Policy impacts of statistical uncertainty and privacy

Reforms could help funding formulas address unequal distribution of uncertainty from data error and privacy protections

By **Ryan Steed,[1] Terrance Liu,[2] Zhiwei Steven Wu,[2] Alessandro Acquisti[1]**

Differential privacy (*1*) is an increasingly popular tool for preserving individuals' privacy by adding statistical uncertainty when sharing sensitive data. Its introduction into U.S. Census Bureau operations (*2*), however, has been controversial. Scholars, politicians, and activists have raised concerns about the integrity of census-guided democratic processes, from redistricting to voting rights. The debate raises important issues, yet most analyses of trade-offs around differential privacy overlook deeper uncertainties in census data (*3*). To illustrate, we examine how education policies that leverage census data misallocate funding due to statistical uncertainty, comparing the impacts of quantified data error and of a possible differentially private mechanism. We find that misallocations due to our differentially private mechanism occur on the margin of much larger misallocations due to existing data error that particularly disadvantage marginalized groups. But, we also find that policy reforms can reduce the disparate impacts of both data error and privacy mechanisms.

Differential privacy is the cornerstone of Census Bureau's updated disclosure avoidance system (DAS) (*2*). Designed to rigorously prevent reconstruction, re-identification, and other attacks on personal data, differential privacy formally guarantees that published statistics are not sensitive to the presence or absence of any individual's data by injecting transparently structured statistical uncertainty (noise) (*1*). But even before differential privacy is applied, estimates from the Decennial Census, surveys like the American Community Survey (ACS), and other Census Bureau data products used for critical policy decisions already contain many kinds of statistical uncertainty, including sampling, measurement, and other kinds of non-sampling error (*4*). Some amount of those errors are quantified, but numerous forms of error are not (*5*), including some non-responses, misreporting, collection errors, and even hidden distortions introduced by previous disclosure avoidance measures such as data swapping (*6*). If quantified and unquantified errors alike are not acknowledged and accounted for, policies that rely on census data sources may not distribute the impacts of uncertainty equally.

In 2021, the U.S. federal government appropriated over $16.5 billion in Title I funds (including several special grants not analyzed here) to distribute to over 13,000 local education agencies (LEAs)—typically school districts—using a formula that takes as input census estimates of the number of children and children in poverty. School districts qualify for Title I grants based on the number or share of children in poverty (*7*). However, the formula does not account for deviations in the poverty estimates that could cause misallocations—cases where the funding amount allocated to a school district differs from its entitlement in an imaginary (*3*), noise-free world.

Researchers have recognized Title I an important case study of policy-relevant privacy-utility trade-offs (*9*), including misallocation after noise injection for differential privacy (*10*). We extend this work by comparing the policy impacts of noise injected for privacy to the impacts of existing statistical uncertainty, contextualizing preliminary error analyses by Census Bureau scientists (*2*). Our results empirically investigate analytical predictions and proposals from previous work on statistical estimation and federal funding formulas (*11, 12, 13*).

We focus specifically on the way Title I implicitly concentrates the negative impacts of statistical uncertainty on marginalized groups. Weakening privacy protection will do little to help the most vulnerable—for these communities, participating in a census survey can be especially risky, despite the benefits of voting rights protection and school funding. Historically, abuse of census data facilitated internment of Japanese Americans and other injustices (*3*). Today, a parent with a restrictive lease may not mention their children to a census worker because they fear being kicked out by their landlord if their responses are re-identified (*14*).

SIMULATING NOISE IN TITLE I ALLOCATIONS

Prior work on differential privacy in the context of Title I is purely analytical, analyzes abstracted components of funding formulas, or focuses only on basic grants (*9*, *10*). In contrast, we fully replicate the Title I provisions for allocating more than $11.6 billion in basic, targeted, and concentration grants using the same data sources and procedures as the Department of Education, which is responsible for calculating the official Title I grant amounts each year (*7*). We measure the impact of data and privacy deviations on the 2021 allocations to 13,190 LEAs across the United States. The primary data input is the Census Small Area Income and Poverty Estimates (SAIPE) from 2019—a table of counts of total population, children, and children in poverty in school districts from all fifty states (excluding Puerto Rico and other territories) that incorporates weighted survey estimates from the ACS (see supplementary materials (SM) §2 for details).

