. As these measures of accuracy and bias show, any residual impact of the types

of systematic bias observed in the early demonstration data will be negligible and well within the normal variance and total error typical for a census.

PROCESS AND TIMELINE MOVING FORWARD

68. The operational delays caused by the global COVID-19 pandemic, and the resulting processing schedule changes for production of the redistricting data product shifted the milestone dates for all the systems necessary to produce the data. While the 2020 Census Disclosure Avoidance System is fully operational, and has already passed the Test Readiness Review (TRR) and Production Readiness Review (PRR) milestones on schedule, we have taken advantage of the additional time before the May 20, 2021 Operational Readiness Review (ORR) to perform additional optimization and testing of the system, and to engage in another round of data user evaluation and feedback.
69. The Census Bureau will release another demonstration product by April 30, 2021 using a higher privacy-loss budget (more accuracy) that better approximates the final privacy-loss budget that will likely be selected for the redistricting data product. These new demonstration data will also reflect system design changes that have been made since the last demonstration data release, along with tuning and optimization of the system that have been done specifically to prioritize population count accuracy and the ability to identify majority-minority districts.⁵⁸ The new release will give users yet another opportunity to let the Census know specifically where the data are (or are not yet) sufficiently accurate to meet their requirements.
70. On March 25, 2021, DSEP approved the privacy-loss budget to be used for the next demonstration product. This privacy-loss budget reflects empirical analysis of over

⁵⁸ Users will be able to see the difference between algorithmic improvements and greater privacy-loss budget. At the same time as the main April 2021 Demonstration Data Product is released, the Census Bureau will also release demonstration data using exactly the same software implementation but setting the global privacy-loss budget to 4.0 for persons, as it was in the four previous demonstration data products.

600 full-scale runs of the Disclosure Avoidance System using 2010 Census data. The Census evaluated these experimental runs using accuracy and fitness-for-use criteria for the redistricting use case informed by the extensive feedback we have received from the redistricting community and the Civil Rights Division at the U.S. Department of Justice. Based on this feedback, the privacy-loss budget for the final demonstration product is set to ensure the accuracy of racial demographics for voting districts as small as 500 individuals. With this tuning, the proportion of the largest racial group within even those small state/local voting districts of 500 individuals will be accurate to within five percentage points of the enumerated value at least 95% of the time. As voting district population size increases to any sort of reasonably anticipated legislative district, the error will be miniscule. For example, Congressional and state legislature districts will have significantly higher accuracy for population counts and voting age population counts.

71. Following the release of the new demonstration data, data users and stakeholders will have about a month to submit additional feedback on their analysis and assessment of these data, before DSEP, in early June 2021, sets the privacy-loss budget and system parameters for the production run of the redistricting data product.
72. The production run for creating the Microdata Detail File (the internal name for the file that contains the privacy-protected data) is scheduled to occur between June 26 and July 18, 2021. This roughly three-week period is similar to the period required to implement disclosure avoidance in prior censuses and is not the cause of the delay in the delivery of the redistricting data.
73. As discussed in more detail below, any court-ordered change in the Census Bureau's implementation of disclosure avoidance would add significant time to this schedule.

BRYAN AND BARBER DECLARATIONS

74. Although I cannot set out all my observations and disagreements with the declarations of Dr. Michael Barber and Mr. Thomas Bryan in this declaration, I want to identify some key areas of dispute.
75. Dr. Barber's expert report does not adequately account for the fact that the Census Bureau's demonstration data products had a privacy-loss budget significantly lower than the expected budget that will be set for the 2020 Census. As I explained above, we purposefully set the budget lower than ones most likely to be finally chosen (set to favor privacy over accuracy), so that we could isolate the distortions and demonstrate the effectiveness of various methodological modifications. One cannot draw conclusions about the accuracy of the data the Census Bureau will release for the 2020 Census based on these demonstration products.
76. Dr. Barber is premature in drawing conclusions about the accuracy of the 2020 redistricting data before the Census Bureau has set a final privacy-loss budget, and he is further incorrect in opining on the accuracy of differential privacy without considering the relative error of alternatives. Dr. Barber focuses most of his report on the possible quality concerns of differentially private 2020 Census data releases with no attention to (1) the demonstrated privacy risks of a 2020 Census protected by legacy methods and (2) the accuracy of alternatives to differential privacy including enhanced swapping or suppression. As I show in this declaration, all disclosure avoidance systems trade-off accuracy for confidentiality protection. They must be compared to each other. Releasing the redistricting data without disclosure avoidance procedures – tabulating the Census Edited File directly – is not an option and was not done for the 1990, 2000, or 2010 Censuses.
77. Dr. Barber relies on external studies that draw incorrect conclusions and use early demonstration data products. In his declaration, Dr. Barber quotes Santos-Lozada, et al. (2020) on page 14 by saying that “[i]nfusing noise in the data, in comparison to the

current disclosure avoidance system, will produce inaccurate patterns of demographic change with higher levels of error found in the calculations for non-Hispanic blacks and Hispanics. At the same time, these counts are bound to impact post-2020 districting for both federal and state elections, as well as evaluations of that redistricting. . . . [T]hese changes in population counts will affect understandings of health disparities in the nation, leading to overestimates of population-level health metrics of minority populations in smaller areas and underestimates of mortality levels in more populated ones.” The Santos-Lozada et al. paper uses the October 2019 Demonstration Data Product. Therefore, its conclusions are only applicable to the state of the algorithms and the overall privacy-loss budget used for that early release. Those were neither the final algorithms nor the final privacy-loss budget. I informed the editors of the Proceedings of the National Academy of Sciences of these defects during the peer-review process. I strongly recommended that the word “will” in the title be changed to “may” for these reasons. There is nothing statistically incorrect in the paper except for the general failure of these demographers to account for estimation error due to disclosure avoidance when doing their statistical analyses as I have noted in my own scholarly work⁵⁹ and other statisticians and computer scientists have also noted.⁶⁰ The fatal error in the Santos-Lozada et al. paper is drawing conclusions from preliminary data generated by an obsolete version of the 2020 Census DAS using obsolete settings for the privacy-loss budget and its allocation. Those conclusions are wrong and so, by extension, are those of Dr. Barber.

⁵⁹ Abowd, John M. and Ian Schmutte “Economic Analysis and Statistical Disclosure Limitation” *Brookings Panel on Economic Activity* (Spring 2015): 221-267. [[download article and discussion](#), open access] [[download preprint](#)].

⁶⁰ Wasserman L. and S. Zhou “A Statistical Framework for Differential Privacy,” *Journal of the American Statistical Association*, Vol. 105, No. 489 (2010):375-389, DOI: [10.1198/jasa.2009.tm08651](https://doi.org/10.1198/jasa.2009.tm08651).

78. Dr. Barber's conclusions do not take into account that if the Census Bureau were forced to hold the number of people in housing units invariant at the block level, that would, in turn, require adding more noise and error to the demographic characteristics of those individuals in an effort to offset what amounts to assigning block-level populations an infinite privacy-loss budget. As I show in my declaration, doing so is unnecessary and harmful to both accuracy and confidentiality protection. The correct procedure is to set accuracy targets for meaningful aggregations then tune the disclosure avoidance procedures to meet them. This procedure is transparent when using differential privacy, but it was also done for the 2010 swapping system albeit in memos that are also protected by 13 U.S. Code §§ 8(b) & 9.
79. Furthermore, Dr. Barber's work draws incorrect conclusions about biases in rural areas and for specific small populations. In his declaration, Dr. Barber states on page 13 that "[p]laces with fewer people (rural locations) and areas with smaller, distinctive populations (minority communities) are more likely to be impacted since these are the places where identification is more concerning, and the application of statistical noise is more likely to have a larger impact on the summary statistics derived from the altered data." He concludes on pages 13 and 14 that "...the process of differential privacy is not applied equally across the entire population. Places with fewer people (rural locations) and areas with smaller, distinctive populations (minority communities) are more likely to be impacted since these are the places where identification is more concerning, and the application of statistical noise is more likely to have a larger impact on the summary statistics derived from the altered data." This conclusion is incorrect. His analysis should say that the privacy-loss of the respondents in these small areas is being treated equally and identically to the privacy-loss of the respondents in large population areas; that is, every single respondent gets the full privacy protection afforded by the DAS—unlike the 2010 system, which only tried to protect certain households. To properly compare urban/rural statistics before and after the

application of disclosure avoidance, regardless of the system, the full algorithm assigning rural/urban status must be used on both the privacy-protected and confidential data. Dr. Barber has not done this.

80. Dr. Barber's work makes incorrect assertions about the non-negativity constraint. In his declaration, Dr. Barber cites Riper, Kugler, and Ruggles (2020) on page 13 stating that "[t]he non-negativity constraint requires that every cell in the final detailed histogram be non-negative. As described above, many of the cells in the noisy household histograms will be negative, especially for geographic units with smaller numbers of households. Returning these cells to zero effectively adds households to these small places, resulting in positive bias." This point is not an accurate description of how non-negativity is being handled in the post-processing of the noisy histogram. The analysis should say that negative values are not simply being returned to zero, but that all blocks with housing units are used to estimate the population counts subject to a non-negativity constraint on the solutions. That is, negative values are not "[r]eturning to zero," the entire 2,016 element matrix (for the redistricting data) is smoothed to a consistent, non-negative matrix for each of the 8,000,000 blocks, 275,000 block groups, 75,000 tracts, 3,143 counties, 51 states (including DC), and the U.S. simultaneously.⁶¹ At the block-level, there are expected to be an average of only 40 people represented across the 2,016 cells. This is the inherent sparsity that any disclosure avoidance system must address. Dr. Barber claims on page 13 that "[t]he combination of the non-negativity constraint and population invariants consistently leads to bias increasing counts of small subgroups and small geographic units and decreasing counts of larger subgroups and geographic units." While the statement is correct in

⁶¹ The matrix is 2,016 elements rather than 252 because there are eight elements in the Group Quarters Table P5 (seven group quarter types and "not a group quarters") that also interact with the other categories. The number of geographic entities at each level is based on approximate values for 2020 tabulation geographies.

principle, the magnitudes shown in his report are not representative of the final re-districting data product. At the levels of privacy-loss budget used for the forthcoming April 2021 Demonstration Data Product, the consequences of the non-negativity constraint were tightly controlled for population areas of at least 500 total persons. The remaining variation in block-level statistics, including small biases, is required to protect locational privacy and deliver consistent data. It is well within the inherent variability of block-level census data, as shown in my declaration.

81. Dr. Barber argues that the amount of error observed in the demonstration files indicates that differential privacy cannot produce data sufficient for important use cases. Mr. Barber's focus on the percentage of blocks in the demonstration data that differ at all from the official 2010 Census data (even if that difference represents the addition or subtraction of a single individual from the block) ignores two important points. First, the entire objective of our implementation of differential privacy is to infuse sufficient noise in block-level data to protect against reconstruction-abetted re-identification attacks while ensuring that when those blocks are aggregated into larger geographies of interest (voting districts, towns, etc.) those relative errors diminish and the accuracy of the tabulations improves. Second, the overall accuracy of the data is a direct consequence of the global privacy-loss budget selected and how it is allocated. The demonstration data used by both Dr. Barber and Mr. Bryan for their analyses, which use a substantially lower privacy-loss budgets than will be used for the final 2020 Census data products, can therefore be expected to be substantially "noisier" than the final data will be. Examples of noise levels in the April 2021 Demonstration Data Product provided in my report and verifiable when those data are released later this month confirm my claims.

82. Mr. Bryan assesses the accuracy of the four Demonstration Data Products (October 2019, May 2020, September 2020 and November 2020) using the percent of blocks with any change at all (pp. 9-13) or percentage errors (pp. 16-19). Both sets of analyses are

based on obsolete versions of the DAS, but they also make serious errors that will still be salient when he uses the April 2021 Demonstration Data Product. The DAS was designed to control the error in counts, not percentages. The basic tables in the P.L. 94-171 Redistricting Data Summary File are counts of resident persons living in specific geographies who have features chosen from the following taxonomy {any age, voting age}, {Hispanic/Latino, not Hispanic/Latino}, and any combination of {Afro-American/Black, American Indian/Alaska Native, Asian, Native Hawai'ian/Pacific Islander, White, Some other race} except "none." The specific aggregate geographies available in the data product are all built from census blocks, but it is the counts of persons in those aggregate geographies, including voting districts, not the block counts themselves that must be accurate enough to be fit for redistricting. Block-level errors, whether in counts or percentages, are irrelevant except to the extent that they are not controlled in larger-population geographies. In 2010, the average population in a block was 28 and the average population in an occupied block was 49. Any block-level variation in one of the 2,016 cells of the redistricting data for total populations this small is going to appear as a "large" percentage error. Indeed, most of those statistics have a base of zero, making percentage variation undefined and meaningless. The DAS must introduce noise into the block-level data to achieve any confidentiality protection at all. This statement is also true for the systems that were used in the 1970 to 2010 Census. The noise from suppression (1970, 1980) is counts that are simply not reported at the block level. The noise from blank and impute (1990) is due to the imputation modeling. The noise from swapping (2000, 2010) is due the exchange of geographic identifiers across blocks. All confidentiality protection applied to block-level redistricting data produces errors of the sort described by Mr. Bryan. Furthermore, many of the supposed DAS errors in Mr. Bryan's analysis cancel out when blocks are aggregated into larger-population geographies like block groups, census tracts, towns, counties, and congressional districts. This is not an accident; it is a carefully

designed feature of the DAS. The tabulation of the protected microdata might miss a person in one block, but have an “excess” person in the neighboring block for a particular characteristic. Because the DAS uses direct measurements from the U.S. all the way down to the block to estimate the counts at every level of geography, whether on- or off-spine, they are all much more accurate than any of the block estimates that comprise them. This is easy to see in any balanced summary of the accuracy of the DAS. Counties and places have far smaller percentage errors than the average percentage error of the blocks that compose them.

CLARIFYING STATEMENT QUOTED IN COMPLAINT

83. Plaintiffs assert, quoting an article in 2018 by the demographer Steven Ruggles and others, that I claimed that database reconstruction does not pose a significant re-identification threat. I made the statement that plaintiffs reference indirectly at the December 14, 2018 meeting of the Federal Economic Statistics Advisory Committee (FESAC) in my own presentation.⁶² Dr. Ruggles was on the FESAC program in the same session. I made the remarks in December 2018 as a report on ongoing research.⁶³ At the February 16, 2019 session of the American Association for the Advancement of Science (AAAS), I retracted my tentative conclusion about re-identification based on additional research reported there. The full text and presentation of the AAAS session are attached as Appendices H and I.⁶⁴ To be clear, the Census Bureau’s simulated recon-

⁶² Federal Economic Statistics Advisory Committee program: [FESAC Meeting Agenda December 2018 \(bea.gov\)](https://www.bea.gov/fesac-meeting-agenda-december-2018).

⁶³ My remarks at the December 18, 2018 FESAC: [Microsoft PowerPoint - Abowd Presentation \(bea.gov\)](https://www.bea.gov/microsoft-powerpoint-abowd-presentation).

⁶⁴ AAAS materials for the February 16, 2019 session area also here: <https://blogs.cornell.edu/abowd/files/2019/04/2019-02-16-Abowd-AAAS-Talk-Saturday-330-500->

struction attack on the 2010 Census data described in this declaration and in the accompanying appendix materials shows there is a significant re-identification risk. However, the Census Bureau’s Data Stewardship Executive Policy Committee (DSEP) acted to adopt differential privacy as soon as that research showed that an accurate microdata reconstruction was feasible. It did not require, nor should it have required, the subsequent demonstration that those reconstructed microdata permit between 52 and 179 million correct re-identifications from the 2010 Census. The reconstructed microdata fail the *2010 Census* microdata disclosure avoidance requirements—the requirements that were in place for that census—because they contain geographic identifiers (the block code) that relate to a minimum population of one rather than the 100,000 person minimum population that contemporary standards required. The reconstructed microdata also did not impose any of the minimum population thresholds required of the tabulation variables, especially age.⁶⁵ These requirements were already in place because it is well understood at the Census Bureau and in the official statistics community worldwide that geographic identifiers for low-population areas, sex, and exact age in microdata files are a major disclosure risk especially in population censuses.

IMPACT OF ANY COURT RULING BARRING USE OF DIFFERENTIAL PRIVACY

84. Were the Court to rule that the Census Bureau was precluded from using differential privacy for the 2020 Census P.L. 94-171 Redistricting Data Summary File, we would be faced with hard choices. The inevitable result would be significant delay in deliv-

[session-FINAL-as-delivered-2jr4lzb.pdf](#) and <https://blogs.cornell.edu/abowd/files/2019/04/2019-02-16-Abowd-AAAS-Slides-Saturday-330-500-session-FINAL-as-delivered-1iqsdg2.pdf>.

⁶⁵ McKenna (2019a).

ery of the already-delayed redistricting data and diminished accuracy. Either the Census Bureau would have to revert to using suppression (as was last used in the 1980 Census) or use enhanced swapping (as was used in the 1990 to 2010 Censuses, but at a much higher rate and with fewer invariants). Either choice would delay results and diminish accuracy.

85. The effect on the schedule for delivering redistricting data would be substantial. The Census Bureau cannot ascertain the length of the delay until it understands any parameters the Court might place on its choice of methodology, but under all scenarios the delay would be multiple months. This delay is unavoidable because the Census Bureau would need to develop and test new systems and software, then use them in production and subject the results to expert subject matter review prior to production of data. The Census Bureau has been developing the systems and software to use differential privacy for several years—the agency has spent millions of dollars purchasing cloud computer capacity and writing and tuning code. The systems and software are ready to go and await only final tuning and a decision on the privacy-loss budget.

86. Even if the agency was ordered to repeat exactly what was done in 2010 (despite the serious risks to privacy the Census has identified), we could not simply “flip a switch” and revert to the prior methodology. Instead, we would need to conduct the requisite software development and testing. The 2020 Census’s system architecture is completely different than that used in the 2010 Census, and it is thus not possible to simply “plug in” the disclosure-avoidance system used in 2010.

