

Assessing the Use of Differential Privacy for the 2020 Census: Summary of What We Learned from the CNSTAT Workshop

Feb 25, 2020

V. Joseph Hotz, Duke University
Joseph Salvo, New York City Department of City Planning

Background

The mission of the Census Bureau is to provide data that can be used to draw a picture of the nation, from the smallest towns and villages to the neighborhoods of the largest cities. Advances in computer science, better record linkage technology, and the proliferation of large public data sets have increased the risk of disclosing information about individuals in the census.

To assess these threats, the Census Bureau conducted a simulated attack, reconstructing person-level records from published 2010 Census tabulations using its previous Disclosure Avoidance System (DAS) that was based in large part on swapping data records across households and localities. When combined with information in commercial and publicly available databases, these reconstructed data suggested that 18 percent of the U.S. population could be identified with a high level of certainty. The Census Bureau concluded that, if adopted for 2020, the 2010 confidentiality measures would lead to a high risk of disclosing individual responses violating [Title 13 of the U.S. Code](#), the law that prohibits such disclosures.

Thus, the Census Bureau was compelled to devise new methods to protect individual responses from disclosure. Nonetheless, such efforts – however well-intentioned – may pose a threat to the content, quality and usefulness of the very data that defines the Census Bureau’s mission and that demographers and statisticians rely on to draw a portrait of the nation’s communities.

The [Census Bureau’s solution](#) to protecting privacy is a new DAS based on a methodology referred to as [Differential Privacy](#) (DP). In brief, it functions by leveraging the same database reconstruction techniques that were used to diagnose the problem in the previous system: the 2020 DAS synthesizes a complete set of person- and household-level data records based on an extensive set of tabulations to which statistical noise has been added. Viewed as a continuum between total noise and total disclosure, the core of this method involves a determination regarding the amount of *privacy loss* or ϵ , that can be accepted without compromising data privacy while ensuring the utility of the data. The key then becomes “where to set the dial” – set ϵ too low and privacy is ensured at the cost of utility, but set ϵ too high and utility is ensured but privacy is compromised. In addition to the overall level of ϵ , its allocation over the content and detail of the census tabulations for 2020 is important. For example, specific block-level tabulations needed for redistricting may require a substantial allocation of the privacy-loss budget to achieve acceptable accuracy for this key use, but the cost is that accuracy of other important data (including for blocks, such as persons per household) will likely be

compromised. Finding ways to resolve these difficult tradeoffs represents a serious challenge for the Census Bureau and users of its data.

The CNSTAT Workshop

In order to test how well this methodology worked in terms of the accuracy of noise-infused data, the [Census Bureau issued special 2010 Census files subject to the 2020 DAS](#). The demonstration files applied the 2020 Census DAS to the 2010 Census confidential data — that is, the unprotected data from the 2010 Census that are not publicly available. The demonstration data permit scientific inquiry into the impact of DP. In addition, the Census commissioned the Committee on National Statistics (CNSTAT) of the National Academies of Sciences, Engineering and Medicine to host a 2-day Workshop on [2020 Census Data Products: Data Needs and Privacy Considerations](#), held in Washington, DC, on December 11-12, 2019. The two-fold purpose of the workshop was:

- To assess the *utility* of the tabulations in the 2010 Demonstration Product for specific use cases/real-life data applications.
- Generate constructive feedback for the Census Bureau that will be useful in setting the ultimate privacy loss budget *and* on the allocation of shares of that budget over the broad array of possible tables and geographic levels.

We both served as the co-chairs of the Committee that planned the Workshop. The Workshop brought together a diverse group of researchers who presented findings for a wide range of use cases that relied on data from past censuses.

These presentations, and the discussions surrounding them, provided a new set of evidence-based findings on the potential consequences of the Census Bureau's new DAS. In what follows, we summarize "what we heard" or learned from the Workshop. This summary is ours alone; we do not speak for the Workshop's Planning Committee, CNSTAT, or the Census Bureau. Nonetheless, we hope that the summary below provides the broader community of users of decennial census data with a better understanding of some of the potential consequences of the new DAS for the utility of the 2020 Census data products. Moreover we hope it fosters an on-going dialogue between the user community and the Census Bureau on ways to help ensure that data from the 2020 Census are of high quality, while still safeguarding the privacy and confidentiality of individual responses.

What We Heard

- **Population counts for some geographic units and demographic characteristics were not adversely affected by Differential Privacy (DP)**. Based on results presented at the Workshop, it appears that there were not, in general, differences in population counts between the 2010 demonstration file at some levels of geography. For the nation as a whole and for individual states, the Census's algorithm, ensured that that counts were exact, i.e., counts at these levels were held invariant by design. Furthermore, the evidence presented also indicated that the counts in the demonstration products and those for actual

2010 data were not very different for geographic areas that received direct allocations of the privacy budget, including most counties, metro areas (aggregates of counties) and census tracts. Finally, for these geographic areas, the population counts by age in the demonstration products were fairly accurate when using broader age groupings (5-10 year groupings or broader ones), as well as for some demographic characteristics (e.g., for non-Hispanic whites, and sometimes for Hispanics).