In a given year, the SAIPE may vary due to a variety of sources of error, including relative error in the county-level estimate, error from other data sources used (*e.g.*, tax data), and errors from raking and recombination methods used to convert county estimates to school district estimates (*4*). To simulate the effects of these "data deviations"—quantified data errors (*8*)—we generate alternative poverty estimates for each school district from a normal distribution around the published estimate of children in poverty in that district from the 2019 SAIPE, following prior work and Census Bureau guidance (*4*) (SM §2).

We then add "privacy deviations"—noise deliberately injected to achieve differential privacy. The Census Bureau has not yet announced any concrete plans for updated disclosure avoidance in the ACS, and the SAIPE currently does not inject noise for privacy on top of its inputs. To illustrate how privacy deviations might affect these and similar products, and to guide policymakers as the Census Bureau develops new disclosure avoidance measures, we follow prior work (*9, 10*) in applying the Laplace mechanism, a commonly used noise-injection procedure which is provably differentially private (*1*). Our hypothetical

[1]Heinz College of Information Systems and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA. [2]School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Email: ryansteed@cmu.edu

mechanism does not include the complex post-processing applied to the discrete Gaussian mechanism used in the Decennial Census; we only round negative numbers to zero (*2*).

The strength of differential privacy (described by the parameter ϵ) determines the magnitude of privacy deviations (lower ϵ implies stronger privacy and generally more noise). ϵ measures how much an individual's decision to respond to a census survey increases their risk of unwanted disclosure. It is not yet clear whether or how privacy deviations would be added to a statistical product like the SAIPE in practice, and because the SAIPE incorporates weighted survey estimates from the ACS, its sensitivity to changes in an individual's response is unclear. Instead, we try several reasonable privacy settings to provide an upper bound on the magnitude of privacy deviations that might be added in practice (*9*) (SM §2). We focus on $\epsilon = 0.1$ and $\epsilon = 1$ (SM §7 additionally varies ϵ from 0.001 to 10). Previous work on Title I (*9*) suggests $\epsilon \geq 2.52$; many applications use similarly high values (*2*), while differential privacy advocates often prefer $\epsilon < 1$.

The Title I legislation includes two post-formula provisions to achieve secondary policy goals. The "hold harmless" provision (20 U.S.C. §6332) limits funding losses to between 5% and 15% per year and the "state minimum" provision (20 U.S.C. §6333) sets a formulaic floor on the total amount received by each state. We treat the allocations generated without these provisions as the official formula-based "entitlements" for each district. Later, we compare these entitlements and the real allocations produced with these provisions. For each privacy setting, we compute the misallocation due to deviations by comparing the simulated allocations after deviations to the official entitlements. We repeat this procedure 1,000 times, drawing new data and privacy deviations in each trial. Our metric of group-weighted misallocation describes the expected misallocation borne by the average formula-eligible child in a given group nationwide, assuming that misallocation to a district is borne equally by all its eligible students.

SUBSTANTIAL MISALLOCATIONS

Of the roughly \$11.7 billion distributed nationwide in 2021, districts in our simulation expect to lose a total of \$1.06 billion (summing all losses in each simulation, then averaging summed losses across 1,000 simulations; SD = \$0.04 billion) in entitlements to other districts due to the Title I formula's handling of existing (before differential privacy) data deviations alone (Fig. 1). The standard deviation in misallocation (computed by averaging over 1,000 trials) is about \$835,000 (the average district receives around \$880,000)—\$237 per student. When we add privacy deviations (for a relatively strong privacy setting $\epsilon = 0.1$), the expected total entitlement loss only increases by \$50 million (4.7%; marginal SD = \$2.9 million). For a less strong privacy setting (smaller privacy deviations; $\epsilon = 1$), the increase is negligible. The marginal impact is small because—as in the 2020 Decennial Census (*2*)—the magnitude of privacy deviations is comparable to the magnitude of data deviations only in the least populous districts, even at a relatively strong privacy setting ($\epsilon = 0.1$) (SM §7).