87. Not only would redistricting data be further delayed, but the resulting data would be less accurate. Both swapping and suppression are blunt instruments for privacy protection. Unlike differential privacy, neither can be effectively tuned to optimize for data accuracy. Knowing that the 2010 Census results were vulnerable to reconstruction, the Census Bureau cannot simply repeat the swapping protocols from the 2010

census, but rather would be forced to fashion appropriate levels of protection for either system. Using an appropriate level of protection for either suppression or swapping would produce far less accurate data than would differential privacy.

88. I would urge any court to be quite wary of opining on the suitability of particular methods for conducting disclosure avoidance, as these decisions are highly technical and can have unanticipated consequences. The only reason the Court knows so much about the proposed methods for the 2020 Census is that transparency does not undermine their confidentiality protections, which is not the case for either swapping or suppression. While we cannot predict the full impact of any change, there is a danger than any change would have cascading effects on data accuracy and privacy, making race and ethnicity data, along with age data, substantially less accurate. Any sort of change in the basic methodology would be minimally tested and would not have the benefit of any input from the user community.

89. In conclusion, it is my professional opinion that the Census Bureau's Data Stewardship Executive Policy Committee should be permitted to control the type and parameters of any disclosure avoidance system used for the 2020 Census, just as it did for the 2010 Census and just as its predecessor committees did for decennial censuses conducted since the passage of the Census Act (13 U.S. Code) in 1954.

I declare under penalty of perjury that the foregoing is true and correct.

DATED and SIGNED:

JOHN ABOWD

Digitally signed by JOHN ABOWD
Date: 2021.04.13 08:45:14 -04'00'

John M. Abowd

Chief Scientist and Associate Director for Research and Methodology

United States Bureau of the Census

John M. Abowd

- Home
- Professional Information
- Courses
- Recent News
- Special Materials

Professional Information

[Updated April 1, 2021]

CONTACT INFORMATION

U.S. Census Bureau
 HQ 8H120 ATTN: Sara Sullivan
 4600 Silver Hill Road
Private delivery services (FedEx, UPS, etc.) physical location: Suitland, MD 20746
USPS mail only: Washington, DC 20233
 Voice: +1.301.763.5880
 Mobile: +1.202.591.0766
 Fax: +1.301.763.8360
 Executive assistant Sara Sullivan: +1.301.763.5116
 E-mail: john.maron.abowd@census.gov

ILR School
USPS mail only (send private delivery service items to the address above):
 275 Ives Hall
 Cornell University
 Ithaca, New York 14853-3901
 Assistant: LDI@cornell.edu
 E-mail: john.abowd@cornell.edu

Webpage: <https://blogs.cornell.edu/abowd/> or <https://www.johnabowd.com>

Twitter: @john_abowd (opinions are my own)

Short biography in PDF format

CURRENT POSITIONS

Chief Scientist and Associate Director for Research and Methodology, U. S. Census Bureau, IPA June 1, 2016 – March 27, 2020; Career Senior Executive Service March 29, 2020 –

Edmund Ezra Day Professor, Department of Economics, Cornell University, July 2011 – currently on leave

Director, Labor Dynamics Institute, Cornell University, October 2011 – currently on leave

Founding member and Professor of Information Science (by courtesy), Faculty of Computing and Information Science, July 2000 – currently on leave

Professor of Statistics and Data Science, September 2013 – currently on leave

Member of the Graduate Fields of Economics, Industrial and Labor Relations, Information Science, and Statistics

Search ...

SEARCH

INSTITUTIONS

- U.S. Census Bureau
- Cornell Economics
- Labor Dynamics Institute
- NCRN node at Cornell
- CISER

OTHER INFORMATION

- Google Scholar
- ORCID
- RePEC/Ideas
- SSRN

Research Associate, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, Massachusetts 02138, September 1983 – (on leave while serving at the U.S. Census Bureau)

Research Affiliate, Centre de Recherche en Economie et Statistique/INSEE, 15, bd Gabriel Péri, 92245 Malakoff Cedex France, November 1997 –

Research Fellow, IZA (Institute for the Study of Labor), P.O. Box 7240 D-53072 Bonn, Germany, June 2002 –

Research Fellow, IAB (Institut für Arbeitsmarkt-und Berufsforschung), Dienstgebäude Weddigenstraße 20-22, 90478 Nürnberg, Germany, January 2013 –

President and Principal, ACES-Research, LLC, john@aces-research.com, July 2007 –

Editor, Journal of Privacy and Confidentiality Online journal

PREVIOUS AND VISITING POSITIONS

Distinguished Senior Research Fellow, United States Census Bureau, September 1998 – May 2016

Associate Chair, Department of Economics, Cornell University, August 2015 – May 2016

Visiting Professor, Center for Labor Economics, University of California-Berkeley, August 2014 – July 2015

Director of Graduate Studies, Economics, July 2010 – June 2014

Professor of Economics and Econometrics, University of Notre Dame, January 2008 – May 2008.

Director, Cornell Institute for Social and Economic Research (CISER), July 1999 – December 2007

Associate Director, Cornell Theory Center (became Cornell University Center for Advanced Computing), September, 2006 – August 2007.

Professor of Labor Economics, Cornell University, January 1990 – October 2001.

Edmund Ezra Day Professor, School of Industrial and Labor Relations, November 2001 –

Associate Director, Cornell Institute for Social and Economic Research (CISER), July 1998 – June 1999.

Chair, Department of Labor Economics, Cornell University, September 1992 – June 1998.

Acting Director, CISER, January 1998-June 1998.

Professeur invité, Laboratoire de Microéconomie Appliquée-Theorie Et Applications en Microéconomie et macroéconomie (LAMIA-TEAM), Université de Paris-I (Panthéon-Sorbonne), May 1998.

Consultant, Centre de Recherche en Economie et Statistique (CREST), Institut National de la Statistique et des Etudes Economiques (INSEE), February 1997.

Professeur invité, ERMES (Equipe de Recherche sur les Marchés, l'Emploi et la Simulation) Université Panthéon-Assas (Paris II), October 1995 – July 1996 (part time).

Professor, Samuel Curtis Johnson Graduate School of Management, Cornell University (adjunct appointment), August 1987 – July 1995.

Chercheur étranger, Institut National de la Statistique et des Etudes Economiques (INSEE), Paris, Department of Research, August 1991 – July 1992, January 1993, January 1994.

Professeur visitant, HEC (Hautes Etudes Commerciales, Paris) Department of Finance and Economics, September 1991 – July 1992 and January 1993, December 1993 – January 1994.

Professeur visitant, CREST (Centre de Recherche en Statistique et Economie, Paris), September 1991 – July 1992, July 1993.

Associate Professor with tenure, Cornell University, August 1987 – December 1989.

Research Associate, Industrial Relations Section, Department of Economics, Princeton University, September 1986 – August 1987.

Visiting Associate Professor of Economics, Department of Economics, Massachusetts Institute of Technology, September 1985 – August 1986.

Associate Professor of Econometrics and Industrial Relations, Graduate School of Business, University of Chicago, September 1982 – August 1986. Assistant Professor, September 1979 – August 1982. Visiting Assistant Professor, September 1978 – August 1979.

Senior Study Director/Research Associate, NORC/Economics Research Center, 6030 Ellis Avenue, Chicago, Illinois 60637, September 1978 – August 1986.

Academic Consultant, Centre for Labour Economics, London School of Economics, January 1979 – April 1979.

Assistant Professor of Economics, Department of Economics, Princeton University, September 1977 – August 1979 (on leave September 1978 – August 1979). Lecturer in Economics, September 1976 – August 1977.

Associate Editor, *Journal of Business and Economic Statistics*, 1983 – 1989.

Editorial Board, *Journal of Applied Econometrics*, 1987 – 1989.

Associate Editor, *Journal of Econometrics*, 1987 – 1989.

EDUCATION

Ph.D. Department of Economics, University of Chicago, December 1977.
Thesis: An Econometric Model of the U.S. Market for Higher Education

M.A. Department of Economics, University of Chicago, March 1976.

A.B. Department of Economics (with highest honors), University of Notre Dame, May 1973.

LANGUAGES

English (native), French

HONORS AND FELLOWSHIPS

Fellow, American Association for the Advancement of Science (elected October 2020)

Julius Shiskin Award, American Statistical Association, Business and Economic Statistics Section (2016)

Cornell University, Graduate and Professional Student Assembly Award for Excellence in Teaching, Advising, and Mentoring (May 2015)

Fellow, *Econometric Society* (elected November 2014)

Roger Herriot Award, American Statistical Association, Government and Social Statistics Sections (2014)

Elected member, *International Statistical Institute* (March 2012)

Council of Sections (2014-2016), Chair (2013) *Business and Economic Statistics Section* (Chair-elect 2012), American Statistical Association

President (2014-2015), *Society of Labor Economists*, President-elect (2013-2014), Vice President (2011-2013)

Fellow, *The American Statistical Association* (elected August 2009)

Fellow, *Society of Labor Economists* (elected November 2006)

La bourse de haut niveau du Ministère de la Recherche et de la Technologie, fellowship for research at the Institut National de la Statistique et des Etudes Economiques (INSEE) awarded by the French Government, September 1991 – February 1992.

National Institute of mental Health postdoctoral fellow at NORC, September 1978 – August 1980.

National Institute of Mental Health pre-doctoral fellow at the University of Chicago, September 1973 – June 1976.

TEACHING EXPERIENCE

Graduate:

- Microeconometrics using *Linked Employer-Employee Data* (CREST-ENSAE)
- Understanding Social and Economic Data* (Cornell, co-instructor: Lars Vilhuber)
- Third-year Research Seminar I and II (Cornell)
- Seminar in Labor Economics I, II, and III (Cornell)
- Microéconomie des Données Appariées (CREST-GENES, in French)
- Microéconomie et Microéconométrie du Travail (Université de Paris I, in French)
- Economie du Travail (Université de Paris II, in French)
- Economics of Compensation and Organization (Cornell)
- International Human Resource Management (Cornell)
- Corporate Finance (Hautes Etudes Commerciales, Paris)
- International Human Resource Management (HEC, Paris)
- Workshop in Labor Economics (Cornell)
- Economics of Collective Bargaining (Cornell)
- Executive Compensation (Cornell)
- Labor Economics (MIT)
- Labor and Public Policy (MIT)
- Applied Econometrics I, II (Chicago)
- Introduction to Industrial Relations (Chicago)
- Econometric Theory I (Chicago)
- Industrial Relations and International Business (Chicago)
- Workshop in Economics and Econometrics (Chicago)
- Econometric Analysis of Time Series (Princeton)
- Mathematics for Economists (Princeton)

Undergraduate:

- Understanding Social and Economic Data* (Cornell, co-instructor: Lars Vilhuber)
- Introductory Microeconomics* (Cornell)
- Economics of Employee Benefits (Cornell)
- Economics of Wages and Employment (Cornell)

Corporate Finance (Cornell)
 Introduction to Econometrics (Princeton)
 Microeconomics (Princeton)

BIBLIOGRAPHY

Books

1. Abowd, John M. and Francis Kramarz (eds.) *The Microeconometrics of Human Resource Management*, special issue of *Annales d'économie et de statistique* 41/42 (Paris: ADRES, January/June 1996).
2. Abowd, John M. and Richard B. Freeman (eds.) *Immigration, Trade and the Labor Market* (Chicago: University of Chicago Press for the National Bureau of Economic Research, 1991).

Articles

1. McKinney, Kevin L., John M. Abowd, and John Sabelhaus, "United States Earnings Dynamics: Inequality, Mobility, and Volatility," In Raj Chetty, John N. Friedman, Janet C. Gornick, Barry Johnson, and Arthur Kennickel, eds., *Measuring the Distribution and Mobility of Income and Wealth*, (Chicago: University of Chicago Press for the National Bureau of Economic Research, 2021), forthcoming. [[download preprint](#)] [[download chapter \(open access\)](#)]
2. Abowd, John M. "Official Statistics at the Crossroads: Data Quality and Access in an Era of Heightened Privacy Risk," *The Survey Statistician*, Vol. 83 (January 2021):23-26. [[download \(open access\)](#)]
3. McKinney, Kevin L., Andrew S. Green, Lars Vilhuber and John M. Abowd "Total Error and Variability Measures for the Quarterly Workforce Indicators and LEHD Origin-Destination Employment Statistics in OnTheMap" *Journal of Survey Statistics and Methodology* (November 2020). [[download arxiv preprint](#)], DOI: <https://doi.org/10.1093/jssam/smaa029>, supplemental online materials DOI: <https://doi.org/10.5281/zenodo.3951670>
4. Abowd, John M., Ian M. Schmutte, William Sexton, and Lars Vilhuber "Why the Economics Profession Must Actively Participate in the Privacy Protection Debate," *American Economic Association: Papers and Proceedings*, Vol. 109 (May 2019): 397-402, DOI:10.1257/pandp.20191106. [[download preprint](#)]
5. Abowd, John M. and Ian M. Schmutte "An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices," *American Economic Review*, Vol. 109, No. 1 (January 2019):171-202, DOI:10.1257/aer.20170627. [AER, [ArXiv preprint](#), [Replication information](#)]
6. Weinberg, Daniel H., John M. Abowd, Robert F. Belli, Noel Cressie, David C. Folch, Scott H. Holan, Margaret C. Levenstein, Kristen M. Olson, Jerome P. Reiter, Matthew D. Shapiro, Jolene Smyth, Leen-Kiat Soh, Bruce D. Spencer, Seth E. Spielman, Lars Vilhuber, and Christopher K. Wikle "Effects of a Government-Academic Partnership: Has the NSF-Census Bureau Research Network Helped Secure the Future of the Federal Statistical System?" *Journal of Survey Statistics and Methodology* (2018) DOI:10.1093/jssam/smy023. [[download](#), [open access](#)] [[download preprint](#)]
7. Abowd, John M., Ian M. Schmutte and Lars Vilhuber "Disclosure Limitation and Confidentiality Protection in Linked Data," in A.Y. Chun, M. Larson, J. Reiter, and G. Durrant (eds.) *Administrative Records for Survey Methodology* (New York: Wiley, forthcoming). [[download preprint](#)]
8. Abowd, John M., Kevin L. McKinney and Ian M. Schmutte "Modeling Endogenous Mobility in Earnings Determination," *Journal of Business and Economic Statistics* Vol. 37, Issue 3 (2019):405-418. DOI: [10.1080/07350015.2017.1356727](https://doi.org/10.1080/07350015.2017.1356727). [[download preprint](#)] [JBES]
9. Abowd, John M., Francis Kramarz, Sébastien Perez-Duarte, and Ian Schmutte "Sorting between and within Industries: A Testable Model of Assortative Matching," *Annals of Economics and Statistics* 129 (March 2018): 1-32. NBER WP-20472. [[download preprint](#)] [[programs](#)] [[data](#)]
10. Abowd, John M., Kevin L. McKinney and Nellie Zhao "Earnings Inequality and Mobility Trends in the United States: Nationally Representative Estimates from Longitudinally Linked Employer-Employee Data," *Journal of Labor Economics* 36, S1 (January 2018):S183-S300 DOI: [10.1086/694104](https://doi.org/10.1086/694104). [[download](#), [not copyrighted](#)] [[download preprint](#)]

11. Abowd, John M. "How Will Statistical Agencies Operate When All Data Are Private?" *Journal of Privacy and Confidentiality*, Vol. 7, Issue 3, Article 1 (2017). [[download](#), [open journal](#)]
12. Haney, Samuel, Ashwin Machanavajjhala, John M. Abowd, Matthew Graham, Mark Kutzbach, and Lars Vilhuber "Utility Cost of Formal Privacy for Releasing National Employer-Employee Statistics," ACM SIGMOD 2017, DOI: 10.1145/3035918.3035940. [[download](#)]
13. Abowd, John M. and Kevin L. McKinney "Noise Infusion as a Confidentiality Protection Measure for Graph-based Statistics" *Statistical Journal of the International Association for Official Statistics* (2016) Vol. 32, No. 1, pp. 127-135, DOI: 10.3233/SJI-160958. [[download article](#), [open access](#)] [[download preprint](#)]
14. Abowd, John M. and Ian Schmutte "Economic Analysis and Statistical Disclosure Limitation" *Brookings Panel on Economic Activity*(Spring 2015): 221-267. [[download article and discussion](#), [open access](#)] [[download preprint](#)]
15. Schneider, Matthew J. and John M. Abowd "A New Method for Protecting Interrelated Time Series with Bayesian Prior Distributions and Synthetic Data," *Journal of the Royal Statistical Society, Series A* (2015) DOI:10.1111/rssa.12100. [[download preprint](#)]
16. Lagoze, Carl, William C. Block, Jeremy Williams, Lars Vilhuber, and John M. Abowd "Data Management of Confidential Data." *International Journal of Digital Curation* 8, no. 1 (2013): 265-278. doi:10.2218/ijdc.v8i1.259. [[download preprint](#)]
17. Abowd, John M. and Martha H. Stinson "Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data," *Review of Economics and Statistics*, Vol. 95, No. 5 (December 2013): 1451-1467. doi:10.1162/REST_a_00352. [[download](#), [not copyrighted](#)]
18. Abowd, John M., Matthew J. Schneider and Lars Vilhuber "Differential Privacy Applications to Bayesian and Linear Mixed Model Estimation," *Journal of Privacy and Confidentiality*: Vol. 5: Iss. 1 (2013): Article 4. [[download](#), [open access](#)]
19. Abowd, John M., Francis Kramarz, Paul Lengermann, Kevin L. McKinney, and Sébastien Roux "Persistent Inter-Industry Wage Differences: Rent Sharing and Opportunity Costs," *IZA Journal of Labor Economics*, 2012, 1:7, doi:10.1186/2193-8997-1-7. [[download](#), [open access](#)] [[online Appendix](#)]
20. Abowd, John M., Lars Vilhuber and William Block "A Proposed Solution to the Archiving and Curation of Confidential Scientific Inputs," in J. Domingo-Ferrer and I. Tinnirello, eds., *Privacy in Statistical Databases* 2012, LNCS 7556, pp. 216-225, (2012). [[download](#), [open access](#)]
21. Abowd, John M. and Lars Vilhuber "Did the Housing Price Bubble Clobber Local Labor Markets When It Burst?" *American Economic Review Papers and Proceedings* Vol. 102, No. 3 (May 2012): 589-93, doi:pdfplus/10.1257/aer.102.3.589. [[download preprint](#)] [[online Appendix](#)] [[data Readme](#)] [[data](#)]
22. Abowd, John M., R. Kaj Gittings, Kevin L. McKinney, Bryce E. Stephens, Lars Vilhuber, and Simon Woodcock "Dynamically Consistent Noise Infusion and Partially Synthetic Data As Confidentiality Protection Measures for Related Time-series," Federal Committee on Statistical Methodology, Office of Management and Budget, 2012 Research Conference Papers. [[download](#), [open access](#), cited on May 21, 2012] [[download archival copy](#)].
23. Abowd, John M. and Matthew Schneider "An Application of Differentially Private Linear Mixed Modeling," ICDMW, pp. 614-619, 2011 IEEE 11th International Conference on Data Mining Workshops, 2011. [[download](#), [open access](#)]
24. Kinney, Satkartar K., Jerome P. Reiter, Arnold P. Reznick, Javier Miranda, Ron S. Jarmin, and John M. Abowd "Towards Unrestricted Public Use Business Micro-data: The Synthetic Longitudinal Business Database," *International Statistical Review*, Vol. 79, No. 2 (December 2011):362-84, doi:10.1111/j.1751-5823.2011.00153.x. [[download](#), [subscription required](#)] [[download preprint](#)]
25. Abowd, John M. and Lars Vilhuber "National Estimates of Gross Employment and Job Flows from the Quarterly Workforce Indicators with Demographic and Industry Detail," *Journal of Econometrics*, Vol. 161 (March 2011): 82-99, doi: 10.1016/j.jeconom.2010.09.008. [[download preprint](#)] [[data](#)]
26. Abowd, John M., Bryce Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators" in T. Dunne, J.B. Jensen and M.J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data* (Chicago: University of Chicago Press for the National Bureau of Economic Research, 2009), pp. 149-230. [[download](#), [not copyrighted](#)] [[archival copy](#)]

