

The Complexities of Differential Privacy for Survey Data

Jörg Drechsler* & James Bailie†

Abstract

The concept of differential privacy (DP) has gained substantial attention in recent years, most notably since the U.S. Census Bureau announced the adoption of the concept for its 2020 Decennial Census. However, despite its attractive theoretical properties, implementing DP in practice remains challenging, especially when it comes to survey data. In this paper we present some results from an ongoing project funded by the U.S. Census Bureau that is exploring the possibilities and limitations of DP for survey data. Specifically, we identify five aspects that need to be considered when adopting DP in the survey context: the multi-staged nature of data production; the limited privacy amplification from complex sampling designs; the implications of survey-weighted estimates; the weighting adjustments for nonresponse and other data deficiencies, and the imputation of missing values. We summarize the project's key findings with respect to each of these aspects and also discuss some of the challenges that still need to be addressed before DP could become the new data protection standard at statistical agencies.

1 Introduction

Differential privacy (DP) (Dwork et al., 2006) has become the quasi-gold standard in recent years for data collection and dissemination whenever privacy or confidentiality is a concern. It offers formal (that is, mathematically quantifiable) privacy guarantees by bounding the influence that any single record of the database can have on the computed outputs. The fundamental difference to earlier privacy frameworks such as k -anonymity is that the guarantees are a property of the mechanism generating the output and not a property of the data. DP specifies how much noise the mechanism needs to introduce to ensure that the probability of obtaining a specific output does not change substantially if one record in the database is changed. In simple examples where we

*joerg.drechsler@iab.de

Institute for Employment Research, Nuremberg, Germany; Ludwig-Maximilians-Universität, Munich, Germany; University of Maryland, College Park, USA.

†jamesbailie@g.harvard.edu

Department of Statistics, Harvard University, Cambridge, USA.

are interested in creating a DP version of an unprotected statistic such as the sample mean, the required amount of noise depends on two components. The first component is the privacy loss parameter ε (also called the privacy loss ‘budget’), which determines how much the probability of obtaining a specific result is allowed to change.¹ The smaller the privacy loss parameter, the more noise needs to be added and the better the level of protection offered. The second component is the sensitivity of the unprotected statistic of interest, which is measured as the maximum possible change of the statistic when changing one record in the database. The higher the sensitivity, the more noise needs to be added.

To illustrate, we can look at one of the classical DP mechanisms that is often used as a building block in more complex algorithms: the Laplace mechanism, which, for any univariate statistic f , ensures ε -DP by adding a random draw from a Laplace distribution centered at zero with scale parameter $b = \Delta f / \varepsilon$. The parameter Δf is the sensitivity of f measured as the maximum absolute distance (the L_1 norm) of the statistic computed over two neighboring datasets, i.e., two datasets that differ only in a single record. With this mechanism, the dependence on the two parameters is obvious: More noise is added for outputs with higher sensitivity and smaller values of ε .

This is one of the attractive properties of DP. The concept is very intuitive and requires only three steps, which in principle seem straightforward to apply: (i) define the maximum privacy loss that is still considered acceptable and select a value for ε accordingly; (ii) identify the sensitivity of the statistic of interest (for example, the sensitivity of a proportion under bounded ε -DP is simply $1/n$, where n is the number of records in the database); and (iii) choose a DP mechanism that infuses the right amount of noise into the reported output based on the parameters from steps (i) and (ii). Of course, in practice all three steps have their challenges. The discussion on how to choose and interpret the privacy loss parameter shows no signs of abating (Abowd & Schmutte, 2019; Bailie et al., 2026b; Drechsler, 2023; Dwork et al., 2019; Hsu et al., 2014; Nanayakkara et al., 2023; Tschantz et al., 2020); the sensitivity of the output is not always easy to compute and can be unbounded without further assumptions (Casacuberta et al., 2022); and finding a suitable DP mechanism can be challenging. Besides, there are often some hidden complications to DP in practice beyond what this three-step process makes apparent (Abowd et al., 2022; Cummings et al., 2024; Seeman & Susser, 2024). (For example, for the same research question there can be multiple choices for which statistic is used in step (ii), and it can be difficult to determine which one leads to the most efficient DP mechanism.) Still, the three components remain the same across applications and at least the general setup is well defined.

