Northwestern | INSTITUTE FOR POLICY RESEARCH

POLICY RESEARCH BRIEF
December 2025

Differential Privacy's Trade-Offs Matter for Policy Use

OVERVIEW

Policymakers and others who make data-driven decisions often need to share summaries of data—census results, for example—without violating the privacy of the individuals in the dataset. While sharing the aggregated data would not seem to put any one person at risk, computer scientists have shown that it is possible to learn things about individuals from published data summaries.

Differential privacy (DP) is a leading technical standard designed to safeguard privacy by injecting statistical noise during computation of outputs—effectively obscuring individual-level information. DP provides strong privacy guarantees, but it can make policy-relevant statistics less accurate.

In a 2024 article, **Priyanka Nanayakkara** (PhD 2024), now at Harvard, and IPR computer scientist **Jessica Hullman** outline the trade-offs that emerge when using differentially private data to evaluate policy. Their guide draws on data-use cases to show that decisions to use DP aren't just technical—they are also shaped by organizational needs, social values, and the challenges of putting the system into practice.

POLICY TAKEAWAYS

- Differential privacy (DP) offers robust protection but may distort policy-relevant statistics, especially for small subgroups.
- Policymakers must balance privacy with data utility, transparency, and fairness.
- Adopting DP requires open and ongoing communication between data scientists and decision-makers.
- Rather than assume DP is always the ethical choice, agencies should evaluate if its use serves the context and values of a given policy domain.

IPR computer scientist **Jessica Hullman** and Harvard postdoctoral fellow **Priyanka Nanayakkara** examine the utility of differential privacy for policymakers.

The authors argue that DP should not be treated as a one-size-fits-all solution. Instead, policymakers must critically assess when DP's protections align with public values and policy goals.

FINDINGS

While DP is often seen as responsible and forward-looking, it introduces uncertainty into data analysis that can hinder policy development. Concerns exist that DP's noise can disproportionately impact small or marginalized populations, affecting equity-related assessments. For instance, data on rural communities or racial minorities may be especially distorted, undermining the very policies aimed at helping them.

In addition, DP can obscure where and how error is introduced, making it difficult to validate analyses. This opacity can erode trust in both the data and the institutions providing them. Additionally, decisions about how much privacy to provide often remain underexplained or undisclosed, leaving public-sector users without the context they need to make informed choices.

These concerns suggest that while DP is technically appealing, its practical deployment must be sensitive to institutional needs, analytic goals, and public accountability.

Differential Privacy (DP) for Policy: Key Considerations

The researchers offer guidance to help policymakers promote four key principles when considering the use of differential privacy.

PRIVACY PROTECTION

Consider whether the definition of DP aligns with **privacy needs** in a particular context. Be aware that DP might conflict with common attitudes toward **binary privacy protection**, interest in absolute risk, and expected rather than worst-case outcomes.

Be aware that the **nature and strength** of privacy protections offered under DP can **vary wildly** depending on implementation choices.

KNOWLEDGE PRODUCTION

Anticipate negative impacts of DP for **knowledge production** by simulating the impact of added statistical noise (at different magnitudes) prior to deployment.

Anticipate slowdowns in the process of **knowledge generation** if DP is introduced in place of no, or other, protection measures. **Allocate resources for training** data users to cope with the additional noise.

TRUST

The **addition of statistical noise** in published data may cause data users and others to **lose trust** in the data, especially if prior privacy protection measures were kept secret.

The **use of DP** may **increase trust** in the data collection process among some data subjects.

TRANSPARENCY

Explore opportunities for making deployment decisions **public**, **auditable**, and **interpretable** beyond the organization deploying DP.

If DP is used, provide **clear** and **interpretable** explanations of DP's guarantees and impacts to data subjects.

If DP is used, **provide documentation** about the addition of statistical noise and **provide guidance** for accounting for noise in analysis to data users.

CASE STUDY

For the 2020 census, the U.S. Census Bureau adopted DP to protect individuals' confidentiality while still publishing usable statistics. Unlike older methods such as swapping data between households, DP uses a mathematical system that limits how much can be learned about any one person from published results.

This change stirred controversy. Some worried that the added noise could reduce accuracy, especially for small towns, minority groups, and funding decisions that depend on precise counts. Others valued DP's transparency, since its algorithms can be shared without compromising privacy.

The Census Bureau's experience shows that protecting privacy today means more than technical fixes: It also requires clear communication, training, and trust-building among the parties involved so data can remain both safe and useful.

REFERENCE

Nanayakkara, P., and J. Hullman. 2024. What to consider when considering differential privacy for policy. *Policy Insights from the Behavioral and Brain Sciences* 11(2): 132–40.