27. Abowd, John M., Kevin McKinney and Lars Vilhuber "The Link between Human Capital, Mass Layoffs, and Firm Deaths" in T. Dunne, J.B. Jensen and M.J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data* (Chicago: University of Chicago Press for the National Bureau of Economic Research, 2009), pp. 447-472. [[download, not copyrighted](#)] [[archival copy](#)]
28. Abowd, John M. and Lars Vilhuber "How Protective are Synthetic Data," in J. Domingo-Ferrer and Y. Saygun, eds., *Privacy in Statistical Databases*, (Berlin: Springer-Verlag, 2008), pp. 239-246. [[download preprint](#)]
29. Abowd, John M., Francis Kramarz and Simon Woodcock "Econometric Analyses of Linked Employer-Employee Data," in L. Mátyás and P. Sevestre, eds., *The Econometrics of Panel Data* (The Netherlands: Springer, 2008), pp. 727-760. [[download preprint](#)]
30. Abowd, John M., John Haltiwanger and Julia Lane "Wage Structure and Labor Mobility in the United States," in E. P. Lazear and K. L. Shaw, eds., *Wage Structure, Raises, and Mobility: International Comparisons of the Structure of Wages within and Across Firms* (Chicago: University of Chicago Press for the National Bureau of Economic Research, 2008), pp. 81-100. [[download](#)] [[download preprint](#)]
31. Machanavajjhala Ashwin, Daniel Kifer, John M. Abowd, Johannes Gehrke, and Lars Vilhuber "Privacy: Theory Meets Practice on the Map," International Conference on Data Engineering (ICDE) 2008: 277-286, doi:10.1109/ICDE.2008.4497436. [[download preprint](#)]
32. Abowd, John M. and Francis Kramarz "Human Capital and Worker Productivity: Direct Evidence from Linked Employer-Employee Data," *Annales d'Economie et de Statistique*, No. 79/80, (Juillet/Décembre 2005): 323-338. [[download preprint](#)]
33. Torra, V. J.M. Abowd and J. Domingo-Ferrer "Using Mahalanobis Distance-Based Record Linkage for Disclosure Risk Assessment," in J. Domingo-Ferrer and Luisa Franconi (eds.) *Privacy in Statistical Databases* (Berlin: Springer-Verlag, 2006), pp. 233-242. [[download preprint](#)]
34. Abowd, John M., Francis Kramarz and Sébastien Roux "Wages, Mobility and Firm Performance: Advantages and Insights from Using Matched Worker-Firm Data," *Economic Journal*, Vol. 116, (June 2006): F245-F285. [[download preprint](#)]
35. Abowd, John M., Francis Kramarz and Sébastien Roux "Heterogeneity in Firms' Wages and Mobility Policies," in H. Bunzel, B.J. Christensen, G.R. Neumann and J-M. Robin, eds., *Structural Models of Wage and Employment Dynamics*, (Amsterdam: Elsevier Science, 2006), pp. 237-268. [[download preprint](#)]
36. Abowd, John M. and Lars Vilhuber "The Sensitivity of Economic Statistics to Coding Errors in Personal Identifiers," *Journal of Business and Economics Statistics*, Vol. 23, No. 2 (April 2005): 133-152, *JBES* Joint Statistical Meetings invited paper with discussion and "Rejoinder" (April 2005): 162-165. [[download preprint](#)].
37. Abowd, John M., John Haltiwanger, Ron Jarmin, Julia Lane, Paul Lengeremann, Kristin McCue, Kevin McKinney, and Kristin Sandusky "The Relation among Human Capital, Productivity and Market Value: Building Up from Micro Evidence," in *Measuring Capital in the New Economy*, C. Corrado, J. Haltiwanger, and D. Sichel (eds.), (Chicago: University of Chicago Press for the NBER, 2005), Chapter 5, pp. 153-198. [[download, not copyrighted](#)] [[download preprint](#)]
38. Abowd, John M. and Simon Woodcock "Multiply-Imputing Confidential Characteristics and File Links in Longitudinal Linked Data," in J. Domingo-Ferrer and V. Torra (eds.) *Privacy in Statistical Databases* (Berlin: Springer-Verlag, 2004), pp. 290-297. [[download preprint](#)]
39. Abowd, John M., John Haltiwanger and Julia Lane "Integrated Longitudinal Employee-Employer Data for the United States," *American Economic Review Papers and Proceedings*, Vol. 94, No. 2 (May 2004): 224-229. [[download preprint](#)]
40. Abowd, John M. and Julia Lane "New Approaches to Confidentiality Protection: Synthetic Data, Remote Access and Research Data Centers," in J. Domingo-Ferrer and V. Torra (eds.) *Privacy in Statistical Databases* (Berlin: Springer-Verlag, 2004), pp. 282-289. [[download preprint](#)]
41. Abowd, John M. and Francis Kramarz "The Costs of Hiring and Separations," *Labour Economics*, Vol. 10, Issue 5 (October 2003): 499-530. [[download preprint](#)]
42. Abowd, John M. "Unlocking the Information in Integrated Social Data," *New Zealand Economic Papers*, 0077-9954, Vol. 36, No. 1 (June 2002): 9-31. [[download preprint](#)]
43. Abowd, John M. and Orley Ashenfelter "Using Price Indices and Sale Rates to Assess Short Run Changes in the Market for Impressionist and Contemporary Paintings" in *The Economics of Art Auctions*, G. Masetto and M. Vecco (eds.), (Milan: F. Angeli Press, 2002). [[download preprint](#)] [[access book](#)]

44. Abowd, John M. and Simon Woodcock "Disclosure Limitation in Longitudinal Linked Data," in *Confidentiality, Disclosure and Data Access: Theory and Practical Applications for Statistical Agencies*, P. Doyle, J. Lane, J. Theeuwes, and L. Zayatz (eds.), (Amsterdam: North Holland, 2001), 215-277. [download preprint]
45. Abowd, John M., Bruno Crépon and Francis Kramarz "Moment Estimation with Attrition: An Application to Economic Models," *Journal of the American Statistical Association*, 96, No. 456 (December 2001): 1223-1231. [download preprint]
46. Abowd, John M., Francis Kramarz, David Margolis, and Kenneth Troske "The Relative Importance of Employer and Employee Effects on Compensation: A Comparison of France and the United States," *Journal of the Japanese and International Economies*. Vol. 15, No. 4, (December 2001): 419-436. [download preprint]
47. Abowd, John M. Julia Lane and Ronald Prevost "Design and Conceptual Issues in Realizing Analytical Enhancements through Data Linkages of Employer and Employee Data" in the *Proceedings of the Federal Committee on Statistical Methodology*, November 2000. [download preprint]
48. Abowd, John M., Francis Kramarz, David Margolis and Kenneth Troske "Politiques salariales et performances des entreprises : une comparaison France/Etats-Unis," *Economie et Statistique*, No. 332-333 (2000): 27-38. [Corporate Wage Policies and Performances: Comparing France with the United States] [download preprint]
49. Abowd, John M. and David Kaplan "Executive Compensation: Six Questions That Need Answering," *Journal of Economic Perspectives*, 13 (1999): 145-168. [Preprint and supplementary materials available at <http://hdl.handle.net/1813/56585>]
50. Abowd, John M., Patrick Corbel and Francis Kramarz "The Entry and Exit of Workers and the Growth of Employment: An Analysis of French Establishments" *Review of Economics and Statistics*, 81(2), (May 1999): 170-187. [download preprint]
51. Abowd, John M. and Francis Kramarz "Econometric Analysis of Linked Employer-Employee Data," *Labour Economics*, 6(March 1999): 53-74. [download preprint]
52. Abowd, John M., Hampton Finer and Francis Kramarz "Individual and Firm Heterogeneity in Compensation: An Analysis of Matched Longitudinal Employer-Employee Data for the State of Washington" in J. Haltiwanger *et al.* (eds.) *The Creation and Analysis of Employer-Employee Matched Data*, (Amsterdam: North Holland, 1999), pp. 3-24. [download preprint]
53. Abowd, John M. and Francis Kramarz "The Analysis of Labor Markets Using Matched Employer-Employee Data," in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, Volume 3(B), Chapter 40 (Amsterdam: North Holland, 1999), pp. 2629-2710. [download preprint]
54. Abowd, John M. Francis Kramarz and David Margolis "High Wage Workers and High Wage Firms," *Econometrica*, 67(2) (March 1999): 251-333. [download preprint]
55. Abowd, John M. Francis Kramarz, Thomas Lemieux, and David Margolis "Minimum Wages and Youth Employment in France and the United States," in D. Blanchflower and R. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries* (Chicago: University of Chicago Press, 1999), pp. 427-472. [download] [download preprint]
56. Abowd, John M. and Francis Kramarz "Internal and External Labor Markets: An Analysis of Matched Longitudinal Employer-Employee Data" in J. Haltiwanger, M. Manser, and R. Topel (eds.) *Labor Statistics and Measurement Issues* (Chicago: University of Chicago Press, 1998), pp. 357-370. [download] [download preprint]
57. Abowd, John M., Francis Kramarz, David Margolis and Kenneth Troske "The Relative Importance of Employer and Employee Effects on Compensation: A Comparison of France and the United States," in *Comparaisons internationales de salaires* (Paris: Ministère du travail et des affaires sociales and INSEE, 1996), pp. 315-327.
58. Abowd, John M. and Laurence Allain "Compensation Structure and Product Market Competition," *Annales d'économie et de statistique*, (January/June 1996, No. 41/42): 207-217. [download preprint]
59. Abowd, John M., Francis Kramarz and Antoine Moreau "Product Quality and Worker Quality," *Annales d'économie et de statistique*, (January/June 1996, No. 41/42): 300-322. [download]
60. Abowd, John M. and Francis Kramarz "The Microeconometrics of Human Resource Management: International Studies of Firm Practices, Introduction and Overview," *Annales d'économie et de statistique*, (January/June 1996, No. 41/42): 1-9 (French), 11-19 (English).
61. Abowd, John M. and Francis Kramarz "Les Politiques Salariales : Individus et Entreprises" (Compensation Policies: Individuals and Firms), *Revue Economique* 47 (May 1996): 611-622.

[\[download preprint\]](#)

62. Abowd, John M. and Francis Kramarz "The Economic Analysis of Compensation Systems: Collective and Individual" in Norman Bowes and Alex Grey, eds. *Job Creation and Loss: Analysis, Policy and Data Development* (Paris: OECD, 1996), pp. 47-54.
63. Abowd, John M. and Michael Bognanno "International Differences in Executive and Managerial Compensation" in R.B. Freeman and L. Katz, eds. *Differences and Changes in Wage Structures* (Chicago: NBER, 1995), pp. 67-103. [\[download\]](#)
64. Abowd, John M. and Thomas Lemieux "The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada," *Quarterly Journal of Economics* 108 (November 1993): 983-1014.
65. Abowd, John M. and Francis Kramarz "A Test of Negotiation and Incentive Compensation Models Using Longitudinal French Enterprise Data," in J.C. van Ours, G.A. Pfann and G. Ridder, eds. *Labour Demand and Equilibrium Wage Formation Contributions to Economic Analysis* (Amsterdam: North-Holland, 1993), pp. 111-46. [\[download preprint\]](#)
66. Abowd, John M. and Richard B. Freeman "Introduction and Summary" in J.M. Abowd and R.B. Freeman, eds. *Immigration, Trade and the Labor Market* (Chicago: NBER, 1991), pp. 1-25. [\[download\]](#)
67. Abowd, John M. and Thomas Lemieux "The Effects of International Competition on Collective Bargaining Outcomes: A Comparison of the United States and Canada," in J.M. Abowd and R.B. Freeman, eds. *Immigration, Trade and the Labor Market* (Chicago: NBER, 1991), pp. 343-67. [\[download\]](#)
68. Abowd, John M. "The NBER Trade and Immigration Data Files," in J.M. Abowd and R.B. Freeman, eds. *Immigration, Trade and the Labor Market* (Chicago: NBER, 1991), pp. 407-21. [\[download\]](#)
69. Abowd, John M. "Does Performance-based Compensation Affect Corporate Performance?" *Industrial and Labor Relations Review* 43:3 (February 1990): 525-735. Reprinted in *Do Compensation Policies Matter?* R.G. Ehrenberg, ed. (Ithaca, NY: ILR Press, 1990), pp. 52-73.
70. Abowd, John M., George Milkovich and John Hannon "The Effects of Human Resource Management Decisions on Shareholder Value," *Industrial and Labor Relations Review* 43:3 (February 1990): 203S-236S. Reprinted in *Do Compensation Policies Matter?* R.G. Ehrenberg, ed. (Ithaca, NY: ILR Press, 1990), pp. 203-236.
71. Abowd, John M. "The Effect of Wage Bargains on the Stock Market Value of the Firm," *American Economic Review* 79:4 (September 1989): 774-800. (working paper title: "Collective Bargaining and the Division of the Value of the Enterprise.")
72. Abowd, John M. and Joseph Tracy "Market Structure, Strike Activity, and Union Wage Settlements," *Industrial Relations* 57:2 (Spring 1989): 227-50.
73. Abowd, John M. and David Card "On the Covariance Structure of Earnings and Hours Changes," *Econometrica* 57:2 (March, 1989): 411-45.
74. Vroman, Wayne and John M. Abowd "Disaggregated Wage Developments," *Brookings Papers on Economic Activity* (1:1988): 313-46.
75. Abowd, John M. and David Card "Intertemporal Labor Supply and Long Term Employment Contracts," *American Economic Review* 77:1 (March 1987): 50-68.
76. Abowd, John M. "New Development in Longitudinal Data Collection for Labor Market Analysis: Collective Bargaining Data," *American Statistical Association 1985 Proceedings of the Business and Economic Statistics Section* (Washington, DC: ASA, 1985). (invited paper)
77. Abowd, John M. and Arnold Zellner "Estimating Gross Labor Force Flows," *Journal of Business and Economic Statistics* 3 (July 1985): 254-283.
78. Abowd, John M. and Arnold Zellner "Application of Adjustment Techniques to U.S. Gross Flow Data," *Gross Flows in Labor Force Statistics*, edited by Paul Flaim and Carma Hogue, Bureau of the Census/Bureau of Labor Statistics Conference Volume (Washington, DC: GPO, 1985).
79. Abowd, John M. and Mark Killingsworth "Employment, Wages, and Earnings of Hispanics in the Federal and Nonfederal Sectors: Methodological Issues and Their Empirical Consequences," in *Hispanics in the U.S. Economy*, edited by G. Borjas and M. Tienda (New York: Academic Press, 1985), pp. 77-125.
80. Abowd, John M. "Economic and Statistical Analysis of Discrimination in Job Assignment," *Industrial Relations Research Association Proceedings of the Thirty-Sixth Annual Meetings* (Madison, WI: IRRRA, 1984), pp. 34-47. (invited paper)
81. Abowd, John M. and Mark Killingsworth "Do Minority/White Unemployment Differences Really Exist," *Journal of Business and Economic Statistics* 2 (January 1984): 64-72.

82. Abowd, John M. and Arnold Zellner "Estimating Gross Labor Force Flows," *American Statistical Association 1983 Proceedings of the Business and Economic Statistics Section* (Washington, DC: ASA, 1983), pp. 162-67.
83. Abowd, John M. and Mark Killingsworth "Sex Discrimination, Atrophy and the Male-Female Wage Differential," *Industrial Relations* 22 (Fall 1983): 387-402.
84. Abowd, John M. and Henry S. Farber "Job Queues and the Union Status of Workers," *Industrial and Labor Relations Review* 35 (April 1982): 354-67. [download]
85. Abowd, John M. and Orley Ashenfelter "Anticipated Unemployment, Temporary Layoffs and Compensating Wage Differentials," in *Studies in Labor Markets*, edited by S. Rosen (Chicago: University of Chicago Press for the NBER, 1981), pp. 141-170. [download]
86. Abowd, John M. "An Econometric Model of Higher Education," in *Managing Higher Education: Economic Perspectives*, A Monograph of the Center for the Management of Public and Nonprofit Enterprises (Chicago: University of Chicago Press, 1981), pp. 1-56.
87. Mulvey, Charles and John M. Abowd "Estimating the Union/Nonunion Wage Differential: A Statistical Issue," *Economica*, 47 (February 1980): 73-79.
88. Abowd, John M. and T. James Trussell "Teenage Mothers, Labor Force Participation, and Wage Rates," *Canadian Studies in Population* (1980): 33-48.

