Comparing Differential Privacy With Older Disclosure Avoidance Methods

The U.S. Census Bureau's methods to protect your responses in published census data have evolved steadily over the decades. In the past century, we moved from a system that relied on "eyeballing" data tables to spot potentially revealing statistics to a system of intricate, statistical techniques to address growing disclosure risks. But the methods used in 2010 and earlier, available at <https://www2.census.gov/about/partners/cac/sac/meetings/2021-05/presentation-research-on-alternatives-to-differential-privacy.pdf>, are no match for the re-identification threats in this era of Big Data and limitless computing power.

About 57 percent of the 2010 Census population were "unique" at the smallest census geography, block level, meaning they were the only people in their block with a specific combination of sex, age (in years), race (any of the 63 possible Office of Management and Budget race combinations), and Hispanic/Latino ethnicity.[1] Those kinds of unique attributes are precisely the vulnerabilities discoverable by today's technology.

The decision to adopt confidentiality protections based on differential privacy for the 2020 Census was based on research that exposed the limits of our previous methods. We conducted experiments to better understand how those techniques would impact census results if applied today, using published 2010 Census results as the basis.

The findings offer stakeholders a tool for comparing the trade-offs between those earlier methods and the new approach designed for application to the P.L. 94-171 redistricting data, the TopDown Algorithm, which is based on the principles of differential privacy.

These findings are summarized below and are also available via a webinar we conducted in June 2021 at <www.census.gov/data/academy/webinars/2021/disclosure-avoidance-series/research-into-alternatives-to-differential-privacy.html>.

DIFFERENTIAL PRIVACY

What Is Differential Privacy and How Does It Work?

Differential privacy, first developed in 2006, is a framework for measuring the precise disclosure risk associated with each release of confidential data. It allows an agency like the Census Bureau to quantify the precise amount of statistical noise required to protect confidentiality. This precision allows us to calibrate and allocate precise amounts of statistical noise in a way that protects confidentiality while maintaining the overall accuracy of the data in the aggregate.

The amount of randomly generated noise that is injected is driven by a tunable, or adjustable, "privacy-loss budget." An algorithm, that is also tunable, determines how much of that noise is injected into individual results and geographies. It is important to note that, since publishing exact counts of people and housing units at low levels of geography is the key to re-identifying the people behind the statistics, the new disclosure avoidance system limits the kinds of statistics that are published as counted. These are called invariants. The 2020 Census publishes exact counts for the total population at the state level, the number and type of occupied group quarters facilities at the block level, and the number of housing units, whether occupied or not, at the block level.

[1] Approved for public release per Disclosure Review Board clearance number: CBDRB-FY21-DSEP-003.

D-FS-GP-EN-0509
August 2021

United States® Census 2020

However, differential privacy injects noise after queries are tabulated directly from as-counted results. The algorithm injects noise to the number that results from each query (e.g., # of Asian Males x Non-Hispanic x Voting Age x Specific Block), centered around a mean of zero noise. This model allows us to fine-tune the injection of noise to reduce distortions and adjust the amount of noise based on use cases our data users provide.

There is no "off-the-shelf" or standard application for meeting the requirements and standards of differential privacy. Each application must be built for and tailored to each specific data set and purpose. The differentially private application for the P.L. 94-171 redistricting data and the Demographic and Housing Characteristics file is called the "TopDown Algorithm." While both of those products will use the TopDown Algorithm, each is tuned separately.

What Impact Does Differential Privacy Have on Confidentiality, Accuracy, and Data Availability?

- **Confidentiality.** Unlike older protection methods, differential privacy provides mathematically provable measures of protection. The privacy-loss budget can range from a value of "zero" (offering zero accuracy) to a value of infinity (offering zero protection against reconstruction and re-identification threats). The chosen privacy-loss budget (represented by "ε," the Greek letter "epsilon") for the P.L. 94-171 redistricting data is ε=17.14 for the persons file and ε=2.47 for the housing unit data. Learn more about the privacy-loss budget and epsilon at <https://content.govdelivery.com/accounts/USCENSUS/bulletins/2e32ea9> and <www.census.gov/data/academy/webinars/2021/disclosure-avoidance-series/differential-privacy-101.html>.
- **Accuracy.** Every disclosure avoidance method imposes a fundamental tradeoff between the degree of privacy protection and the resulting accuracy or usefulness of the data. The TopDown Algorithm was tuned specifically to meet or exceed accuracy targets critical to core redistricting needs, determined based on stakeholder feedback and consultations with the Department of Justice. These include the ability to accurately identify communities of interest in voting districts when enforcing Section 2 of the Voting Rights Act.