- **Concerns with data for small geographic areas and units and certain population groups.** At the same time, evidence presented at the Workshop indicated that most data for small geographic areas – especially census blocks – are not usable given the privacy-loss level used to produce the demonstration file. With some exceptions, applications demonstrated that the variability of small-area data (i.e., blocks, block groups, census tracts) compromised existing analyses. Many Workshop participants indicated that a larger privacy loss budget will be needed for the 2020 Census products to attain a minimum threshold of utility for small-area data. Alternatively, compromises in the content of the publicly-released products will be required to ensure greater accuracy for small areas.

The Census did not include a direct allocation of the privacy-loss budget 2010 demonstration file to all geographic areas, such as places and county subdivisions, or to detailed race groups, such as American Indians. As noted by numerous presenters, these units and groups are very important for many use cases, as they are the basis for political, legal, and administrative decision-making. Many of these cases involve small populations and local officials rely on the census as a key benchmark; in many cases, it defines who they are.

- **Problems for temporal consistency of population counts.** Several presentations highlighted the problem of temporal inconsistency of counts, i.e., from one census to the next using DP. The analyses presented at the Workshop suggested that comparisons of 2010 Census data under the old DAS to 2020 Census data under DP may well show inexplicable trends, up or down, for small geographic areas and population groups. (And comparisons of 2030 data under DP with 2020 data under DP may also show inconsistencies over time). For example, when using counts as denominators to monitor disease rates or mortality at finer levels of geography by race, by old vs young, etc., the concern is that it will be difficult to determine real changes in population counts, and, thus, real trends in disease or mortality rates, versus the impact of using DP.
- **Unexpected issues with the post-processing of the proposed DAS.** The Top-Down algorithm (TDA) employed by the Census Bureau in constructing the 2010 demonstration data produced histograms at different levels of geography that are, by design, unbiased — but they are not integers and include negative counts. The post-processing required to produce a microdata file capable of generating tabulations of persons and housing units with non-negative integer counts produced biases that are responsible for many anomalies observed in the tabulations. These are both systematic and problematic for many use cases. Additional complications arise from the need to hold some data cells invariant to change (e.g., total population at the state level) and from the separate processing of person and housing unit tabulations.

The application of DP to raw census data (the Census Edited File [CEF]) produces estimates that can be used to model error, but the post-processing adds a layer of complexity that may be very difficult to model, making the creation of “confidence intervals” problematic.

- **Implications for other Census Bureau data products.** Important parts of the planned 2020 Census data products cannot be handled by the current 2020 DAS and TDA approach. They will be handled using different but as-yet-unspecified methods that will need to be consistent with the global privacy-loss budget for the 2020 Census. These products were not included in the demonstration files and were out of scope for the Workshop. Nonetheless, as noted by several presenters and participants in the Workshop, these decisions raise important issues for many users and use cases going forward. To what extent will content for detailed race/Hispanic/nationality groups be available, especially for American Indian and Alaska Native populations? To what degree will data on household-person combinations and within-household composition be possible under DAS?

For example, while the Census Bureau has stated that 2025 will be the target date for the possible application of DP to the ACS, they indicated that the population estimates program will be subject to DP immediately following 2020. These estimates would then be used for weighting and post-stratification adjustments to the ACS.

- **Need plan to educate and provide guidance for users of the 2020 Census products.** Regardless of what the Census Bureau decides with respect to ϵ and how it is allocated across tables, the Workshop participants made clear that a major re-education plan for data users’ needs to be put in place, with a focus on how best to describe key data and the shortcomings imposed by privacy considerations and error in general. Furthermore, as many at the Workshop voiced, such plans must be in place when the 2020 Census products are released to minimize major disruptions to and problems with the myriad uses made of these data and the decisions based on them.
- **Challenging privacy concerns and their potential consequences for the success of the 2020 Census.** Finally, the Workshop included a panel of experts on privacy. These experts highlighted the disclosure risks associated with advances in linking information in public data sources, like the decennial census, with commercial data bases containing information on bankruptcies and credit card debt, driver licenses, and federal, state and local government databases on criminal offenses, public housing, and even citizenship status. While there are federal and state laws in place to protect the misuse of these governmental databases as well as the census (i.e., Title 13), their adequacy is challenged by advances in data linkage technologies and algorithms. And, as several panelists noted, these potential disclosure risks may well undercut the willingness of members of various groups – including immigrants (whether citizens or not), individuals violating public housing codes, or those at risk of domestic violence – to participate in the 2020 Census.

The Census Bureau has [recently stated](#) that it plans to have CNSTAT organize a follow-up set of expert meetings to “document improvements and overcome remaining challenges in the 2020 DAS.” In our view, such efforts, however they are organized, need to ensure meaningful involvement and feedback from the user community. Many within that community remain skeptical of the Bureau’s adoption of Differential Privacy and its consequences for their use

cases. So, not only is it important that Census try to address the various problems identified by Workshop presenters and others who evaluated the 2010 demonstration products, it also is essential that follow-up activities are designed to involve a broader base of user communities in a meaningful way.

We encourage members of the census data user community to become engaged in this evaluation process, agreeing, if asked, to become involved in these follow-up efforts. Such efforts will be essential to help ensure that the Census Bureau meets its dual mandate of being the nation's leading provider of quality information about its people and economy while safeguarding the privacy of those who provide this information.