Monographs

1. Abowd, John M., Martha H. Stinson and Gary Benedetto *Final Report to the Social Security Administration on the SIPP/SSA/IRS Public Use File Project*, November 2006. [download archival copy and Excel tables at <http://hdl.handle.net/1813/43929>]
2. Abowd, John M. and Michael Bognanno "The Center for Advanced Human Resource Studies Managerial Compensation Database: User's Guide," March 1991.
3. Abowd, John M. and Michael Bognanno "The Center for Advanced Human Resource Studies Managerial Compensation Database: Technical Guide," March 1991.
4. Abowd, John M. *An Econometric Model of the U.S. Market for Higher Education* (New York: Garland Press, 1984).
5. Abowd, John M. and Mark Killingsworth "Employment, Wages, and Earnings of Hispanics in the federal and Nonfederal Sectors," in *Hispanics in the Labor Force: A Conference Report*, edited by G. Borjas and M. Tienda. Final Report to the National Employment Policy Commission (Washington, DC: GPO, 1982).
6. Abowd, John M. "Program Evaluation: New Panel Data Methods for Evaluating Training Effects," in *Program Evaluation Final Report to the U.S. Department of Labor* (Contract No. 23-17-80-01) (Washington, DC: NTIS, 1983).
7. Abowd, John M. and Mark Killingsworth "Employment, Wages, and Earnings of Hispanics in the federal and Nonfederal Sectors," in *Hispanics in the Labor Force: A Conference Report*, edited by G. Borjas and M. Tienda. Final Report to the National Employment Policy Commission (Washington, DC: GPO, 1982).
8. Abowd, John M. "Minority Unemployment, Compensating Differentials and the Effectiveness of the EEOC," in *Issues in Minority and Youth Unemployment final Report to the U.S. Department of Labor* (Contract No. 20-17-80-44) (Washington, DC: NTIS, 1982)
9. Abowd, John M. and Mark Killingsworth "Structural Models of the Effects of Minimum Wages on Employment by Age Groups," *Final Report of the Minimum Wage Study Commission*, Volume 5 (Washington, DC: GPO, 1981).
10. Abowd, John M. and Mark Killingsworth "An Analysis of Hispanic Employment, Earnings and Wages with Special Reference to Puerto Ricans," *Final Report to the U.S. Department of Labor* (Grant 21-36-78-61) (Washington, DC: NTIS, 1981).

Miscellany

1. Abowd, John M., Ian M. Schmutte, William Sexton, and Lars Vilhuber, Introductory Readings in Formal Privacy for Economists (May 8, 2019, updated regularly). [read, download]
2. Abowd, John M., "The Census Bureau Tries to Be a Good Data Steward in the 21st Century" International Conference on Machine Learning (ICML) 2019 keynote address. [video, start at minute 18:00] [slides]
3. Garfinkel, Simson L., John M. Abowd, and Christian Martindale, "Understanding Database Reconstruction Attacks on Public Data," *ACMQueue*, Vol. 16, No. 5 (September/October 2018): 28-53. [download, not copyrighted]

4. Garfinkel, Simson L., John M. Abowd and Sarah Powazek "Issues Encountered Deploying Differential Privacy," *WPES'18 Proceedings of the 2018 Workshop on Privacy in the Electronic Society*, Ontario, CA (October 2018): 133-137, DOI:10.1145/3267323.3268949. [ArXiv preprint]
5. Abowd, John M. "The U.S. Census Bureau Adopts Differential Privacy," KDD '18 Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, London, UK (August 2018): 2867, DOI:10.1145/3219819.3226070. [download, subscription required], [archival copy] [video]
6. Abowd, John M., Lorenzo Alvisi, Cynthia Dwork, Sampath Kannan, Ashwin Machanavajjhala, and Jerome Reiter "Privacy-Preserving Data Analysis for Federal Statistical Agencies," Computing Community Consortium White Papers (January 2017). [CCC white paper archive; ArXiv preprint]
7. Abowd, John M. "Why Statistical Agencies Need to Take Privacy-loss Budgets Seriously, and What It Means When They Do," presented to the Federal Committee on Statistical Methodology, Policy Conference, December 7-8, 2016. [download]
8. Vilhuber, Lars, John M. Abowd and Jerome P. Reiter "Synthetic Establishment Microdata around the World," *Statistical Journal of the International Association for Official Statistics*, Vol. 32 (2016): 65-68. [download, open access] [download preprint]
9. Abowd, John M. "Synthetic Establishment Data: Origins and Introduction to Current Research," *Statistical Journal of the International Association for Official Statistics*, Vol. 30, No. 2 (Summer 2014): 113-115. [download, subscription required] [download preprint]
10. Benedetto, Gary, Martha H. Stinson and John M. Abowd "The Creation and Use of the SIPP Synthetic Beta," U.S. Census Bureau Technical Paper (April 2013). [download]
11. Abowd, John M. and Lars Vilhuber "Science, Confidentiality, and the Public Interest," *Chance*, Vol. 24, No. 3 (Fall 2011): 58-62. [download]
12. Abowd, John M. "OnTheMap: Block-level Job Estimates Based on Longitudinally Integrated Employer-Employee Micro-data," *Association of Public Data Users Newsletter* Vol. 33, No. 2 (March/April 2010): 10-19. [download]
13. Abowd, John M. Kobbi Nissim and Chris Skinner "First Issue Editorial" *Journal of Privacy and Confidentiality*, Vol. 1, No. 1 (2009): 1-6. [download]
14. Abowd, John M. "Comments on 'Regional difference-in-differences in France using the German annexation of Alsace-Moselle in 1870-1918' by Matthieu Chemin and Etienne Wasmer" *NBER International Seminar on Macroeconomics* (2008): 306-309. [download]
15. Abowd, John M. and Julia Lane "The Economics of Data Confidentiality," *ICP Bulletin*, Volume 4, No. 2 (August 2007): 18-21. [download preprint]
16. Abowd, John M. "Rapporteur comments: International Symposium on Linked Employer-Employee Data, Econometric Issues" *Monthly Labor Review* 121:7 (July, 1998): 52-53.
17. Abowd, John M. "Discussion of 'How much do immigration and trade affect labor market outcomes' by Geroge J. Borjas, Richard B. Freeman and Lawrence F. Katz." *Brookings Papers in Economic Activity* (1997:1): 76-82.
18. Abowd, John M. "Discussion of Gross Worker and Job Flows in Europe by M. Burda and C. Wyplosz." *European Economic Review* (1994): 1316-1320.
19. Abowd, John M. "Discussion of 'The Quality Dimension in Army Retention' by Charles Brown." in A. Meltzer (ed.) *The Carnegie-Rochester Conference on Public Policy* 33 (1990).
20. Abowd, John M. "Immigration, Trade, and Labor Markets in Australia and Canada," in *Immigration, Trade, and the Labor Market*, edited by R.B. Freeman (Cambridge, Mass: NBER, 1988), pp. 29-34.
21. Abowd, John M. "Discussion of 'Public Sector Union Growth and Bargaining Laws: A Proportional Hazards Approach with Time-Varying Treatments' by c. Ichniowski." in *Public Sector Unionism*, edited by R. Freeman (Chicago: University of Chicago Press for the NBER, 1988).
22. Abowd, John M., Ross Stolzenberg and Roseann Giarusso "Abandoning the Myth of the Modern MBA Student," *Selections The Magazine of the Graduate Management Admission Council* (Autumn 1986): 9-21.
23. Abowd, John M., Brent Moulton and Arnold Zellner "The Bayesian Regression Analysis Package: BRAP User's Manual Version 2.0," H.G.B. Alexander Research Foundation, Graduate School of Business, University of Chicago, 1985.
24. Abowd, John M. and Mark R. Killingsworth "The Minimum Wage Law Winners and Losers," *The Wall Street Journal* (August 1981).

Working and Unpublished Papers

1. McKinney, Kevin L. and John M. Abowd, "Male Earnings Volatility in LEHD before, during, and after the Great Recession," (August 2020). [[download preprint](#)]
2. Abowd, John M., Gary L. Benedetto, Simson L. Garfinkel et al. "The Modernization of Statistical Disclosure Limitation at the U.S. Census Bureau," (August 2020). [[download preprint](#)]
3. Abowd, John M., Ian M. Schmutte, William Sexton, and Lars Vilhuber "Suboptimal Provision of Privacy and Statistical Accuracy When They are Public Goods," (June 2019). [[download preprint](#)]
4. Abowd, John M., Joelle Abramowitz, Margaret C. Levenstein, Kristin McCue, Dhiren Patki, Trivellore Raghunathan, Ann M. Rodgers, Matthew D. Shapiro, Nada Wasi, 2019. "Optimal Probabilistic Record Linkage: Best Practice for Linking Employers in Survey and Administrative Data," Working Papers 19-08, Center for Economic Studies, U.S. Census Bureau, handle: RePEc:cen:wpaper:19-08. [[download preprint](#)]
5. McKinney, Kevin L. Andrew Green, Lars Vilhuber, and John M. Abowd "Total Error and Variability Measures with Integrated Disclosure Limitation for Quarterly Workforce Indicators and LEHD Origin Destination Employment Statistics in On The Map" (December 2017). [[download preprint](#)]
6. Abowd, John M. and Ian Schmutte "Revisiting the Economics of Privacy: Population Statistics and Confidentiality Protection as Public Goods" (April 2017), [[download preprint](#)], published as Abowd, John M. and Ian M. Schmutte "An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices," *American Economic Review*, Vol. 109, No. 1 (January 2019):171-202, DOI:10.1257/aer.20170627. [[AER](#), [ArXiv preprint](#), [Replication information](#)]
7. Abowd, John M. "Where Have All the (Good) Jobs Gone? (May 2014) Society of Labor Economists Presidential Address. [[download preprint](#)] [[accompanying audio](#)]
8. Abowd, John M., John Haltiwanger, Julia Lane, Kevin McKinney and Kristin Sandusky "Technology and Skill: An Analysis of Within and Between Firm Differences" (March 2007) NBER WP-13043. [[download preprint](#)]
9. Abowd, John M., Francis Kramarz, David N. Margolis, and Thomas Philippon "Minimum Wages and Employment in France and the United States" (February 2006). [[archival download](#)]
10. Abowd, John M., Paul Lengerhmann and Kevin L. McKinney "The Measurement of Human Capital in the U.S. Economy," (March 2003) [[download Census](#), cited on September 1, 2015] [[archival download](#)]
11. Abowd, John M., Robert Creecy and Francis Kramarz "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," (March 2002). [[download Census](#), cited on September 1, 2015] [[archival download](#)] [[Fortran source](#)] [[Support files](#)] [[VirtualRDC archive](#)]

MAJOR GRANTS AND RESEARCH CONTRACTS

1. Associate Director for Research and Methodology and Chief Scientist U.S. Census Bureau, Intergovernmental Personnel Act (IPA) with Cornell University, June 1, 2016—March 27, 2020.
2. Research and Methodology Support Services, U.S. Census Bureau contract with Cornell University, June 1, 2015—May 31, 2016, \$268,897.
3. The Economics of Socially Efficient Privacy and Confidentiality Management for Statistical Agencies, Alfred P. Sloan Foundation awarded to Cornell University, April 1, 2015—March 31, 2019, \$535,970. (co-PIs Lars Vilhuber and Ian Schmutte)
4. RCN: Coordination of the NSF-Census Research Network, National Science Foundation SES 1237602 awarded to the National Institute of Statistical Sciences, July 15, 2012—June 30, 2017, transferred to Cornell University, September 2014, \$748,577. (PI Lars Vilhuber, other co-PIs Alan Karr, Jerome Reiter)
5. NCRN-MN: Cornell Census-NSF Research Node: Integrated Research Support, Training and Data Documentation, National Science Foundation Grant SES 1131848 awarded to Cornell University, October 1, 2011—September 30, 2016, \$2,999,614. (with William Block, Ping Li, and Lars Vilhuber)
6. A Census-Enhanced Health and Retirement Study: A Proposal to Create and Analyze an HRS Dataset Enhanced with Characteristics of Employers, Alfred P. Sloan Foundation grant awarded to the Institute for Social Research, University of Michigan with a subcontract to

- Cornell University, September 1, 2011–August 31, 2016, Cornell component \$349,608. (PI: Margaret Levenstein; other co-PIs: Matthew Shapiro, Kristin McCue and David Weir)
7. **Synthetic Data User Testing and Dissemination**, National Science Foundation Grant SES 1042181 awarded to Cornell University, September 15, 2010 to September 14, 2013, \$197,170. (Co-PI Lars Vilhuber)
 8. **CDI-Type II: Collaborative Research: Integrating Statistical and Computational Approaches to Privacy**, National Science Foundation Grant BCS 0941226 awarded to Cornell University, September 1, 2010–August 31, 2014, \$409,296. (Other PIs: Aleksandra B Slavkovic, Stephen E. Fienberg, Sofya Raskhodnikova, and Adam Smith)
 9. **TC:Large: Collaborative Research: Practical Privacy: Metrics and Methods for Protecting Record-level and Relational Data**, National Science Foundation Grant TC 1012593 awarded to Cornell University, July 15, 2010 to July 14, 2015, \$1,326,660. (Other PIs: Johannes Gehrke, Gerome Miklau, and Jerome Reiter)
 10. **The Longitudinal Employer-Household Dynamics Program**, U.S. Bureau of the Census, Interagency Personnel Act (IPA) with Cornell University, September 18, 1998 – September 17, 2000, \$260,000; renewed September 14, 2000–September 13, 2002, \$320,000; contract renewed as consultant September 14, 2002–September 13, 2003 (\$120,000); renewed as IPA September 15, 2003 – September 14, 2005 (\$384,590); renewed as IPA September 15, 2005–September 14, 2007 (\$425,215); new September 15, 2008–September 14, 2010 (497,897); renewed September 15, 2010–September 14, 2012 (532,893); continued as a contract with ACES-Research, LLC (September 17, 2012–September 16, 2013); re-established as IPA October 1, 2013–September 30, 2014 (\$231,757); re-established as IPA November 14, 2014 –May 31, 2015 (\$229,095).
 11. **Social Science Gateway to TeraGrid**, National Science Foundation Grant SES 0922005 awarded to Cornell University, July 1, 2009 to June 30, 2012, \$393,523. (Co-PI Lars Vilhuber) [Cornell Chronicle Article] [ILR News Release]
 12. **Joint NSF-Census-IRS Workshop on Synthetic Data and Confidentiality Protection**, July 2009 Washington, DC, National Science Foundation Grant SES 0922494 awarded to Cornell University, July 1, 2009 to June 30, 2010, \$18,480. (Co-PIs Lars Vilhuber, Jerome Reiter, and Ron Jarmin)
 13. **The Economics of Mass Layoffs: Displaced Workers, Displacing Firms, Causes and Consequences**, National Science Foundation Grant SES-0820349 awarded to Cornell University, October 1, 2008 to September 30, 2010, \$245,950. (Co-PI Lars Vilhuber)
 14. **LEHD Developmental and Confidentiality Research**, Census Bureau Contract to Abt Associates with subcontract awarded to Cornell University, August 1, 2007 to September 30, 2008, \$358,270.
 15. **CT-T: Collaborative Research: Preserving Utility While Ensuring Privacy for Linked Data**, National Science Foundation Grant CNS-0627680 awarded to Cornell University, September 5, 2006 to August 31, 2009, \$488,950. (PI Johannes Gehrke)
 16. **LEHD Confidentiality Research**, Census Bureau Contract to Abt Associates with subcontract awarded to Cornell University, October 1, 2004 to September 30, 2005, \$230,155.
 17. **ITR-(ECS+ASE)-(dmc+int): Info Tech Challenges for Secure Access to Confidential Social Science Data**, National Science Foundation Grant SES-0427889 awarded to Cornell University, October 1, 2004 to September 30, 2007, \$2,938,000. (Co-PIs Matthew D. Shapiro, Ronald Jarmin, Stephen F. Roehrig, and Trivellore Raghunathan) [Cornell Chronicle article]
 18. **EITM: Developing the Tools to Understand Human Performance: An Empirical Infrastructure to Foster Research Collaboration**, National Science Foundation Grant SES-0339191 awarded to Cornell University, October 1, 2004 to September 30, 2007, \$337,455 (Co-PIs John Haltiwanger and Ron Jarmin)
 19. **The New York Research Data Center**, National Science Foundation Grant SES-0322902 awarded to the NBER, August 1, 2003 to July 31, 2004, \$300,000. (PI Neil G. Bennett, Other co-PIs Bart Hobijn, Erica L. Groshen, Robert E. Lipsey)
 20. **Workshop on Confidentiality Research**, National Science Foundation Grant SES-0328395 awarded to the Urban Institute, June 1, 2003 – May 31, 2004, \$43,602. (Co-PI Julia Lane)
 21. **Firms, Workers and Workforce Quality: Implications for Earnings Inequality and Economic Growth**, Alfred P. Sloan Foundation Grant 22319-000-00 awarded to the Urban Institute, January 2003–January 2006, \$1,400,000. (Co-PIs John Haltiwanger, Julia Lane, J. Bradford Jensen, Fredrick Knickerbocker, and Ronald Prevost)
 22. **The Demand for Older Workers: Using Linked Employer-Employee Data for Aging Research**, National Institute on Aging, R01-AG18854-01 to Cornell University, July 1, 2002 – April 30,

- 2007, \$1,753,637. (Co-PIs John Haltiwanger, Andrew Hildreth, and Julia Lane)
23. Workers and Firms in the Low-wage Labor Market: Interactions and Long Run Dynamics, Russell Sage Foundation, Rockefeller Foundation, and Department of Health and Human Services (ASPE) to the Urban Institute \$700,000, September 1, 2001 August 31, 2003. (Co-PIs John Haltiwanger, Harry Holzer, and Julia Lane)
 24. From Workshop Floor to Workforce Clusters: A New View of the Firm, Alfred P. Sloan Foundation, 99-12-12 to the Urban Institute, March 1, 2000 – March 31, 2002, \$314,604. (Co-PIs John Haltiwanger and Julia Lane)
 25. Dynamic Employer-Household Data and the Social Data Infrastructure, National Science Foundation, SES-9978093 to Cornell University, September 28, 1999 – September 27, 2005, \$4,084,634. (Co-PIs John Haltiwanger and Julia Lane)
 26. The Longitudinal Employer-Household Dynamics Program, National Institute on Aging, interagency funding to the United States Census Bureau, September, 1999 – August, 2001, \$490,000. Renewed September 2001– August 2004, \$750,000 (Co-PIs John Haltiwanger and Julia Lane) [Cornell Chronicle article]
 27. Individual and Firm Heterogeneity in Labor Markets: Studies of Matched Employee-Employer Data, National Science Foundation SBR 9618111 to the NBER, March 15, 1997 – February 28, 2002, \$243,361.
 28. Creation of an Employer Identification Link File and Addition of Employer Information to the National Longitudinal Survey of Youth 1979 Cohort, Bureau of Labor Statistics (subcontracted by NORC, University of Chicago, Chicago, IL 60637), July 1, 1995 – December 31, 1997, \$82,946.
 29. Employment and Compensation Policies: Studies of American and French Labor Markets Using Matched Employer-Employee Data, National Science Foundation SBR 9321053 to the NBER, July 1, 1994 – June 31, 1997, \$ 185,257. (Co-PIs David Margolis and Kenneth Troske)
 30. Compensation System Design, Employment and Firm Performance: An Analysis of French Microdata and a Comparison to the United States, National Science Foundation, SBR 9111186 to Cornell University, July 1, 1991 – December 30, 1994, \$174,565.
 31. The Effects of Collective Bargaining and Threats of Unionization on Firm Investment Policy, Return on Investment, and Stock Valuation, National Science Foundation, SES 8813847 to the NBER, July 1, 1988 – June 30, 1990, \$81,107.
 32. Improving the Scientific Research Utility of Labor Force Gross Flow Data, National Science Foundation, SES 85-13700 to the NBER, April 15, 1986 – March 31, 1988, \$69,993.
 33. Program Evaluation: New Panel Data Methods for Evaluating Training Effects, U.S. Department of Labor Contract 23-17-80-01 to NORC at the University of Chicago, 1983.
 34. Minority Unemployment, Compensating Differentials and the Effectiveness of the EEOC, U.S. Department of Labor Contract 20-17-80-44 to NORC at the University of Chicago, 1982.
 35. An Analysis of Hispanic Employment, Earnings and Wages with Special Reference to Puerto Ricans, U.S. Department of Labor Grant 21-36-78-61, 1981.