  Internal research concluded that the TopDown Algorithm met accuracy targets for all congressional and state legislative districts. A detailed description of the accuracy targets used and performance against those targets is available by viewing our working pager at <www.census.gov/library/working-papers/2021/adrm/SSS2021-01.html>, newsletter at <https://content.govdelivery.com/accounts/USCENSUS/bulletins/2e2545b>, and webinar at <www.census.gov/data/academy/webinars/2021/disclosure-avoidance-series/demonstration-data-for-redistricting-and-voting-rights-act-use-cases.html>.
- **Data availability.** In and of itself, differential privacy does not impact data availability. However, across data products, the 2010 Census released over 150 billion statistics on the 308,734,538 people it counted. Each statistic provided a clue to the identities of the people behind it. With the feedback of stakeholders, we are making difficult but data-driven decisions about balancing the level of detail we can provide in our published 2020 Census statistics, especially for smaller geographic areas and population groups, while protecting the privacy of individuals.

Will Differential Privacy Work for the 2020 Census?

Yes. Differential privacy is the only framework that can provably protect 2020 Census data against known and emerging re-identification threats while producing quality, fit-for-use data. More information is available at <https://www.census.gov/library/video/2021/protecting-privacy-in-census-bureau-statistics.html>.

United States® Census 2020

D-FS-GP-EN-050
August 2021

SUPPRESSION

What Is Suppression and How Does It Work?

In 1980 and earlier, the primary mechanism that the Census Bureau used to protect the confidentiality of individual census responses was to withhold publication ("suppress") of any whole tables of data or cells of data within the tables 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 there were at least five individuals in each reported race category.

The tables that were published were affected by additional individual cell suppression requirements. Not only were cells that failed thresholds suppressed (called primary cell suppression), but cells that could be used to deduce the value of those cells were also suppressed (called complementary cell suppression). Suppression is still used in various Census Bureau products today, including those from the Economic Census.

These suppression routines helped to protect confidentiality by reducing the detail of data published about individuals who were relatively unique within their communities.

What Impact Does Suppression Have on Confidentiality, Accuracy, and Data Availability?

- **Confidentiality.** Tables and cells that aren't published retain privacy.
- **Accuracy.** Suppression results in highly accurate data for those geographies that aren't suppressed.
- **Data availability.** Data users were dissatisfied with the amount of suppression required in the 1980 Census (the last census for which this method was used). Recent experiments by our disclosure avoidance team revealed that, if used today, suppression would result in blocking publication of far more tables and cells than when it was last used.

  If suppression rules as we last used them were applied to 2010 Census data, more than 80 percent of block, block group, and tract level geographies would not have any data for two key Census Redistricting (P.L. 94-171) Data Summary File tables: the population 18 years and over by race, and the population 18 years and over by race and Hispanic origin.

  Those tables that could have been published would have also been affected by additional individual cell suppression requirements. Not only would cells that failed thresholds be suppressed (primary cell suppression), but cells that could have been used to deduce the value of those cells would have also been suppressed (complementary cell suppression). About eight percent of cells for block-level data would have been suppressed using primary cell suppression. That number would have been higher factoring in complementary cell suppression.

Could Suppression Work Today for the Decennial Census?

No. An unacceptably high number of whole tables and individual data cells would not be published given current re-identification threats.

SWAPPING

What Is It and How Does It Work?

Between 1990 and 2010, the Census Bureau used a form of noise infusion, "swapping," to safeguard respondent confidentiality. This method adds statistical "noise" (uncertainty) to the data by swapping perceived "outlier" households between blocks, block groups, tracts, or counties.

D-FS-GP-EN-0509
August 2021

United States® Census 2020

A confidential "key" determined the rules by which households were subject to swapping. The key matched households on several characteristics, including household size and the population over the age of 18, keeping those characteristics invariant (unchanged). Both the rate of swapping and the key itself were confidential by design.

What Impact Does Swapping Have on Confidentiality, Accuracy, and Data Availability?

- **Privacy.** In our recent experiments looking at the potential impact of older disclosure avoidance methods if used today, we found that even very high swapping rates had essentially no impact on re-identification outcomes. This was true even though we slightly altered the population size of up to half of the households and moved swap pairs beyond tracts for up to 70 percent of housing units (the "SwapHigh" experiment). A "SwapLow" experiment, with a basic 5 percent swap rate and no other alterations, yielded re-identification rates comparable to the published 2010 Census results. More information on the experiments is available at <www.census.gov/data/academy/webinars/2021/disclosure-avoidance-series/research-into-alternatives-to-differential-privacy.html>.

  Given that most of the population has a unique combination of attributes at the census block level, increasing the swapping rate and relaxing the household and location invariants would be required.
- **Accuracy.** Results from the SwapHigh experiment showed significant distortions in population and race. Note that because swaps occur prior to query tabulation, fine-tuning towards greater accuracy is not possible. The SwapLow experiment yielded accuracy rates comparable to the published 2010 Census results.
- **Data Availability.** Swapping does not limit the availability of data, just the accuracy and confidentiality of the data.

Could Swapping Work Today for the Decennial Census?

No. Our research makes clear that the relatively low swapping rate used in the 2010 Census does not protect respondent confidentiality. Higher swapping rates severely compromise data accuracy.

D-FS-GP-EN-050
August 2021

United States® Census 2020