These costs are geographically asymmetrical. Certain population-sparse school districts, especially in the Northwest, benefit greatly on average from data deviations (SM Fig. 3a)—their small sample sizes induce proportionally larger data deviations, and, because of their low absolute numbers of children in poverty (though poverty rates may still be high), they have more room to gain funding than to lose funding. Then, because the federal appropriation is fixed and allocations are zero sum, more populous districts, especially in the Southeast, pay for that proportional increase in funding with a small "tax" (*10*). Less populous districts gain even more as they qualify for new grants (*11*) (SM Fig. 4). Notably, while less populous, usually rural districts gain funding on average from data deviations, their allocations are more volatile (*10*) (SM Fig. 7).

When we add privacy deviations (for relatively strong privacy, $\epsilon = 0.1$), gains by small districts are even more exaggerated (SM Fig. 3b). Unlike data deviations, where the absolute variance increases with population size, our privacy deviations have the same variance in every district, exceeding data deviations in magnitude only in the least populous districts. Still, the marginal increase in cost to districts due to privacy deviations is much less than the base-level misallocations resulting from data deviations, and the marginal change reduces total misallocation about half the time.

DIVERSION FROM MARGINALIZED GROUPS

Due to Title I's distribution of quantified data deviations alone, Black students and Asian students can expect to lose around \$5 and \$8 per eligible student respectively, while White students gain over \$2 per eligible child on average (Fig. 2). (The average district receives \$1,120 per eligible student.) Likewise, school districts with large Cuban, Puerto Rican, and other Hispanic communities expect to lose funding (between \$3 and \$14 per eligible student) while non-Hispanic districts gain (SM §4, Fig. 9). For a child in a particular district in an unlucky year, the disparity may be worse. Whether a demographic group loses funding depends on whether its members tend to live in high- or low- poverty districts. Often, this happens because the poverty rate in the group itself is high. Groups that tend to live in denser, usually urban districts with more children in poverty lose out, while groups that live in sparse, often rural districts with fewer children in poverty (though the rate of poverty may be higher) gain. Geographically concentrated groups—such as tribal nations or racial subgroups (SM §4)—experience more volatility in outcomes across trials, which depend on the population density and poverty rates where they live.

In a relatively strong privacy setting ($\epsilon = 0.1$), our differential privacy mechanism aggravates these disparities, especially for Black students, who lose more than twice as much funding on average after noise is injected—possibly because Black students are more likely to attend populous school districts where the costs of privacy deviations accumulate. But in less strong privacy settings ($\epsilon \geq 1$), disparities change very little from the *status quo* when privacy deviations are added (SM §7).

To assess the impacts on non-categorical demographics, we also fit a generalized additive model (GAM) to the school district-level combined misallocations ($\epsilon = 0.1$) using district population characteristics: population density, median household income, proportion White, proportion Hispanic, proportion renter-occupied housing, and racial homogeneity (the Herfindahl-Hirschman index). Fitting the GAM on a sample of 100 trials, we find that districts with a median income between approximately \$25,000 and \$75,000 (about 56% of districts) can expect to lose out due to deviations, while most other districts gain (SM Fig. 6). The 40% most population-dense districts can also expect to lose funding. Conversely, districts that are less than 5% Hispanic tend to benefit from data and privacy deviations.

SIMPLE REFORMS

Simple changes to the formula—including additional provisions currently required by law—can alleviate or aggravate disparities. For example, adding the hold harmless provision reduces the standard deviation in misallocation (relative to the formula entitlement) but drastically increases disparities in outcomes for racial minorities (Fig. 2). Hold harmless prevents small districts from losing funding to data or privacy deviations, thereby increasing the tax on more populous districts and their non-White residents. The state minimum provision has a similar but smaller effect. Typically received by low population states, the state minimum slightly increases the amount of grants to low population districts, exacerbating disparities.