PROFESSIONAL SERVICE, SURVEYS, AND DATA COLLECTION

1. Canadian Research Data Centre Network Inaugural Board 2017-2019.
2. American Economic Association, Committee on Economic Statistics (AEAWeb) 2013-2018.
3. National Academy of Sciences, Committee on National Statistics (CNSTAT) 2010-2013; reappointed 2013-2016.
4. National Academy of Sciences, CNSTAT, Panel on Measuring and Collecting Pay Information from U.S. Employers by Gender, Race, and National Origin, (Chair) 2011-2012.
5. National Academy of Sciences, CNSTAT, Panel on Measuring Business Formation, Dynamics and Performance, 2004-2007.
6. National Academy of Sciences, CNSTAT, Panel on Data Access for Research Purposes, 2002-2005.
7. Executive Committee, Conference on Research in Income and Wealth 2002-.
8. Distinguished Senior Research Fellow, LEHD Program, U.S. Census Bureau 1998-2016.
9. Social Science and Humanities Research Council (Canada), Major Collaborative Research Initiatives review panel, 1997, 1998.
10. Technical Advisory Board for the National Longitudinal Surveys of the Bureau of Labor Statistics, 1988-1990, 1992-2001, Chair 1999-2001.
11. National Science Foundation, Economics Panel, 1990-91, 1992-93; KDI Panel 1999; Infrastructure Panel 2000; CDI Panel 2008; CDI Panel 2009.

12. Principal Investigator for The Center for Advanced Human Resource Studies Managerial Compensation Data Base. sponsored by the Cornell University Center for Advanced Human Resource Studies, 1989-1994.
13. Principal Investigator for A Longitudinal Data Base of Collective Bargaining Agreements. Sponsored by the Bureau of National Affairs and the University of Chicago Graduate School of Business, 1985.

PROFESSIONAL ORGANIZATIONS

1. American Economic Association
2. American Statistical Association
3. Econometric Society
4. Society of Labor Economists
5. International Statistical Institute
6. International Association for Official Statistics
7. National Association for Business Economics
8. American Association of Wine Economists
9. American Association for Public Opinion Research
10. Association for Computing Machinery
11. American Association for the Advancement of Science

PERSONAL INFORMATION

United States citizen

Personal email: john.abowd@gmail.com

APPENDIX B – 2010 RECONSTRUCTION-ABETTED RE-IDENTIFICATION SIMULATED ATTACK

1. This appendix provides a high-level summary of the reconstruction-abetted re-identification attack simulation that the Census Bureau conducted on the released 2010 Census data. To assess the risk of a reconstruction-abetted re-identification attack, the Census Bureau conducted a series of statistical exercises to quantify the contemporaneous and future risk that individual responses could be disclosed. The Census Bureau has completed two simulated attacks that address the re-identification risk of a 100% microdata file (a file with detailed, individual-level records for every person enumerated in the census) reconstructed from the published Summary File 1 data. The 2010 Summary File 1, usually called SF1, includes the 2010 P.L. 94-171 Redistricting Data Summary File, the 2010 Advanced Group Quarters Data Summary File, and the bulk of the demographic and housing characteristics released from the 2010 Census in tabular format.¹ The fundamental structure of these simulations is as follows.

SIMULATED RECONSTRUCTION ATTACK

2. Database reconstruction is the process of statistically re-creating the individual-level records from which a set of published tabulations was originally calculated. That is, database reconstruction attempts to “reverse engineer” the confidential input data used in a statistical tabulation system.
3. The Census Bureau released over 150 billion statistics as part of the 2010 Census. The simulated reconstruction attack used as its input a small fraction of those statistics—approximately 6.2 billion statistics contained in the following published SF1 tables from the 2010 Census:

P001 (Total Population by Block)
P006 (Total Races Tallied by Block)
P007 (Hispanic or Latino Origin by Race by Block)

¹ See the technical documents in [Summary File 1 Dataset \(census.gov\)](https://www.census.gov/data/tables/2010/sf1.html).

P009 (Hispanic or Latino, and Not Hispanic or Latino by Race by Block)
 P011 (Hispanic or Latino, and Not Hispanic or Latino by Race for the Population 18 Years and Over by Block)
 P012 (Sex by Age by Block)
 P012A-I (Sex by Age by Block, iterated by Race)
 P014 (Sex by Single-year-of-age for the Population under 20 Years by Block)
 PCT012A-N (Sex by Single-year-of-age by Tract, iterated by Race)

4. The reconstruction of the 2010 Census microdata for the sex, age, race, Hispanic/Latino ethnicity, and census block variables was carried out by constructing a system of equations consistent with the published tables listed above that, once solved, could then be converted into microdata. This system of equations was solved using commercial mixed-integer linear programming software (Gurobi).
5. Because the parameters of the 2010 Census swapping methodology included invariants on total population and voting age population at the block level, the reconstruction was able to exactly reconstruct all 308,745,538 million records with correct block location and voting age (18+). Then, leveraging the race (63 categories), Hispanic/Latino origin, sex, and age (in years) data from the specified tables, the simulated attack was able to further reconstruct those variables on the individual-level records.
6. To assess the accuracy of these reconstructed individual-level records, the team performed exact record linkage of the five variables in the reconstructed microdata to the same five variables in the Census Edited File (CEF, the confidential data) and Hundred-percent Detail File (HDF, the confidential swapped individual-level data before tabulation). The results are summarized in Table 1. The “left” file of the record linkage is in the first column. The “right” file is the reconstructed microdata from SF1.

Left file	Record Counts		Agreement Rates		
	In Left	In Reconstructed	Exact	Fuzzy Age	One error
CEF	308,745,538	308,745,538	46.48%	70.98%	78.31%
HDF	308,745,538	308,745,538	48.34%	73.33%	80.39%

DRB clearance number CBDRB-FY21-DSEP-003

7. The agreement rates shown in Table 1 include block (which was never wrong), sex, age (in years), race (63 OMB categories), and Hispanic ethnicity and are computed as a percentage of the total population. Exact agreement means all five variables agreed precisely bit for bit. Fuzzy-age agreement means that block, sex, race, and Hispanic ethnicity agreed exactly, but age agreed only +/- 1 year (e.g., age 25 on the CEF is in fuzzy-age agreement with ages 24, 25, and 26 on the reconstructed data). The one-error agreement rate allows one variable – sex, age (outside +/- one year), race or ethnicity to be wrong.
8. Most errors in the reconstructed file are that the age variable is off by +/- 2 years rather than +/- 1 year. This error is the balance of the width of the 5-year categories used in the block-level summaries. Hence, even though the disclosure avoidance requirement for the 2010 Census SF1 tabular summaries specified block-level aggregation to 5-year bins for those age 20 and over, the effective aggregation was far less.
9. Figure 1 shows the distribution of agreement rates by block size. Agreement rates are only substantially lower than the population averages shown in Table 1 for blocks with populations between 0 and 9 people, which is where the Census Bureau has said it concentrated the swaps.² However, uniqueness on sex, age, race, and ethnicity is not limited to small population blocks. *This is one of the principal failures of the 2010 tabular disclosure avoidance methodology – swapping provided protection for households deemed “at risk,” primarily those in blocks with small populations, whereas for the entire 2010 Census a full 57% of the persons are population uniques on the basis of block, sex,*

² McKenna, L. (2018), “Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing,” <https://www.census.gov/content/dam/Census/library/working-papers/2018/adrm/Disclosure%20Avoidance%20for%20the%201970-2010%20Censuses.pdf>, p. 8.

age (in years), race (OMB 63 categories), and ethnicity. Furthermore, 44% are population uniques on block, age and sex.³




Figure 1 Block-level agreement rates between the reconstructed 2010 Census micro-data and the 2010 Census Edited File by population in the block
DRB clearance number CBDRB-FY21-DSEP-003.

10. Although there are no recent re-identification studies for decennial Public Use Micro-data Samples (PUMS) with geography coded to the Public Use Microdata Area (PUMA), the Census Bureau continues to use 100,000 persons as the minimum population threshold for such areas and has coded geography on the 2010 PUMS and all American Community Survey (ACS) PUMS using these PUMAs. Since sex and age (single years) are population uniques at the tract level for only 0.18% of persons, this may still be justifiable for a 10% sample of 2010 Census records, but the potential re-

³ The statistics in this paragraph are cleared for public release by the Census Bureau Disclosure Review Board (CBDRB-FY21-DSEP-003).

identification rate for a 100% public-use microdata file geocoded to the block level is certainly quite large.

11. The reconstruction experiment demonstrated that existing technology can convert the Census Bureau's traditional tabular summaries of Census data which was released in 2010 into a 100% coverage microdata file geocoded to the block level with very limited noise which was not released in 2010. This microdata file contains so much detail that it would have been deemed "unreleasable" if it had been proposed in conjunction with the original 2010 Census data products.
12. The ability to reconstruct the microdata means that there is now a significant disclosure risk for the 2010 Census Summary Files 1 and 2 (SF1, SF2) and the American Indian Alaska Native Summary File (AIANSF) data. There are approximately 150 billion statistics in the SF1, SF2, and AIANSF summaries (recall that the 2010 P.L. 94-171 Redistricting Data Summary File and the 2010 Advanced Groups Quarters Summary File are part of SF1). Because of the features noted above, releasing this many very accurate statistics made the ensemble of those publications equivalent to releasing the 2010 Hundred-percent Detail File (HDF), the swapped version of and the 2010 Census Edited File (CEF). There can be no uncertainty about this: *the 2010 Census tabular publications were equivalent to releasing every tabulation variable in the 2010 HDF in universe public-use microdata files without the hierarchical structure--person and household records can be fully reconstructed, but not directly linked to each other.* The team that demonstrated this vulnerability stopped after reconstructing person-level records for block, sex, age (in years), race (63 OMB categories), and Hispanic ethnicity because the vulnerability had been fully exposed mathematically and demonstrated empirically.
13. There are 308,745,538 (U.S. only) person records and 131,704,730 housing unit records in both the 2010 HDF and CEF, linked in their correct hierarchy. For the unswapped records in HDF, the images are identical to their CEF counterparts. For the swapped household records, the block identifier, household size, adult (age 18+) household

size, occupancy, and tenure variables are identical to their unswapped counterparts and on the person record the voting-age variable is identical to the unswapped counterpart.

14. As the documentation in McKenna (2018, 2019a) makes clear, a public-use microdata file containing the 308,745,538 person records in the HDF including only the five tabulation variables block, sex, age (in years), race (63 OMB categories), and Hispanic ethnicity is so disclosive that it would not have passed the disclosure avoidance criteria used for the 2010 Census Public-Use Microdata Sample.⁴ Furthermore, the same file would not have passed the disclosure avoidance criteria applied to SF1 itself.⁵ The official 2010 PUMS had a geographic population threshold of 100,000, collapsed categories to national population thresholds of 10,000, used partially synthetic data for the group quarters population, and “topcoding, bottom-coding, and noise infusion for large households.” The PUMS was sampled from the swapped version of the 2010 HDF, not the Census Edited File.

15. The additional disclosure avoidance methods used for the 2010 PUMS are explicitly noted on pages 2-1 and 2-2 of its technical documentation. The definition of a Public Use Microdata Area also explicitly references its confidentiality protection purpose:

“The Public Use Microdata Sample (PUMS) files contain geographic units known as Public Use Microdata Areas (PUMAs). To maintain the confidentiality of the PUMS data, a minimum population threshold of 100,000 is set for PUMAs. Each state is separately identified and may be comprised of one or more PUMAs. PUMAs do not cross state lines. (page 1-2, emphasis added)”

⁴ McKenna, L. (2019a) “Disclosure Avoidance Techniques Used for the 1960 Through 2010 Decennial Censuses of Population and Housing Public Use Microdata Samples,” Research and Methodology Technical Report available at [Disclosure Avoidance Techniques Used for the 1960 Through 2010 Census](#).

⁵ McKenna, L. (2018)

16. This failure to apply microdata disclosure avoidance matters because the reconstructed 2010 microdata for block, sex, age (in years), race (63 OMB categories), and Hispanic ethnicity are a very accurate image of the HDF, and the HDF is a very accurate image of the CEF, which is the reason that it is also confidential. Consequently, the new technology-enabled possibility of accurately re-constructing HDF microdata from the published tabular summaries and the fact that those reconstructed data do not meet the disclosure avoidance standards established at the time for microdata products derived from the HDF demonstrate that the swapping methodology as implemented for the 2010 Census no longer meets the acceptable disclosure risk standards established when that swapping mechanism was selected for the 2010 Census.
17. Having demonstrated that a 100% microdata file can be successfully reconstructed from the published 2010 Census tabulations, the Census Bureau proceeded to use these reconstructed microdata to simulate a re-identification attack on those data.

DE-IDENTIFICATION ATTACK SIMULATION

18. The simulated re-identification attack proceeds as follows. Identify a person-level data source file that contains name, address, sex, and birthdate (e.g., commercially available data). Convert the names and addresses to their corresponding Census Bureau Protected Identification Key (PIK). Identify the corresponding census block for every address in the source file. Then, looping through all the records in the reconstructed microdata file produced from the reconstruction, find the first record in the source file that matches exactly on block, sex, and age. Once this step is completed, run through the remaining unmatched records from the reconstructed microdata and find the first unmatched record from the source file that matches exactly on block and sex, and matches on age plus or minus 1 year.
19. When both steps have been completed, output the records with successful matches from these two passes. These are called *putative re-identifications* because they appear

to link the reconstructed microdata to a real name and address associated with the block, sex, age, race, and ethnicity on the reconstructed microdata. These are the records the hypothetical attacker thinks are re-identified.

20. Putative re-identifications are not necessarily correct. An external attacker would have to do extra field work to estimate the *confirmation rate* – the percentage of putative re-identifications that are correct. An external attacker might estimate the confirmation rate by contacting a sample of the putative re-identifications to confirm the name and address. An external attacker might also perform more sophisticated verification using multiple source files to select the name and address most consistent with all source files and the reconstructed microdata.
21. At the Census Bureau we usually estimate the confirmation rate as a percentage of the total population, not as a percentage of the putative re-identifications, by performing a similar record linkage exercise of the putative re-identifications against the CEF, looking for exact matches on all variables (including PIK, block, sex, age, race, and ethnicity), followed by a second pass looking for exact matches except age, which is allowed to vary by plus or minus 1 year. Once these two passes have been completed, the matched records are the confirmed re-identifications, using exact match on PIK, block, sex, race (63 OMB categories), and ethnicity and match on age +/- 1 year as the definition of correct. The remaining unmatched records from the putative re-identifications of the reconstructed data are the unconfirmed re-identifications.
22. Table 2 shows the results of two such re-identification confirmation exercises. The first of these uses the combined commercial databases from Experian Marketing Solutions Incorporated, Infogroup Incorporated, Melissa Data Corporation, Targus Information Corporation, and VSGI LLC as the source file for name, address, sex, and age. This exercise simulates data quality circa 2010 for an external attacker relying on the consumer information in these databases. These results are in the row labeled “Commercial.” This re-identification experiment was the basis for the statistics released at the

American Association for the Advancement of Science 2019 annual meeting. Putative re-identifications were 138 million (45% of the 2010 Census resident population of the U.S.). Confirmed re-identifications were 52 million (17% of the same population).

