"DIFFERENTIAL PRIVACY" IN THE U.S. CENSUS:

Protecting Individual Data Confidentiality; Preserving Accuracy

October 2025

The U.S. Constitution gives Congress responsibility for overseeing a census of population every ten years, for the purpose of apportioning seats in the U.S. House of Representatives (Article I, sec. 2) and ensuring fair representation under the Fourteenth Amendment. To carry out this duty, Congress has enacted a set of laws collectively known as the "Census Act," codified in Title 13, United States Code.

A hallmark of our census taking is the strict confidentiality Title 13 requires for all data the U.S. Census Bureau collects.[1] Significant research and experience — in the U.S. and around the world — show unequivocally that confidentiality protections are essential for convincing people that it is safe and beneficial to participate in the census. This trust is critical to ensuring a fair and accurate census.

The Census Bureau approaches the confidentiality of personal and household responses in two ways:

- First, *no one* outside of the Census Bureau can access personal, identifiable census responses, *for any reason* — not law enforcement agencies (at any level of government), landlords, courts of law, or other federal agencies.[2] Further, the law prohibits the use of census responses for any purpose other than tabulating anonymized statistics.
- Second, in order to secure the confidentiality of responses, the Census Bureau has adopted steps to ensure that the anonymous data it publishes at small geographic levels (such as block-level data for redistricting) will not allow re-identification of any person or household. It is that latter goal that has become more difficult in recent decades, with rapid advances in computer sciences and the growing volume and availability of other data sources that could make it possible to reconstruct the underlying information and identify individuals from even "anonymous" statistics.

It was this rapid evolution of confidentiality threats that led the Census Bureau to research and develop for census use, including during the first Trump administration, a newer method for protecting published census data from unintentionally disclosing sensitive personal information**.** Career Census Bureau experts had come to believe that previous traditional methods for obscuring small-area statistics from disclosure of personal information would not be adequate to ensure confidentiality for the 2020 Census and beyond.[3] ***Differential privacy***, they concluded, would help meet emerging technology challenges.

Differential privacy is a mathematical framework for adding carefully calibrated amounts of random statistical "noise" (i.e. imprecision) to published statistics, so that no one's information can be identified and, possibly, misused; however, overall statistical patterns remain accurate. While the technique itself is complex, the concept is not. Fundamentally, its application involves a trade-off between confidentiality and data usefulness. The fact is, ***differential privacy is not applied to*** **state population totals, including those used to apportion seats in the House of Representatives.**

---

[1] 13 U.S.C. §8 and §9.

[2] An individual, with proper documentation, may access their own census responses, for example, to prove their date of birth or place of residence. Census records are kept confidential for 72 years after the date of collection, after which the raw data are made available through the National Archives.

[3] Previous disclosure avoidance methods included *swapping* some responses between nearby households and replacing some responses with statistically estimated values,

Here's the bottom line:

- ➔ *Differential privacy* is a proven statistical technique. It is used to protect against the disclosure of confidential information about every person and household counted in the census, as federal law requires.
- ➔ The Census Bureau held numerous public briefings and stakeholder discussions, and invited public feedback on the application of *differential privacy* to different levels of geography, before making final decisions on applying the method to 2020 Census data before publication.
- ➔ *Differential privacy* does **not** affect the state population totals used to reapportion seats in Congress. Those are published without the injection of any statistical noise.
- ➔ *Differential privacy* can make data for small geographic areas and/or population groups less precise, in order to protect people's data. But it does not move people around from one place to another.

The Census Bureau takes seriously the trade-off between protecting the confidentiality of Americans' personal data and producing useful statistics to inform policy making and support the fair allocation of political representation, as the Constitution envisions. It should continue to work closely with stakeholders, outside experts, and policymakers to strike the best balance possible to meet all of those important goals.

---

[4] Here's a good two-page explanation of *differential privacy* from the Georgetown Center on Poverty and Inequality.