This result illustrates a tension in evidence-based formula funding: because estimates for less populous geographies have higher variance in both privacy and data deviations relative to their populations and entitlements, measures that overwhelmingly benefit those small areas burden larger areas. We tested proposed policy changes that could alleviate this tension (SM §6). We find that using multi-year averages with windows of increasing size decreases both overall misallocation and outcome disparities compared to when we use the averaged poverty estimates as a baseline (SM Fig. 14 and Fig. 15). In general, using an average diminishes both data deviations and the privacy deviations required to achieve differential privacy, limiting both increases in expected funding for less populous districts and alleviating worst-case outcomes. Averaging may even be just as effective at stabilizing funding year-to-year as the hold harmless provision (*11*). We also tested requiring repeated years of ineligibility before disqualifying districts from funding, which did not change overall misallocation—likely because it permits more marginally wealthy districts to receive funding—but did reduce disparities (SM Fig. 14 and Fig. 15).

PAYING FOR (PRIVATE) DATA

Simple policy changes can alleviate disparities in the impact of statistical uncertainty, but precisely targeted funding formulas will still have costs. Policymakers could ensure that no school district expects to lose money due to the underlying data deviations quantified in our simulation by assigning just \$107 million (SD = \$31 million) in targeted payments to individual districts that lose funding on average across 1,000 simulations. The cost of stronger privacy (using our simplified mechanism) could be much less: to compensate districts for only the expected additional lost funding due to privacy deviations, policymakers need only distribute an extra \$41 million (SD = \$3.8 million) for stronger privacy ($\epsilon = 0.1$), or \$1.7 million (SD = \$601,000) for less added privacy ($\epsilon = 1$) (SM §7). Still, a district's actual loss in any given year often greatly exceeds its expected loss, especially for less populous districts. To compensate districts for both data and privacy deviations all but the worst 5% of our simulations, an additional \$4.7 billion would be needed in the stronger privacy setting ($\epsilon = 0.1$). The cost is greater if policymakers wish to also compensate for the many other forms of error not quantified here, or for a stronger privacy mechanism.

It may be difficult to justify or legislate funding increases to just the districts expected to lose funding. Simply increasing the total federal appropriation to Title I (benefiting all districts unequally) by \$135 million (the combined total expected loss) would only compensate for about half of expected losses. However, a \$4.7 billion increase (95% loss coverage) would compensate for nearly all total expected losses and cut total 5% quantile misallocation roughly in half. The White House's proposed 2022 allocation—a \$20 billion increase, since reduced to \$1 billion in Congress—would completely compensate for privacy and data deviations incurred under the 2019 budget, but inequalities would remain. An overall budget increase would provide "no-penalty" compensation (*10*) for data and privacy deviations, but would not solve issues of relative equity (though budget increases do reduce the number of held harmless districts).

DISCUSSION

The addition of noise for differential privacy exposes epistemic issues with formula design predicted by early work on census-guided federal funding even before differential privacy was first proposed (*3*, *11*, *12, 13*). In fact, our results suggest that the impacts of differential privacy relative to other sources of error in census data could be minimal. But current legislation holds few allowances for the impacts of statistical uncertainty. Use of census data for the Title I formula is mandated "unless the Secretary and the Secretary of Commerce determine that some or all of those data are unreliable or. . . otherwise inappropriate" (20 U.S.C. §6333). National Research Council studies, commissioned by the Department of Education before ACS estimates were first incorporated in the SAIPE after 2005, warned against hard thresholds and hold harmless provisions (*12, 13*)—but these provisions are still in effect. Recently, the Biden administration proposed a new Title I budget that includes funding to improve the poverty estimates—but there are still no measures to update the formula to handle uncertain inputs. Simply acknowledging the effects of data error could improve future policy design for both formula funding and disclosure avoidance.