23. Using the commercial data as the source for name, address, sex, and age is, as discussed in the main declaration, a best-case assumption. We know that these data exist and were available circa 2010 because that is when the Census Bureau acquired them. An external attacker, using the versions that the Census Bureau acquired and the relatively straightforward methodology above, would succeed at least as often as we did. This means that at least 52 million persons enumerated during the 2010 Census could be correctly re-identified using the attack strategy outlined here.
24. Suppose the external attacker had name, address, sex, and age of much better quality than the five commercial sources above. How much better could that attacker do using exactly the same strategy? This question can be answered by substituting the name, address, sex, and age from the 2010 CEF as the source file in the putative re-identification simulation. This is not cheating because no extra information in the CEF such as race, ethnicity or household structure is used for the source file. Hence, it is a proper worst-case scenario, and the one historically used by the Census Bureau in assessing microdata re-identification risk (see McKenna 2019b). If the external data on name, address, sex, and age are comparable to the 2010 Census, then the attacker will putatively re-identify 238 million persons (77% of the 2010 Census resident U.S. population). Confirmed re-identifications will be 179 million (58% of the same population). This means that with the best quality external data, relative to the 2010 Census, as many as 179 million persons could be correctly re-identified using the attack strategy outlined here.

PIK, Block, Age, Sex Record Linkage Source	Available Records	Records with PIK, Block, Sex, and Age	Putative Re-identifications using Source	Confirmed Re-identifications
Commercial	413,137,184	286,671,152	137,709,807	52,038,366
CEF	308,745,538	279,179,329	238,175,305	178,958,726
DRB clearance number CBDRB-FY21-DSEP-003.				

25. The record linkage results reported in Table 2 can be interpreted using two additional statistical quality measures: the *recall rate* and the *precision rate*. Taken together, these measures assess how successful an attacker can be at re-identifying records and how confident the attacker would be in those re-identifications.
26. *Recall rate*. The recall rate is the percentage of available source records that are correctly re-identified. Its numerator is the same as the confirmation rate, but its denominator is the number of records in the source file with sufficient information to perform the putative re-identification record linkage. For the two source files analyzed in these experiments, Table 2 shows the denominators for the recall rate in the column “Records with PIK, Block, Sex, and Age,” which gives the count of records with sufficient information to generate a putative match. Table 3 shows the recall rates for the two experiments. Both are greater than the respective confirmation rate because both the commercial data and the CEF have fewer usable records than the U.S. resident population. A critical result is the recall rate of 64% when the CEF is used as the source file. This result means that an external attacker with high quality name, address, sex, and age information succeeds in re-identification almost two times in three.

Table 3 Confirmation and Recall Rates		
Source	Percentage of U.S. Resident Population (Confirmation Rate)	Percentage of Complete Data Population (Recall Rate)
Commercial	16.85%	18.15%
CEF	57.96%	64.10%
DRB clearance number CBDRB-FY21-DSEP-003.		

27. *Precision rate.* Precision is the ratio of confirmed to putative re-identifications. It answers the question “How often is the attacker’s claimed re-identification correct as a percentage of the names the attacker attached to reconstructed census microdata?” Table 4 summarizes the precision rates for the two experiments. The precision of the experiment reported in February 2019 was 38% (first row of Table 4). The precision of the worst-case experiment is 75% (second row of Table 4). *This result means that an attacker using high-quality name, address, sex, and age data is correct three times out four.*

Table 4 Precision Rates	
Source	Confirmed Percentage of Putative Re-identification (Precision Rate)
Commercial	37.79%
CEF	75.14%
DRB Clearance number CBDRB-FY21-DSEP-003.	

28. To be successful, an attacker does not have to be a commercial entity, nor does a successful attack need to use commercially available data. Many agencies of federal, state and local governments in the U.S. now possess high-quality data on name, address,

sex, and age. When preparing public-use microdata files that contain variables that other agencies can access exactly, it has long been the practice to coarsen such data to prevent non-statistical uses by other agencies (see McKenna 2019b). Applying such precautions to decennial census data products would imply severe limitations on the variables published at the block level, even in the presence of swapping.

29. In conclusion, the Census Bureau's simulated reconstruction-abetted re-identification attack definitively established that the tabular summaries from the 2010 Census could be used to reconstruct individual record-level data containing the tabulation variables with their most granular definitions. Such microdata violated the disclosure avoidance rules that the Data Stewardship Executive Policy Committee had established for the 2010 Census and would not have been released had they been proposed as an official product because they posed too great a disclosure risk. The disclosure risk presumed by the 2010 standards recognized the excessive risk of re-identification if block geographic identifiers were placed on a 100% enumeration microdata file along with age (in years) and sex. The Census Bureau believed in 2010 that the minimum population that the geographic identifier could represent in such microdata is 100,000 persons – the size of a Public-Use Microdata Area. That belief was strongly confirmed by the simulated re-identification attack. Somewhere between 52 and 179 million person who responded to the 2010 Census can be correctly re-identified from the re-constructed microdata.

FINAL 2/16/10**DSEP Meeting Record**

Topic: Updates

Meeting Date: 1/14/10

Attendees:

<i>Position</i>	<i>Attending for Position</i>
Deputy Director (Chair)	Tom Mesenbourg
AD, Administration	Ted Johnson
AD, Decennial	David Whitford
AD, Demographic	Cheryl Landman
AD, Economic	Harvey Monk
AD, Field	Marilia Matos
AD, IT	Brian McGrath
AD, Strategic Planning	Nancy Gordon
Rep. for Statistical Methodology	Tommy Wright
Senior Advisor for Data Management	Teresa Angueira
Chief, ITSO	
Chief, OAES	Kathleen Styles
Chief Privacy Officer	Mary Frazier
Also Attending:	Carol Comisarow, Ron Jarmin, David Raglin, Sharon Stern, Laura Zayatz, Michael Hawes

UPDATES

Background

[REDACTED]

Disclosure Review Board

- The DRB has been using enhanced disclosure avoidance procedures and methods for projects involving sensitive topics and/or sensitive populations. These procedures were implemented in response to an August 2004 memo from Director Kincannon. Though the memo was superseded by the Custom Tabulations policy, the DRB was not informed of this, and has not changed its procedures for sensitive topics and populations.
- Laura Zayatz also voiced the DRB's concern about planning for the 2020 Census and continuing to release data at the block level, as block populations continue to decrease (e.g. 40% of blocks in North Dakota have only 1 household in them)

[REDACTED]

Action Items

1. The DRB will develop recommendations for addressing the issue of disclosure review for sensitive populations. They will present their recommendations to DSEP once they have been vetted at the staff level.
2. DSEP agrees that the problem of block population size and disclosure avoidance is real, and that it deserves attention in the context of 2020 planning.

[REDACTED]

DSEP Meeting Record

Topics:



Public Use File Reidentification Threats Update

Meeting Date: February 5, 2015

Attendees:

<i>Position</i>	<i>Attending for Position</i>
Deputy Director (Chair)	<i>absent</i>
AD, Administration	<i>absent</i>
AD, Decennial	Lisa Blumerman
AD, Demographic	Enrique Lamas
AD, Economic	<i>absent</i>
AD, Field	Jay Keller
AD, IT	Avi Bender
AD, Research and Methodology	Tom Louis
AD, 2020 Census	Lisa Blumerman
AD, Communications	Kim Collier
AD, Performance Improvement	Susan Reeves
Chief, PCO/ Chief Privacy Officer	Robin Bachman
Chief Demographer	Howard Hogan
Chief Information Security Officer	Tim Ruland
Asst. Director, Research and Methodology	Ron Jarmin
Also Attending:	Barbara Downs, Randy Becker, Byron Crenshaw, Laura McKenna, Heather Madray, Raj Dwivedy, Julie Atwell, Mike Castro

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]




[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Public Use File Reidentification Threats Update

Background and Discussion:

On December 11, 2014, DSEP discussed a reidentification issue that occurred involving a Public Use File (PUF) produced as part of the New York City Housing Vacancy Survey (NYCHVS). At that meeting, DSEP commissioned a team to pursue the recommendations presented to DSEP in July 2014 in the paper titled “PUMS File Re-identification Threats and Potential Solutions for Mitigating those Threats.”

After discussing the logistics with some key stakeholders, and the difficulties of managing so many different angles on one team, DSEP approved a two-pronged approach to pursuing the paper’s recommendations.

The Center for Disclosure Avoidance Research (CDAR) has recently received authorization to hire new staff to focus primarily on synthetic data and reidentification research. This group is preparing a research proposal that focuses on the disclosure avoidance side of the PUF reidentification issue.

In addition to these efforts, the Demographic Programs Directorate (ADDP) will charter a team that focuses on the broader future of PUFs as well as some of the non-technical means of disclosure avoidance discussed in the July 2014 paper. This team will discuss Terms of Use for PUFs, restricted access, and other methods. This team will have representation from all of the impacted directorates. DSEP also recommended the team engage with external researchers on some of these ideas, and address their concerns.

Action Items:

- CDAR will prepare a research proposal to outline future Census Bureau efforts in Synthetic Data and Reidentification Research.
- ADDP will charter a team to evaluate the future of PUFs and explore some of the non-technical solutions outlined in the July 2014 paper.

DSEP Meeting Record

Topics: Initial Request for DSEP Determination on Disclosure Avoidance for the 2018 End-to-End Test of the 2020 Census of Population and Housing (John Abowd, ADRM)

Record-level Re-identification Linkages for Evaluating the 2010 and 2020 Census Disclosure Avoidance Systems (John Abowd, ADRM)



Meeting Date: May 10, 2017

<i>Position</i>	<i>Attending for Position</i>
Deputy Director (Chair)	Laura Furgione
CAO	David Ziaya
CFO	Joanne Crane
AD, Decennial	Lisa Blumerman
AD, Demographic	Karen Battle
AD, Economic	Ron Jarmin
AD, Field	Joan Hill
AD, IT	Nitin Naik
AD, Research and Methodology	John Abowd
AD, 2020 Census	Lisa Blumerman
AD, Communications	Stephen Buckner
AD, Performance Improvement	Ted Johnson
Chief, PCO/ Chief Privacy Officer	Robin Bachman
Chief Demographer	Howard Hogan
Senior Advisor Designee from the Director's Office	<i>absent</i>
Chief Information Security Officer	<i>absent</i>

Asst. Director, Research and Methodology	John Eltinge
Also Attending:	Simson Garfinkel, Byron Crenshaw, Eloise Parker, Ashley Landreth, Mike Castro, Harold Saintelien, Janean Darden, Julie Atwell

Initial Request for DSEP Determination on Disclosure Avoidance for the 2018 End-to-End Test of the 2020 Census of Population and Housing

Background:

The Census Bureau's Research and Methodology Directorate (ADRM) is researching and developing disclosure avoidance methods and systems to replace those used for Census 2000 and the 2010 that were not designed to protect against database reconstruction attacks. ADRM is establishing the 2020 Disclosure Avoidance System (DAS), a formally private system based on the theoretical model known as differential privacy. This is the available technology for controlling reconstruction attacks.

The 2020 DAS team is working to establish adjustable formal privacy parameters for the 2018 End-to-End test. They are seeking DSEP concurrence with the Disclosure Review Board's (DRB's) April 10, 2017 determination that six data elements of PL 94-171 can continue to be published as enumerated. The team will test methods and systems with these elements published as enumerated for the 2018 End-to-End with the goal of making sound recommendations to DSEP for the full 2020 DAS. These elements to be published as enumerated are:

- the number of occupied housing units per block,
- the number of vacant housing units per block,
- the number of households per block,
- the number of adults (age 18+) per block (where the definition of an adult is inferred from the structure of the PL94-171 age categories),
- the number of children (age less than 18) per block (where the definition of a child is also inferred from the structure of the PL94-171 age categories),
- and the number of persons per block.

ADRM expects to perform follow-up analyses of the test products developed for the End-to-End Test. Because there is no national sample in 2018, some aspects of the differentially private system cannot be implemented in the End-to-End Test. They will have to be simulated from the 2010 Census data. This means that the demonstration data from the test can be made as noisy as DSEP wishes. However, there is only time to implement algorithms that maintain confidentiality with the six data elements in the 2010 PL94-171 redistricting data. There will be both policy and disclosure avoidance issues surrounding how broadly those products can be disseminated. Those issues will be brought to the DRB in a timely fashion.

ADRM also notes that DSEP will be asked to assume a formal policy consultant role for setting the confidentiality protection parameters for the final 2018 End-to-End Test and the 2020 DAS. The charter for DSEP currently delegates the authority to set disclosure avoidance standards to the DRB, with review by DSEP if necessary. However, these parameters now must be public in a formal privacy system. Furthermore, they, like any other operational decision need to be

discussed and set in a manner consistent with their importance in the publication of results from the 2020 Census. The privacy-loss setting recommended by DRB and DSEP, and accepted by the Director, will be implemented in the production system.

Requests to DSEP:

Request 1: Concurrence with the DRB's decision on the PL 94-171 file items that can be published as enumerated.

In order to meet the timeline for the 2018 End-to-End Test, the version of the DAS under development for the test is limited in scope to the PL94-171 redistricting data. ADRM will not have time to experiment with a suite of potential implementations. And, in particular, ADRM will not have time to modify certain implementation decisions. They will be put back on the table for the full 2020 Disclosure Avoidance System and the decision on these six specific items may be revisited.

Request 2: Concurrence with Change to DRB Operating Principles Related to 2020 Census

The second request is for DSEP concurrence on a change in the operating principles of the DRB for issues related to disclosure avoidance in the 2020 Census of Population. Because the differentially private disclosure avoidance methods operate on the ensemble of proposed publications, DSEP is asked to concur that any disclosure avoidance request for publications from 2020 Census data be routed to the 2020 DAS team first. Those requests should not be considered by the DRB until the 2020 DAS team supplies a memo stating that the requested publication can or cannot be incorporated into the total privacy-loss accounting.

This is not a request for a moratorium on approvals for decennial data releases or design. The privacy-loss budget itself and its allocation to various components of the publication system are policy decisions that the 2020 Disclosure Avoidance System team will not make. Those decisions will ultimately be made in a manner consistent with the charters of the DRB and DSEP, and defended by the Director.

There is very little historical guidance for this process. We need to develop practical use cases that illustrate the consequences of publication decisions under alternative privacy-loss scenarios. We need to document the extent to which a best-effort reconstruction of the 2010 Hundred-percent Detail File (HDF) is correlated with the actual HDF. This is going to take some time. In the interim, ADRM is asking the DRB to take a leadership role in making these important choices by enabling the development of technologies better adapted to global risk management.

Discussion:

DSEP recognized the value in ADRM's efforts to assemble a skilled team of experts in an effort to modernize Census Bureau disclosure avoidance techniques using formal privacy methods.

This is essential in light of research that demonstrates that we must protect against database reconstructions that could lead to re-identification.

DSEP discussed the details of the six data elements from PL 94-171 and considered the necessity of including all of these in the proposed 2020 DAS research. ADRM requested that all elements remain available for the 2018 test research with a reconsideration for the full 2020 DAS, once the Census Bureau understand the outcomes. Conversations with the Department of Justice for Voting Rights Acts requirements with PL 94-171 will also play a part in future decisions about published enumerations.

DSEP recognized the need to develop ways to communicate with state stakeholders and the public about data protections that based on 2020 DAS methods. Our messaging will have to provide some simpler description of how the methods make changes to the attributes of the people in block counts, but still provide accurate and usable data.

DSEP noted that The National Conference of State Legislators (NCSL) will be expecting updates from Decennial based on 2018 testing outcomes in anticipation of 2020 releases of PL 94-171. It will be important to engage NCSL in discussions about 2020 DAS methods.

DSEP acknowledged that this and other details from ADRM's research were scheduled for discussion at the May 10, 2017 meeting of the 2020 Census Portfolio Management Governing Board (PMGB). DSEP postponed further discussion on this project and requests, pending any feedback from the presentation on this topic to the 2020 PMGB.

Post Meeting Notes:

DSEP revisited this topic at the beginning of the May 11, 2017 meeting.

Regarding issues of surrounding Voting Rights Acts Requirements, DSEP recognized that Decennial would need to talk to Justice if we were to alter any of the 6 constraints from PL 94-171 for 2020.

DSEP noted that the 2020 PMGB is supportive of the efforts of the 2020 DAS to optimize output noise infusion methods while publishing the most accurate data possible. There was unanimous support from 2020 PMGB for DRB's determination that the six data elements from PL 94-171 should be published as enumerated and form the base for the 2018 End-to-End testing research with the 2020 DAS.

DSEP agreed that the DRB should require that any request for disclosure avoidance of proposed publications for the 2020 Census be routed to the 2020 DAS team before going to the DRB.

Decision:

Request 1: DSEP approves publication of the six data elements from PL 94-171 as enumerated for the 2018 End-to-End test. Based on lessons learned, the use of these constraints for the PL 94-171 will be revisited for 2020.

Request 2: DSEP agreed that the DRB should require that any request for disclosure avoidance of proposed publications for the 2020 Census be routed to the 2020 DAS team before going to the DRB.

Record-level Re-identification Linkages for Evaluating the 2010 and 2020 Census Disclosure Avoidance Systems

Background:

The DAS team is attempting a database reconstruction using data from the 2010 PL94-171 and SF1 tabulations. The next step is to link those reconstructed microdata to commercial name and address files obtained in support of post-2010 research meant to represent the type of publically available file an attacker might potentially acquire. These files include Experian, InfoGroup, Melissa, Targus, TransUnion, and VSGL. This linkage involves the use of name and address data.

The final step is to compare the fully reconstructed microdata, including the commercially supplied names and address, to the name and address data on the 2010 Census Unedited File (CUF). Following accepted disclosure avoidance evaluation practices on re-identification, the 2020 DAS team would report to DRB and DSEP the putative re-identification rate (percentage of the records in the reconstructed microdata that could be linked to name and address information in the commercial files) and the proportion of putative re-identifications that were correct (proportion of reconstructed data records with putative re-identifications that were correctly linked to 2010 Census responses, including name and address).

Discussion:

DSEP recognized that the project proposal meets Data Linkage Policy requirements and involves sensitive but critical work that will allow the 2020 DAS subteam to understand the degree of risk of re-identification and database reconstruction with Census files.

DSEP noted that the subteam assembled for this research is composed of federal employees and one SSS individual.