¹For simplicity we limit our exposition to the classical bounded ε -DP setting, where ‘bounded’ means that neighboring datasets are defined as datasets that can be obtained by replacing a single record with another record without changing the size of the database. Similar arguments would apply for other variants, such as (ε, δ) -DP, ρ -zero-concentrated DP, or f -DP, and for other definitions of neighboring datasets, such as unbounded DP for which a neighboring database is obtained by adding or removing a single record.

However, when working with survey data, there are additional complexities that typically do not arise in other settings. Moreover, the implications of using DP in the context of surveys have received little attention in the DP literature until recently. This led the U.S. Census Bureau to conclude in 2022 that “the science does not yet exist” to implement DP in its American Community Survey (U.S. Census Bureau, 2022c). An expert panel convened by the National Academies of Sciences, Engineering, and Medicine reached a similar conclusion with respect to the Survey of Income and Program Participation (National Academies of Sciences, Engineering, and Medicine, 2024).

Given its commitment to formal privacy for all its data products, including its surveys (U.S. Census Bureau, 2018), the U.S. Census Bureau sponsored a research project from 2020 to 2025 to better understand the complexities that arise when adopting DP in the survey context. In this paper, we will summarize some of the key findings of this project and also discuss some of the challenges that still need to be addressed. Overall, we identify five aspects that need to be considered when implementing DP in the survey context:

- Data production is a multistage process. As such, there are various options for how and where to integrate DP in this pipeline, each of which come with their own advantages and disadvantages.
- Previous studies found that sampling can amplify DP’s privacy guarantees. However, these amplification effects do not necessarily hold for the complex sampling designs used by statistical agencies.
- These complex sampling designs need to be incorporated into any survey statistic and hence must also be incorporated into any DP mechanism.
- Weighting adjustments are routinely used to account for unit nonresponse and to benchmark to known population totals. As these adjustments can substantially increase the sensitivity of the survey statistic, there is a need to develop robust adjustment strategies that are congenial to DP.
- Item nonresponse is often addressed using imputation, but similar to weighting adjustments some standard imputation techniques can greatly inflate the sensitivity of the resulting statistic. Ongoing research is currently investigating the feasibility of differentially private imputation techniques.

We will discuss each of these aspects in the remainder of this paper.

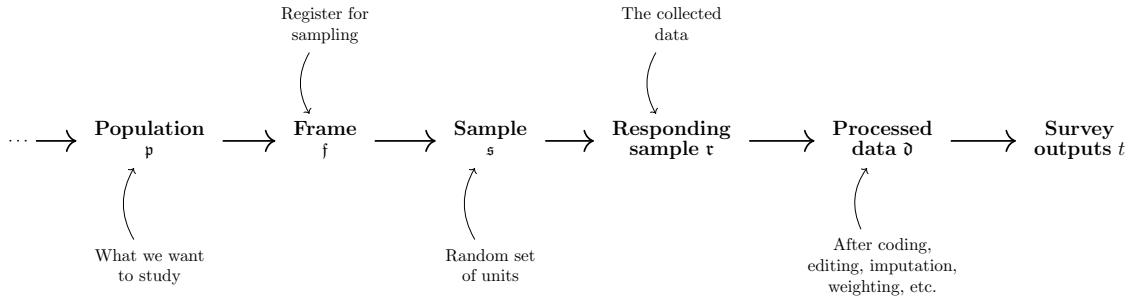


Figure 1: A *survey pipeline* consists of multiple steps, of which some of the most important are: determining the target population to be studied; constructing the frame; drawing the sample; collecting survey data from the responding units; processing the data (including coding free-form responses; editing inconsistent or improbable data; imputing missing records or variables; calculating the survey weights; and injecting privacy-protecting noise); and computing the survey outputs. There are of course additional steps to a survey pipeline after the survey outputs are released (such as data analysis) but, as they are not important to this paper’s subject, we exclude these steps from discussion. While not shown in this figure, it should be noted that data from previous stages of a pipeline are often used in later stages. (For example, the frame is usually used in computing the survey weights during the production of the processed data.)