Our findings come with limitations. Injected noise is just the tip of the iceberg: many other unquantified forms of statistical uncertainty—including previous disclosure avoidance methods—affect poverty estimates in different ways (*5*). No confidentiality measures are directly applied to the SAIPE, but its inputs (mainly ACS and IRS data) may have hidden or unintended distortions due to swapping and other ad-hoc disclosure avoidance techniques (*6*). By replacing other methods of disclosure avoidance, differential privacy could even reduce the amount of overall misallocation due to uncertainty. Lacking an alternative source of poverty data, we do not assess the impacts of systematic biases, including under-counts of marginalized groups. Our analysis of the Title I allocation process also leaves out several elements that could affect the applicability of our findings to the real-world distribution of funds, including small district appeals (20 U.S.C. §6333) and district-level heterogeneity in use of funds. Temporal trends in funding, in combination with provisions like hold harmless, could compound the effects of deviations (*11*).

Data error—from under-counts to sampling error to noise injection—will always affect evidence-based policy to some degree. In 2017, 316 federal spending programs relied on U.S. census data to distribute over \$1.5 trillion in federal funding across states, cities, and school districts (*15*). Uncertainty in census data—including intentionally-added error for privacy—will incur costs for stakeholders in those programs. But, at least the quantifiable portion of those costs can be mitigated with uncertainty-aware policy design and budget increases—an avenue for compromise between targeted policy, equity, and also additional privacy.

REFERENCES AND NOTES

1. C. Dwork, F. McSherry, K. Nissim, A. Smith, Calibrating Noise to Sensitivity in Private Data Analysis, *Theory of Cryptography*, S. Halevi, T. Rabin, eds., Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2006), pp. 265–284.
2. J. Abowd, *et al.*, The 2020 Census Disclosure Avoidance System TopDown Algorithm, Harvard Data Science Review (2022).
3. d. boyd, J. Sarathy, Differential Perspectives: Epistemic Disconnects Surrounding the U.S. Census Bureau's Use of Differential Privacy, *Harvard Data Science Review* (2022).
4. U. C. Bureau, Quantifying Relative Error in the School District Estimates (2020).
5. R. M. Groves, L. Lyberg, Total Survey Error: Past, Present, and Future, *Public Opinion Quarterly* **74**, 849–879 (2010).
6. M. Christ, S. Radway, S. M. Bellovin, Differential Privacy and Swapping: Examining De-Identification's Impact on Minority Representation and Privacy Preservation in the U.S. Census, *2022 IEEE Symposium on Security and Privacy (SP)* (IEEE Computer Society, 2022), pp. 1564–1564. ISSN: 2375-1207.
7. T. Snyder, R. Dinkes, W. Sonnenberg, S. Cornman, Study of the Title I, Part A Grant Program Mathematical Formulas, *Statistical Analysis Report 2019-016*, U.S. Department of Education (2019).
8. B. D. Spencer, Statistical Aspects of Equitable Apportionment, *Journal of the American Statistical Association* **80**, 815–822 (1985). Publisher: [American Statistical Association, Taylor & Francis, Ltd.].
9. J. M. Abowd, I. M. Schmutte, An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices, *American Economic Review* **109**, 171–202 (2019).
10. D. Pujol, *et al.*, Fair decision making using privacy-protected data, FAT* 2020 - Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (Association for Computing Machinery, Inc, 2020), pp. 189–199.
11. A. M. Zaslavsky, A. L. Schirm, Interactions Between Survey Estimates and Federal Funding, *Journal of Official Statistics* **18**, 371 (2002).
12. National Research Council, *Small-Area Income and Poverty Estimates: Priorities for 2000 and*

Beyond (National Academies Press, Washington, D.C., 2000). Pages: 9957.
13. National Research Council, *Statistical Issues in Allocating Funds by Formula* (National Academies Press, Washington, DC, 2003).
14. Panel Discussion on Key Privacy Issues: Privacy and Census Participation, *2020 Census Data Products: Data Needs and Privacy Considerations: Proceedings of a Workshop*, D. L. Cork, C. F. Citro, N. J. Kirkendall, eds. (The National Academies Press, Washington, DC, 2020), pp. 129–133.
15. A. Reamer, Comprehensive Accounting of Census-Guided Federal Spending (FY2017), *Brief 7A*, George Washington Institute of Public Policy (2020).