Decision:

DSEP approved this project.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

DSEP Meeting Record

Topics: 2020 Decennial Record Linkage Test (Ned Porter, CSRM)



Meeting Date: May 11, 2017

<i>Position</i>	<i>Attending for Position</i>
Deputy Director (Chair)	Ron Jarmin
CAO	David Ziaya
CFO	Joanne Crane
AD, Decennial	Al Fontenot
AD, Demographic	Karen Battle
AD, Economic	Ron Jarmin
AD, Field	Joan Hill
AD, IT	Nitin Naik
AD, Research and Methodology	John Abowd
AD, 2020 Census	Al Fontenot
AD, Communications	Stephen Buckner
AD, Performance Improvement	Ted Johnson
Chief, PCO/ Chief Privacy Officer	Robin Bachman
Chief Demographer	Howard Hogan
Senior Advisor Designee from the Director's Office	<i>absent</i>
Chief Information Security Officer	Tim Ruland
Asst. Director, Research and Methodology	John Eltinge

Also Attending:	Simson Garfinkle, Tommy Wright, Eloise Parker, Ned Porter, Bill Winkler, Christa Jones, Letitia McKoy, Melissa Creech, Hampton Wilson, Ashley Landreth, Mike Castro, Janean Darden, Julie Atwell
-----------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Administrative Notes:

At the beginning of the meeting, DSEP resumed their discussion and made a final decision on the topic: *Initial Request for DSEP Determination on Disclosure Avoidance for the 2018 End-to-End Test of the 2020 Census of Population and Housing*. The summary of that discussion and decision is in the May 10, 2017 meeting record.

2020 Decennial Record Linkage Test

Background:

Identifying duplicate records in the decennial census is critical to providing a more accurate count. One of the areas of research for improving the Decennial Matching methodology is improving the computer matching in the Duplicate Person Identification (DPI) process. This research will use the 2010 Census Unedited File (CUF) as well as data from Census Coverage Measurement (CCM). In addition, the research will determine if it is possible to increase the proportion of records receiving Personal Identification Keys (PIKs).

This research requires access to PIKs and complete name data on the files. This access will be limited to only five Census Bureau researchers as well as the Center for Statistical Research and Methodology's Data Steward. The data will be restricted to only authorized clusters.

Discussion:

DSEP acknowledged that research into deduplication methods is a routine and critical part of Census operations. DSEP further acknowledged that while this research project will use new technology and methods, it is fundamentally the same as research that happened in previous censuses.

Decision:

DSEP approved the project.

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

DSEP Meeting Record

Topics:



Database Reconstruction Issue Mitigation (John Abowd, ADRM)

Meeting Date: February 15, 2018

<i>Position</i>	<i>Attending for Position</i>
COO (Chair)	Enrique Lamas
ADDC	Albert Fontenot
ADDP	Karen Battle
ADEP	Nick Orsini and Ron Jarmin
ADFO	Tim Olson
ADITCIO	Nitin Naik
ADRM	John Abowd
ADCOM	Stephen Buckner
CAO	David Ziaya
CFO	Joanne Crane
Asst. DRM	John Eltinge
Chief PCO/ Chief Privacy Officer	Robin Bachman
S.A. Director's Office	Douglas Clift
CISO	<i>Absent</i>
At-Large	Howard Hogan
At-Large	Frank Vitrano
Also Attending:	William Samples, David Waddington, Burton Reist, Victoria Velkoff, Robert Sienkiewicz, Jim Treat, Cynthia Hollingsworth, Clifford Jordan, Julia Naum, Jim Dinwiddie, Simson Garfinkel, Melissa Creech, Pat Cantwell, Byron Crenshaw, Hampton Wilson, Ashley Landreth, Mike Castro, Julie Atwell, Michael Snow

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

- [REDACTED]
- [REDACTED]

Database Reconstruction Issue Mitigation

Background

The Census Bureau’s Operating Committee (OPCOM), serving as the Enterprise Risk Review Board, elevated the enterprise risk of database reconstruction to an enterprise issue based on the results of a database reconstruction attack research effort the Census Bureau launched to understand that risk better. When an enterprise risk is elevated to an enterprise issue, the risk owner must implement an active mitigation plan to mitigate the risk. To that end, the Research and Methodology Directorate presented six recommendations to help manage the Census Bureau’s publication strategy in ways that will protect its databases from reconstruction attacks.

NOTE: presenters and DSEP recognized that implementing several of the recommendations will require decisions on budget and staffing resources and that those decisions would need to be handled by other bodies at the Census Bureau. DSEP confined its discussion to establishing policy in response to the recommendations.

The following 6 recommendations were presented to DSEP:

- 1. Suspension until September 30, 2019 of ad hoc releases of sub-state geography from any confidential source unless vetted differential privacy tools, or a DRB-approved noise-infusion alternative, have been used to produce the publication. This applies to all research projects whether they are external or internal. It does not apply to scheduled publications from sponsored survey clients for whom there is already an approved DRB protocol. Those clients should be put on notice for subsequent contracts. The complete list of approved exceptions, including sponsored survey products, is provided in 20180215b-External_Internal_Substate_Geography.xlsx. The suspension will be reviewed prior to September 30, 2019.**

NOTE: This suspension does not apply to state and national publications. It also does not apply to already scheduled publications from regular production activities. Program areas provided ADRM a list of those scheduled publications that should be exempted from the suspension. ADRM proposed ending those exemptions by September 30, 2019 even for those publications if they were not being produced using formally private systems by that point.

Discussion: DSEP recognized the need to modernize the Census Bureau's disclosure avoidance systems. DSEP acknowledged that by approving a list of exemptions they are agreeing to hold elevated levels of risk of database reconstruction associated with all of these data products. However, DSEP acknowledged the Census Bureau is obligated to provide the data the public needs for decision making and some of the release dates are required by law.

DSEP also acknowledged the need to set a target date for making these changes. While the ultimate goal is to make the publications of all of our programs formally private, that likely will not happen by September, 2019. However, in the meantime significantly improved noise infusion methods will be put in place to mitigate reconstruction risk.

DSEP members expressed concern that the list of already scheduled publications presented might be incomplete and asked for additional time for program areas to review the list and submit updates. DSEP agreed that the Center for Disclosure Avoidance Research (CDAR) should continue to accept submissions and finalize the list in advance of the next DSEP meeting. DSEP will formally approve the list at that point.

Decision: DSEP will finalize their approval of this recommendation at the March 15 DSEP meeting once the list of excepted publications has been finalized.

Action Items: Program areas will send updates on the table of exempted data releases to the Chief of CDAR by February 23. The Chief of CDAR will redistribute the combined list to all contributors by February 28. CDAR will finalize the list of approved exceptions for distribution before DSEP's meeting on March 15.

- 2. Suspension of all proposed tables in Summary File 1 and Summary File 2 for the 2020 Census at the block, block-group, tract, and county level except for the PL94-171 tables, as announced in Federal Register Notice 170824806–7806–01 (November 8, 2017, pp. 51805-6). To add a summary file table at any level of geography, racial/ethnic subpopulation other than OMB aggregate categories as specified in the 1997 standard (Federal Register October 30, 1997, pp. 58782-90), or group quarters type below the 2010 P42 seven categories, an affirmative case must be made for that table, use cases identified, and suitability for use standards developed. In addition, we recommend that the voting-age invariant in PL94-171 be removed, so that voting-age would be protected. DSEP will be asked to approve the SF1 and SF2 table specifications once they have cleared 2020 governance.**

NOTE: The PL94-17 tables from the 2018 End-to-End Census Test have been designed with a formally private system already and will be published, with the voting-age invariant, as planned.

Discussion: DSEP recognized that the SF1 and SF2 involved a very detailed set of tables that had been created to suit a wide set of data users. These tables were created, as a rule, to produce as much highly accurate data as possible within the existing disclosure avoidance framework. However, DSEP acknowledged that these data in many cases were accurate to a level that was not supported by the actual uses of those data, and such an approach is simply untenable in a formally private system.

DSEP acknowledged a fundamental need to take stock of what data the Census Bureau is required to publish, both by statute and the needs of our data users, and at what level of accuracy. This is not an activity that should be done by our Disclosure Review Board. Program areas have to make the case of what the data will be used for, and the actual minimum level of accuracy needed for those uses, so that CDAR and the DRB can build the system to allocate the privacy-loss budget according to those use cases.

A redesign of SF1 and SF2 based on formally articulated use cases will take a tremendous amount of effort but cannot be done in a vacuum. Program areas will have to reach out to data-user communities on developing the use cases for the needed data accuracy and levels of geography.

NOTE: DSEP discussed but tabled until later any decision on changing the voting-age invariant for the PL94-171 table produced as part of the 2020 Census.

Decision: DSEP approved this recommendation. For the 2020 Census, SF1 and SF2 will be rebuilt based on use cases.

Action Items: DCMD, POP, and ADDC divisions will work with the relevant program management governing board (PMGB) to establish a plan to execute this redesign.

3. Immediate review of all sub-state geography scheduled publications from the American Community Survey (ACS) to determine which ones can be delayed until there is a formally private publishing system for ACS.

Discussion: DSEP acknowledged that many of the ACS tables are already in production and that production needs to move forward. DSEP acknowledged that there are likely no publications currently suitable for delay, however they emphasized that ACSO needs to ensure that all exceptions are added to the list.

Decision: DSEP approved this recommendation.

Action Items: ACSO will verify that they have included all of the necessary publications on the list of exempted data releases.

4. Consideration of postponing ACS PUMS releases indefinitely.

NOTE: DSEP recognized that all of the publication systems and methods for the Census of Island Areas are identical to the ACS. DSEP emphasized that any changes made to the ACS should also reflect consideration of the needs of the Island Areas.

Discussion: DSEP acknowledged that while the threat of database reconstruction and reidentification attacks applies to all of the Census Bureau's data products, should the ACS data be subject to a reidentification attack, from a public perception standpoint, our continued publication of the ACS PUMS files would appear to be an egregious mistake.

However, DSEP also acknowledged that the ACS PUMS is a heavily used dataset for research and recognized that discontinuing this publication could generate a great deal of traffic for the FSRDCs. DSEP acknowledged that, before the Census Bureau restricts use the ACS PUMS to the FSRDCs, it needs to verify that they can handle the increased workload. Additionally, at present there are no FSRDCs that are readily accessible from the Island Areas.

DSEP recognized that immediate suspension of the ACS PUMS would cause a great deal of concern among data users and others. DSEP discussed the need to work on messaging around

any suspension and to brief the Department of Commerce before the Census Bureau implements the suspension.

Decision: DSEP deferred for one month any decisions to suspend release of the ACS PUMS pending further consideration of the ability of the FSRDC network to support increased demand, the impact on the data needs of the Island Areas, and development of a messaging plan.

Action Items: ADRM will prepare an assessment of the potential increased demand on the FSRDC network, and Decennial will prepare an assessment of the impact of suspending this publication on the Island Areas. ADCOM will work on a messaging plan.

5. Mandate for the 2022 Economic Censuses to use formally private publication systems for all tables.

Discussion: DSEP recognized that it is too late to begin creating a formally private system for data releases from the 2017 Economic Census. DSEP additionally discussed how modernizing disclosure avoidance systems will involve much more than just budgeting extra funds. It also will require having the adequate number of people with the right skills to do the work.

DSEP recognized that program areas will have to involve their PMGB in setting resources, budgets, and timelines and that it should be feasible to put formally private systems in place in time for the 2022 Economic Census.

Decision: DSEP approved this recommendation. The Census Bureau will move forward with designing and implementing formally private systems for the 2022 Economic Census.

6. Mandate to the Demographics Directorate to begin negotiations with survey clients for increased use of restricted-access microdata protocols and formally private table publication systems.

POST MEETING NOTE: a member in attendance recommended that there should also be outreach to reimbursable clients for the Economic Directorate.

Discussion: DSEP recognized the need to begin discussions with sponsors of Census Bureau surveys but determined that the Census Bureau should have a communications plan in place before mandating that the Demographic Directorate speak to sponsors.

Decision: DSEP will reconsider in one month whether to mandate conversations with survey and report sponsors.

Consolidated Action items:

- Program areas will send updates on the table of exempted data releases to the Chief of CDAR by February 23.
- The Chief of CDAR will redistribute the combined list to all contributors by February 28.
- DCMD, POP, and the ADDC will work with the relevant PMGBs to establish a plan to execute the redesign of SF1 and SF2 based on use cases.
- ACSO will work to determine that all ACS data releases in production are listed on the spreadsheet of exceptions to the suspension.
- ADRM will prepare an assessment of the potential increased demand on the FSRDC network from suspension of the ACS PUMS.
- ADCOM will work on a messaging plan related to the suspension of the ACS PUMS.
- Decennial will prepare an assessment of the impact of suspending publication of the ACS PUMS on the Island Areas.

Staring Down the Database Reconstruction Theorem

John M. Abowd

Chief Scientist and Associate Director for Research and Methodology
U.S. Census Bureau

American Association for the Advancement of Science
Annual Meeting Saturday, February 16, 2019 3:30-5:00



U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov


The views expressed in this talk are my own
and not those of the U.S. Census Bureau.

The challenges of a census:

1. collect all of the data necessary to underpin our democracy;
2. protect the privacy of individual data to ensure trust and prevent abuse.

- Too many statistics
- Noise infusion is necessary
- Transparency about methods helps rather than harms

Fundamental Tradeoff between Accuracy and Privacy Loss



Good science and privacy protection are partners

OnTheMap

[LEHD Home Help and Documentation Reload Text-Only](#)

Start Base Map Selection Results

Distance/Direction Analysis

Work to Home

▼ Display Settings

Labor Market Segment: **All Workers**

Filter:

Year: **2015**

▼ Map Controls

Color Key: Blue

Thermal Overlay:

Point Overlay:

Selection Outline:

Identify: Zoom to Selection:

Clear Overlays: Animate Overlays:

▼ Report/Map Outputs


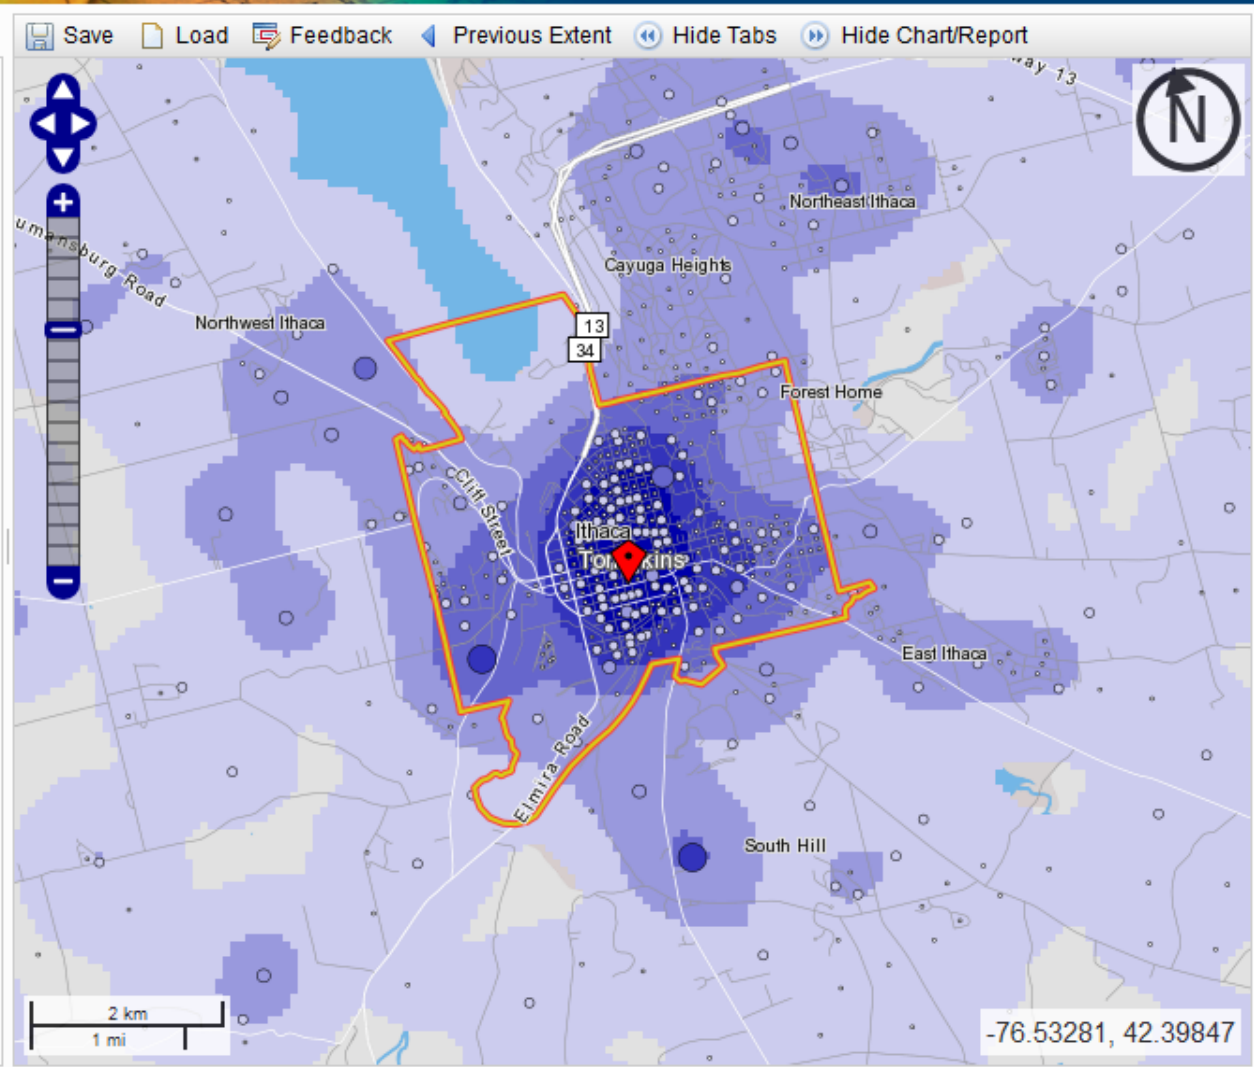
Detailed Report:

Export Geography:

Print Chart/Map:

▼ Legends

[Change Settings](#)



Jobs by Distance - Work Census Block to Home Census Block

	2015	
	Count	Share
Total Primary Jobs	12,260	100.0%
Less than 10 miles	5,949	48.5%
10 to 24 miles	2,987	24.4%
25 to 50 miles	1,451	11.8%
Greater than 50 miles	1,873	15.3%

What we did

- Database reconstruction for all 308,745,538 people in 2010 Census
- Link reconstructed records to commercial databases: acquire PII
- Successful linkage to commercial data: putative re-identification
- Compare putative re-identifications to confidential data
- Successful linkage to confidential data: confirmed re-identification
- Harm: attacker can learn self-response race and ethnicity

What we found

- Census block correctly reconstructed in all 6,207,027 inhabited blocks
- Block, sex, age, race, ethnicity reconstructed
 - Exactly: 46% of population (142 million of 308,745,538)
 - Allowing age +/- one year: 71% of population (219 million of 308,745,538)
- Block, sex, age linked to commercial data to acquire PII
 - Putative re-identifications: 45% of population (138 million of 308,745,538)
- Name, block, sex, age, race, ethnicity compared to confidential data
 - Confirmed re-identifications: 38% of putative (52 million; 17% of population)
- For the confirmed re-identifications, race and ethnicity are learned exactly, not statistically

We fixed this for the 2020 Census by implementing differential privacy

Acknowledgments

- The Census Bureau's 2020 Disclosure Avoidance System incorporates work by Daniel Kifer (Scientific Lead), Simson Garfinkel (Senior Scientist for Confidentiality and Data Access), Rob Sienkiewicz (ACC Disclosure Avoidance, Center for Enterprise Dissemination), Tamara Adams, Robert Ashmead, Michael Bentley, Stephen Clark, Craig Corl, Aref Dajani, Nathan Goldschlag, Michael Hay, Cynthia Hollingsworth, Michael Ikeda, Philip Leclerc, Ashwin Machanavajjhala, Christian Martindale, Gerome Miklau, Brett Moran, Edward Porter, Sarah Powazek, Anne Ross, Ian Schmutte, William Sexton, Lars Vilhuber, Cecil Washington, and Pavel Zhuralev

Thank you.