2 DP and the Multistage Process of Data Production

2.1 The Survey Pipeline

The production of survey data is a complex multistage process (Figure 1). The design of a survey typically begins by conceiving the *target population*: the set of units that one wants to study. Usually, the target population is not actually specified as a concrete list of units. Instead, it is defined conceptually: “all adults in Massachusetts” or “all businesses in Hawaii.” Once the target population has been defined, the *frame* is sourced. The frame is a register of units from which the sample will eventually be drawn. It must include sufficient contact information so that the sampled units can be surveyed. The frame should align with the target population as much as possible. However, perfect alignment is not possible in most cases, even when the target population and the frame have the same inclusion criteria, because errors will typically be made in the frame’s construction. These errors will result in overcoverage (including units that are not in the target population) and undercoverage (not including units that are in the target population).

A *sample* is randomly drawn from the frame according to the survey’s *sampling design*: the probability distribution that specifies for every potential sample the chance that sample is selected. After sample selection, the statistical agency will solicit survey data from the sampled units. Most surveys, especially modern ones, suffer from nonresponse. This means only a subset of the sampled units will respond and the agency will not obtain survey data from the other units. Data collected

from the responding sample, along with the frame and some auxiliary information (such as data from administrative records or from previous censuses or surveys), are passed through a number of complex data processing steps before the survey outputs are computed and released. These data processing steps often include editing survey responses to correct errors in data recording; coding each free-form answer into a categorical variable; imputing missing answers to individual survey questions (“item nonresponse”) or to the entire survey questionnaire (“unit nonresponse”); and calculating multiple sets of survey weights for each record—to account for unequal probabilities of selection in the sampling design, to mitigate bias due to nonresponse patterns, and to calibrate survey data to auxiliary sources of information. Finally, we note that data may be deliberately injected with artificial noise at any point in the survey pipeline, so that releasing the survey outputs does not breach the privacy of the data subjects.

2.2 DP in the Survey Pipeline

DP is a criterion applied to *data release mechanisms*: algorithms that take data as input and produce a set of outputs that will then be published (that is, “released”). (For a more extensive, yet still approachable, discussion on what DP is, we direct the reader to the companion article Bailie et al., 2026a, in this volume.) Implementing DP involves both designing a data release mechanism that is compliant with DP, as well as integrating that mechanism into the relevant data pipeline. Both tasks are crucial for successfully producing outputs with high accuracy and good privacy protection.

There are two important considerations when integrating a DP mechanism into a data pipeline. Firstly, at what point in the pipeline should the DP mechanism *start*? And secondly, which of the earlier stages of the data pipeline should be considered *invariant* – i.e., should be treated as fixed – by DP? With survey pipelines, there are a number of possible options with respect to both considerations. In the option most commonly seen in the DP literature, the data release mechanism starts at the end of the pipeline and performs just the last step – computing the survey outputs from the processed data – and none of the previous steps are taken as invariant (Figure 2a). However, a mechanism could conceivably start at any point of the survey pipeline and incorporate all the steps that follow. For example, it could take as input the frame, execute the sampling step, process the data and finally compute the survey outputs (Figure 2b). Furthermore, any of the steps before the mechanism starts could conceivably be taken as invariant. (For explanations of how steps of the data pipeline can be taken as invariant by DP, or how a mechanism’s starting point can be encoded into DP, we again direct the reader to Bailie et al., 2026a.) In the rest of this section, we will explore these two considerations in turn.

Throughout this paper, we assume that the data release mechanism always includes the final step