Acknowledgments: We thank Claire McKay Bowen, danah boyd, Aloni Cohen, Bill Eddy, Zeid El-Kilani, Rayid Ghani, Priyanka Nanayakkara, David Pujol, Aaron Roth, Ian Schmutte, Jayshree Sarathy, seminar participants at Carnegie Mellon, Columbia, the Simons Institute at Berkeley and the Census Bureau, and anonymous referees for their very helpful insights and feedback on this research. We also thank Jerry Maples, Bill Sonnenberg, and Todd Stephenson for their help with replicating the poverty estimation and Title I allocation processes.

Funding: National Science Foundations Grant # 1939606 (SW, TL); National Science Foundations Grant # 2120667 (SW, TL); Alfred P. Sloan Foundation Grant G-2015-14111 (AA); MacArthur Foundation Grant 22-2203-156318-TPI (AA).

Citation: This is the accepted version of the following article: R. Steed, T. Liu, Z. S. Wu, A. Acquisti, Policy impacts of statistical uncertainty and privacy, *Science* **377**, 928–931 (2022), which has been published in final form at science.org/doi/10.1126/science.abq4481.

10.1126/science. abq4481

Supplementary Materials
science.org/doi/10.1126/science.abq4481

Fig. 1. Expected sum of lost entitlements in basic, concentration, and targeted grants over 1,000 trials—with quantifiable data error, or data deviations, alone (purple) and with the addition of injected noise for privacy, or privacy deviations (purple plus orange). Noise is injected with the ε-differentially private Laplace mechanism. The margins of error at 99% confidence are too small to be depicted—less than $4 million for all three bars. Note that for $\epsilon = 1.0$, the additional funding loss due to privacy deviations falls within the 90% margin of error for the impact of data deviations alone.