John.Maron.Abowd@census.gov



U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

More Background on the 2020 Census Disclosure Avoidance System

- September 14, 2017 CSAC (overall design)
<https://www2.census.gov/cac/sac/meetings/2017-09/garfinkel-modernizing-disclosure-avoidance.pdf?#>
- August, 2018 KDD'18 (top-down v. block-by-block)
<https://digitalcommons.ilr.cornell.edu/ldi/49/>
- October, 2018 WPES (implementation issues)
<https://arxiv.org/abs/1809.02201>
- October, 2018 *ACMQueue* (understanding database reconstruction)
<https://digitalcommons.ilr.cornell.edu/ldi/50/> or
<https://queue.acm.org/detail.cfm?id=3295691>
- December 6, 2010 CSAC (detailed discussion of algorithms and choices)
<https://www2.census.gov/cac/sac/meetings/2018-12/abowd-disclosure-avoidance.pdf?#>

Selected References

- Dinur, Irit and Kobbi Nissim. 2003. Revealing information while preserving privacy. In Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS '03). ACM, New York, NY, USA, 202-210. DOI: 10.1145/773153.773173.
- Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. In Halevi, S. & Rabin, T. (Eds.) Calibrating Noise to Sensitivity in Private Data Analysis Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings, Springer Berlin Heidelberg, 265-284, DOI: 10.1007/11681878_14.
- Fellegi, Ivan P. 1972. On the Question of Statistical Confidentiality. Journal of the American Statistical Association, Vol. 67, No. 337 (March):7-18, stable URL <http://www.jstor.org/stable/2284695>.
- Ganda, Srivatsava, Shiva Kasiviswanathan and Adam Smith. 2008. Composition Attacks and Auxiliary Information in Data Privacy. In Knowledge, Discovery and Datamining, Las Vegas, NV, doi:10.1145/1401890.1401926.
- Machanavajjhala, Ashwin, Daniel Kifer, John M. Abowd, Johannes Gehrke, and Lars Vilhuber. 2008. Privacy: Theory Meets Practice on the Map, International Conference on Data Engineering (ICDE) 2008: 277-286, doi:10.1109/ICDE.2008.4497436.
- McKenna, Laura. 2018. Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing, Working Papers 18-47, Center for Economic Studies, U.S. Census Bureau, Handle: RePEc:cen:wpaper:18-47.
- Ramachandran, Aditi, Lisa Singh, Edward Porter, and Frank Nagle. 2012. Exploring Re-Identification Risks in Public Domains, Tenth Annual International Conference on Privacy, Security and Trust, IEEE, doi:10.1109/PST.2012.6297917.
- U.S. Census Bureau. 2019. LEHD Origin-Destination Employment Statistics (2002-2015) [computer file]. Washington, DC: U.S. Census Bureau, Longitudinal-Employer Household Dynamics Program [distributor], accessed on February 15, 2019 at <https://onthemap.ces.census.gov>.

[Slide 1] [Before I start, I want to remind members of the audience that, while I am appearing in my official capacity as the Chief Scientist of the U.S. Census Bureau, I am presenting a summary of research findings. The views expressed in this talk are my own, not those of the Census Bureau.]

Staring Down the Database Reconstruction Theorem

[Slide 2] The 2020 Census will be the safest and best-protected ever. This is not nearly as easy as it sounds.

Throughout much of the history of the decennial census, our country has struggled with two challenges:

- 1) collect all of the data necessary to underpin our democracy;
- 2) protect the privacy of individual data to ensure trust and prevent abuse.

The first obligation derives directly from the Constitution, of course. As for the privacy requirement, Section 9 of the Census Act (Title 13 of the U.S. Code) prohibits making “any publication whereby the data furnished by any particular establishment or individual under this title can be identified.” In fact, the Census Bureau is about the only organization operating under a blanket U.S. legal requirement never to release data that can be tied back to individuals or companies no matter what.

The Census Bureau has always been committed to meeting both of its obligations; that is, providing population statistics needed by decision-makers, scholars, and businesses while also protecting the privacy of census participants.

A paper by Laura McKenna (2018), who supervised the confidentiality protection systems used by the Census Bureau for more than 15 years, catalogued the public information about the technical systems used for protection of publications from decennial censuses since 1970.

As McKenna noted, beginning with the 1990 Census, the primary confidentiality protection method employed was household-level swapping of geographic identifiers—moving an entire household from one location to another—prior to tabulating the data. The goal was to introduce uncertainty about whether households allegedly re-identified from the published data were correct.

Essentially the same methods were used for the 2000 and 2010 Censuses but with refinements that recognized the changing external environment.

The discipline of statistics has evolved over the last century. So too has the widespread availability of data. With each new development, the Census Bureau must ask how the current state of affairs will affect the production of the statistical products that it releases to the public so as to be both useful and privacy-preserving.

Sixteen years ago, two computer scientists, Irit Dinur and Kobbi Nissim (2003), wrote a seminal article proving a “database reconstruction theorem,” which is also known as the “fundamental law of information recovery.”

Three years later, Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith (2006) provided a mathematical foundation for what we now call “differential privacy.” In short, they explained how to quantify the limits on the accuracy of answers to queries based on the confidential data and the privacy-loss to the entities in those data, when the queries are answered publicly. More importantly, they provided a technique for enhancing privacy that goes far beyond the swapping approach that many statisticians have been using for years.

[Slide 3] The full implications of database reconstruction were not understood in 2003, but over the next several years a scientific consensus emerged in the data privacy community that:

- **Too many statistics**, published too accurately, expose the confidential database with near certainty (Dinur and Nissim 2003).
- **A necessary condition for controlling privacy loss** against informed attackers is to add noise to every statistic, calibrated to control the worst-case disclosure risk, which is now called a privacy-loss budget (Dwork, McSherry, Nissim and Smith 2006; Ganta, Kasiviswanathan, and Smith 2008).
- **Transparency about methods helps rather than harms**, Kerckhoff’s principle, applied to data privacy, says that the protections should be provable and secure even when every aspect of the algorithm and all of its parameters are public. Only the actual random number sequence must be kept secret (Dwork, McSherry, Nissim, and Smith 2006).

If you curate confidential data, then you can use those data for two competing goals:

- You can publicly and precisely answer statistical queries about the data.
- You can preserve and protect the privacy of those whose information is in the data.

You can do some of both.

[Slide 4] But if you do all of one, you can't do any of the other.

Period.

This trade-off is one of the hardest lessons to learn in modern information science. It is a lesson about data generally, not about counting people. And it is a mathematical theorem, not an opinion or implementation detail.

[Slide 5] This transformation in the fields of statistics and computer science is truly mind-blowing. It's at the heart of the science that we're here to celebrate. Cryptographers usually study the safety of methods for encrypting information about private data. Now their insights show us safe ways to publish information from private data. The cryptographic approach shows that some new methods can provably protect privacy, and some old methods provably do not. But the safe methods only work if we accept the inherent limitations on the accuracy of those publications that the cryptographers have highlighted.

Specifically, technical advances revealed a new vulnerability, allowing people to reconstruct data from tables that were previously assumed to be privacy-preserving, given the available computing resources. But other technical advances have also enabled a new form of privacy protection that is not only more sophisticated but also mathematically grounded in a way that allows statisticians to fully understand the limits of what they can make available and what kind of privacy they can provably offer. This dual breakthrough is transforming how we protect data today.

Good science and real privacy protection turn out to be partners, not competitors, in the efforts to modernize the methods data analysts use. For this reason, we have seen many companies, like Google, Microsoft, and Apple, turn to differential privacy to secure data and make guarantees about the privacy of

statistical tables. But it was actually the Census Bureau who first recognized the power of this method at scale.

[Slide 6] In 2008, the Census Bureau implemented an early version of differential privacy on data that display the commuting patterns of people based on where they live and work (Machanavajjhala et al. 2008; U.S. Census Bureau 2019).

Working with statisticians and computer scientists, we have collectively advanced the state of differential privacy such that we are going to implement it at scale as part of the 2020 Census. While I will talk about what that looks like in more detail tomorrow at 8:00AM, today I want to explain why we absolutely must implement differential privacy in order to protect the privacy of those participating in the census.

Starting in 1972, researchers began highlighting how it was possible to combine statistical tables and use differencing techniques to identify which census respondents provided the associated data (Fellegi 1972). As the market for detailed data grew and evolved, researchers also began highlighting how combining commercial data with census tables could introduce new vulnerabilities. While external users could not provably know whether or not their reconstructions were accurate, the Census Bureau recognized that it was critical to know the potential vulnerability of census data.

We acted proactively, as the Census Bureau has done for many decades. We designed our own internal research program to assess the current state of this vulnerability without waiting for a specific external threat. I'm now going to explain what we found.

[Slide 7] Here are the steps we followed:

- Using only published contingency tables (summary statistics), we applied the database reconstruction theorem to construct record-level images for all 308,745,538 persons enumerated in the 2010 Census. A record-level image is a row in the reconstructed database with the same variables that were used in publications from the confidential database. There is no traditional PII (personally identifiable information) on these reconstructed records.

- Using only the information in the reconstructed data records, we linked those records to commercial databases to acquire name and address information. This information would have been available to an external attacker, circa 2010.
- When the record linkage operation is successful, the PII from the commercial data are attached to the reconstructed census record. We call the reconstructed record, now laden with PII, “putatively re-identified,” which means that an attacker might think that the attack was successful.
- We then compared the putatively re-identified census records to the real confidential census records. When this comparison matched on all variables, including the PII and those variables not available in the commercial data, we called this a “confirmed re-identification.”
- The harm from such re-identifications, in the 2010 Census, is that the attacker learns the self-reported race and ethnicity on the confidential census record. Those data are not available in identifiable form to any commercial or governmental agency except the Census Bureau.

[Slide 8] Here are the basic results:

- In the reconstructed data, certain variables are always correctly reconstructed—meaning that the value in the reconstructed variable always matches its value in the confidential data. The census block, where the person lived on April 1, 2010, is always correctly reconstructed. This is true for every one of the 6,207,027 inhabited blocks in the 2010 Census.
- All the variables we studied: block, sex, age in years, race, and ethnicity are exactly correct in the reconstructed records for 46% of the population (142 million of 308,745,538 persons)—meaning that the reconstructed record exactly matches the confidential record on the value of all five variables. This result is salient because in the confidential data, more than 50% of the records are unique in the population—the only instance of this combination of values observed in the census (the exact percentage is confidential). If we allow the age to vary by plus or minus one year, then the number of reconstructed records that match the confidential data on these five variables rises to 71% (219 million of 308,745,538 persons).
- When we use the reconstructed block, sex and age to link each reconstructed record to the records harvested from commercial data

acquired at the time of the 2010 Census, we putatively re-identify 45% of the total population (138 million of 308,745,538 persons). That means that we were able to attach a unique name and address to 45% of the reconstructed records from the 2010 Census. The match is exact for block and sex. Age is allowed to vary by plus or minus one year.

- When we compared the unique name, block, sex, age, race, and ethnicity on the putative re-identifications to the same variables on the 2010 Census confidential data, we confirmed 38% of these matches (52 million of 308,745,538 persons, or 17% of the total population).

The putative re-identifications probably have a recall rate (or sensitivity) of at least 45%. Neither the attacker nor the Census Bureau have PII on all 308,745,538 persons enumerated in the 2010 Census, so the correct recall rate denominator is certainly less than the total population.

The precision of the record linkage is 38%, which means that the attacker would be correct between one-quarter and one-half of the time.

And both of these estimates (45% putatively re-identified; 38% of which are correct) are really lower bounds for other reasons: our experiments didn't use all of the information that the Census Bureau published from the 2010 Census. For example, we didn't use any information on household composition, which means that potential harm from discovering other features of households, like same-sex unions and adoptions, is still unquantified. We also made no use of the 2010 Public-Use Microdata Sample.

To further put these results in context, the last time the Census Bureau released results for a re-identification study, which did not use database reconstruction (Ramachandran et al. 2012), the putative re-identification rate was 0.017% (389 persons of 2,251,571) and the confirmation rate was 22% (87 of 389).

[Slide 9] All of us—the entire scientific community—have an obligation to examine the methods we use in light of the cryptographic critique of the privacy protections those methods offer. We must also recognize that these developments are sobering to everyone.

This is not just a challenge for statistical agencies or Internet giants, although those institutions have been in the vanguard of this movement.

It's a challenge for Internet commerce, because recommendation systems expose private data.

It's a challenge for bioinformatics, because summaries of genomes expose private data.

It's a challenge for commercial lenders, because benchmark risk assessments expose private data.

It's a challenge for nonprofit survey organizations, because their research reports expose private data.

Regardless of what anyone says, people want to be assured that their data are private. They want to know that we can't use statistical magic to re-identify information that they thought was private. They want to know that statistical tables can't come back to haunt them.

That's why I'm so grateful that the data we are showing today aren't the end of the story. They simply show that we cannot accept the status quo. We cannot presume that what worked a decade ago will work again in 2020. We have to innovate. And that's what we are doing.

In 2016, the Census Bureau acknowledged that database reconstruction was a vulnerability of the methods traditionally used to protect confidentiality in decennial census publications.

What we showed today is that we have a clear understanding of how it's possible to reconstruct 2010 Census data for block, sex, age, race and ethnicity. But this understanding isn't in vain. This understanding gave us the information we needed to develop techniques to make sure this isn't possible in 2020.

We are going into the 2020 Census confident that we can protect the privacy of all who participate. We have to make some important decisions about what statistics should be made available and how to weigh public data interests with our commitment to keep individual data private from reconstruction. But we know where the vulnerabilities are and we have the tools to make certain that what I showed today can't happen in the future.

The publications of the 2020 Census will be protected by differential privacy because it's imperative above all else that we ensure the trust of the American people.

The exact algorithms, and all parameters, will also be publicly released well in advance of the tables because it is imperative that we be accountable to the scientific community and the public at large.

[Slide 10] Statistics has evolved significantly over the last century. I'm honored to be a part of a statistical agency with a long tradition of implementing cutting-edge knowledge on the behalf of the American people. And I'm deeply grateful to the amazing team at the Census Bureau for identifying the challenges we face and ensuring that we can meet those challenges.

I promise the American people that they will have the privacy they deserve.

For those who would like to know more about how we are implementing differential privacy in the 2020 Census, please join me tomorrow at 8:00 AM where I will present our methods in more detail.

References

- Dinur, Irit and Kobbi Nissim. 2003. Revealing information while preserving privacy. In Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS '03). ACM, New York, NY, USA, 202-210. DOI: 10.1145/773153.773173.
- Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. In Halevi, S. & Rabin, T. (Eds.) Calibrating Noise to Sensitivity in Private Data Analysis Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings, Springer Berlin Heidelberg, 265-284, DOI: 10.1007/11681878_14.
- Fellegi, Ivan P. 1972. On the Question of Statistical Confidentiality. *Journal of the American Statistical Association*, Vol. 67, No. 337 (March):7-18, stable URL <http://www.jstor.org/stable/2284695>.
- Ganda, Srivatsava, Shiva Kasiviswanathan and Adam Smith. 2008. Composition Attacks and Auxiliary Information in Data Privacy. In *Knowledge, Discovery and Datamining*, Las Vegas, NV, doi:10.1145/1401890.1401926.
- Machanavajhala, Ashwin, Daniel Kifer, John M. Abowd, Johannes Gehrke, and Lars Vilhuber. 2008. Privacy: Theory Meets Practice on the Map, International Conference on Data Engineering (ICDE) 2008: 277-286, doi:10.1109/ICDE.2008.4497436.
- McKenna, Laura. 2018. Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing, Working Papers 18-47, Center for Economic Studies, U.S. Census Bureau, Handle: RePEc:cen:wpaper:18-47.

Abowd, AAAS presentation Saturday, February 16, 2019, 3:30-5:00

FINAL Page 9

Ramachandran, Aditi, Lisa Singh, Edward Porter, and Frank Nagle. 2012. Exploring Re-Identification Risks in Public Domains, Tenth Annual International Conference on Privacy, Security and Trust, IEEE, doi:10.1109/PST.2012.6297917.

U.S. Census Bureau. 2019. LEHD Origin-Destination Employment Statistics (2002-2015) [computer file]. Washington, DC: U.S. Census Bureau, Longitudinal-Employer Household Dynamics Program [distributor], accessed on February 15, 2019 at <https://onthemap.ces.census.gov>.