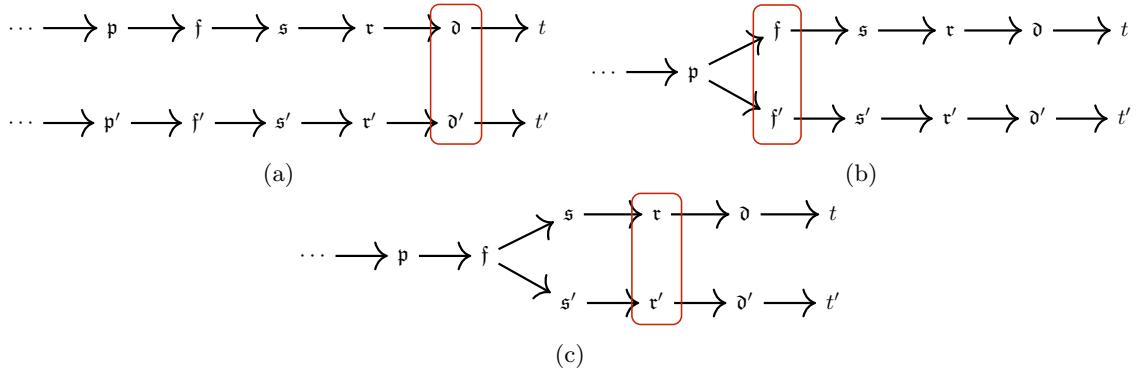


Figure 2: Three examples of where to start the data release mechanism (circled in red) in the survey pipeline and which of the previous stages to take as invariant (those stages before the pipeline branches). Recall from Figure 1 that p denotes the population, f the frame, s the sample, r the responding sample, d the processed data and t the survey outputs. The apostrophe ' indicates an alternative realisation of the associated variable. Figure (a) illustrates the standard approach in which there are no invariants and the data release mechanism only executes the final step of the survey pipeline—transforming the processed data into the survey outputs. In Figure (b), the mechanism begins with the frame and includes the sampling, responding and processing steps. The population is considered invariant. In Figure (c), the mechanism takes as input the responding sample. Both the population and the frame are taken as invariant, so that DP only compares samples from the same frame. This reduces the sensitivity of weighted estimators at the expense of less privacy (Section 4).

of the survey pipeline, the computation of the survey outputs.² Under this assumption, a survey pipeline can be split into those steps that are executed before the data release mechanism starts and those steps that are executed by the mechanism. Yet choosing where to make this split is not a simple matter. In fact, there are a number of complexities associated with starting the data release mechanism earlier or later in the pipeline. We identify five.

Firstly, starting the DP mechanism earlier can complicate the computation of the cumulative privacy loss across multiple data release mechanisms because DP’s composition theorems³ are not applicable when there is dependence between the mechanisms’ noise terms (which can happen, for example, when their sampling designs are dependent) (Bailie & Drechsler, 2026).

Secondly, as we will describe in Section 3, including the sampling step within the data release mechanism can amplify DP’s privacy guarantees without degrading data utility. However, this privacy amplification can be nullified if the attacker knows that the record they are attacking is in the sample (Bailie & Drechsler, 2026). More generally, if the attacker has knowledge about information intermediary to the DP mechanism (that is, information that is conditionally dependent on confidential data, or on the artificial noise introduced by the mechanism, conditioning on the output of the mechanism), the privacy guarantees afforded by DP can be weakened. For this reason, DP prohibits the direct release of such information. Therefore, because the choice of the sampling design is often dependent on data in the frame, the sampling design cannot be directly made public but instead can only be released by including it in the set of DP-protected survey outputs.

However, defining a data release mechanism – let alone one that satisfies DP – that releases the sampling design is challenging due to the third complexity we identify: Incorporating existing steps of a survey pipeline into a data release mechanism can be difficult. A data release mechanism is an algorithm that must be fully specified in order to be analysed by DP; hence any stage of the survey pipeline must first be fully “algorithmized” (that is, the process by which each of the stage’s

²Technically, a data release mechanism is simply an algorithm that takes data as input and outputs some (possibly noisy) transformation of that data. So, in principle, a data release mechanism could be incorporated into a survey pipeline even if it ends before the final step of the pipeline. (And such a mechanism could still be compliant with DP.) In this case, the survey pipeline includes additional post-processing steps after the data release mechanism ends but before the computation of the outputs that will be published. Such post-processing steps are usually included to improve the utility, usability or accessibility of the survey outputs. On the other hand, any data release mechanism can always be extended to one that ends with the final step of the survey pipeline, and any DP guarantees afforded to the original mechanism automatically carry over to the extended one by the post-processing theorem. (The post-processing theorem states that any function of a DP mechanism’s output – i.e. any “post-processing” – also satisfies DP with at most the same privacy loss.) Therefore, we do not gain anything by considering DP mechanisms that end before the survey pipeline’s final step.