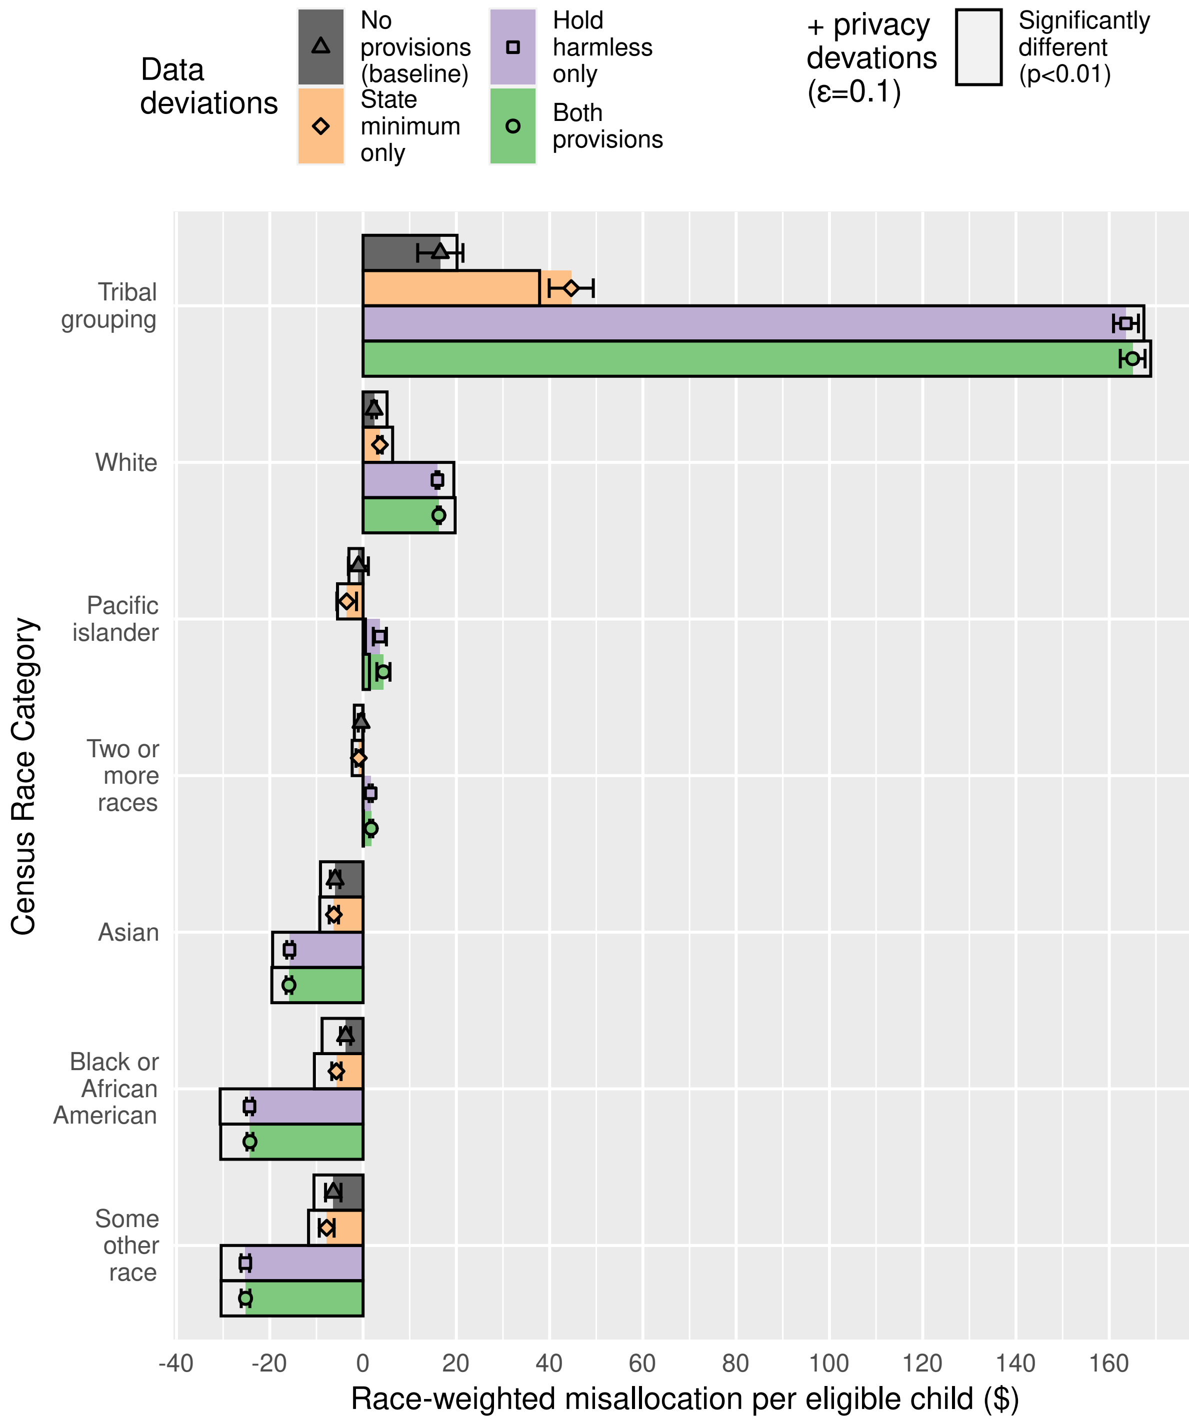

Fig. 2. Expected misallocation borne by the average formula-eligible child in a given census racial group nationwide (assuming each child in a district is affected by misallocation equally). Specifically, bars depict the nationwide sum of each district's misallocation multiplied by the proportion of respondents of a given census single race category in that district, divided by the total nationwide number of eligible children of that race (SM §2). Averaged over 1,000 trials. The colored bar and point indicate the race-weighted misallocation due to data deviations (data error) alone, with an error bar spanning a 90% normal confidence interval for this quantity. The black-outlined bar indicates the race-weighted misallocation due to combined data deviations and privacy deviations (noise injected for privacy, drawn from Laplace mechanism$\epsilon = 1.0$). The additional impact of privacy deviations is significant ($p < 0.01$) for all groups, according to a two-sample z-test.