³A composition theorem describes how to bound the total privacy loss incurred by multiple DP data releases that are all based on the same confidential dataset. For example, the composition theorem for pure ϵ -DP states that: if there are K mechanisms M_1, \dots, M_K , which all satisfy pure ϵ -DP and all have the same input dataset, then the total privacy loss – that is, the privacy loss of the mechanism that publishes all the outputs of M_1, \dots, M_K together – is bounded by the sum $\sum_k \epsilon_k$ over the privacy losses ϵ_k of each mechanism M_k . Existing composition theorems assume that the noise added by each mechanism is “fresh”, i.e., independent of everything else.

possible inputs is transformed into one of its outputs must be completely and programmatically specified) before it can be included in a mechanism.⁴ A survey pipeline often includes a number of complex, ill-defined and human-intensive tasks, such as building the frame, choosing a sampling design, coding and editing. Because these tasks all usually require a degree of human judgment, they would be difficult to algorithmize. Moreover, including these procedures – or other procedures often found in a survey pipeline – in a data release mechanism can add difficulties to making the mechanism compliant with DP. (In later sections, we will discuss some such difficulties as they relate to the weighting and imputation procedures.)

Fourthly, even if a data release mechanism begins later in the survey pipeline so that some steps of the pipeline do not have to be incorporated in the mechanism, implementing DP still requires understanding those steps' effect on the mechanism's input data. For example, some imputation techniques replace missing records with copies of non-missing donor records. This means an individual survey respondent can contribute to multiple records in the post-imputation dataset. This complicates the appropriate definition of neighboring datasets, since there is no longer an exact correspondence between the dataset's records and the real-world entities (the individual respondents) that should be protected: In the post-imputation dataset, changing a single record does not correspond to changing the data of one entity. Hence, naively applying DP to the post-imputation dataset will not provide a donor record with the expected level of protection; that is, the privacy guarantees for a donor record will be weaker than those for a post-imputed record. In general, the later the DP mechanism begins, the more difficult it is to determine an appropriate notion of neighbors since steps earlier in the pipeline may introduce dependencies between dataset records, thereby complicating the relationship between records and data subjects.

Fifthly, and most fundamentally, the starting point of the data release mechanism determines what form of the data is protected by that mechanism (Bailie et al., 2026b). For example, if a DP mechanism begins after data processing, then it is the processed data – and not, for example, the raw responses from the data providers – that are protected by that mechanism. That is to say, DP guarantees implicitly assume that the attacker is interested in inferring the data that is input into a DP mechanism. Measures of protection are in terms of the attacker's ability to learn this input data – and not the data at other points in the pipeline. If the DP mechanism takes the processed data as input, then the DP guarantees apply to the processed data and do not necessarily carry over to the responding sample data. In order to have guarantees for the responding sample, the

⁴The post-processing theorem provides an exception to this general rule. If the preliminary steps of a data release mechanism, taken on their own, satisfy DP, then the later steps of the mechanism do not need to be algorithmized, because the post-processing theorem ensures that the mechanism as a whole always satisfies DP regardless of what the later steps do. All that must be checked is that the later steps only use the DP outputs from the preliminary steps, and not some other data. However, this exception does not apply to the survey pipeline steps under discussion (choosing a sampling design, coding and editing) because these steps are typically applied before – not after – privacy protection.

statistical agency must show that the pipeline from the responding sample to the survey outputs (considered as a data release mechanism) also satisfies DP.

These five complexities demonstrate that there can be conflicting demands in deciding where a DP mechanism should start within the survey pipeline. For example, suppose a statistical agency wants to protect the unprocessed survey responses. Then either the coding and editing steps will need to be included in the agency’s mechanism (which may be difficult because these steps could be hard to algorithmize) or these steps will need to be removed from the survey pipeline (which could decrease the quality of the survey outputs).

We now return to the question of which steps of the survey pipeline should DP take as invariant. DP assesses the privacy of a data release mechanism by comparing the survey outputs’ distribution under pairs of counterfactual input datasets. These input datasets are generated by counterfactual runs of the initial steps of the pipeline, up until the data release mechanism begins. By taking some of these steps as invariant, DP’s counterfactual comparisons are reduced to only those pairs of input datasets that share the same realization of the invariant steps. For example, suppose the steps in the survey pipeline that generate the population and the frame are taken as invariant and the data release mechanism starts with the responding sample (Figure 2c). Then DP only compares those responding samples (i.e. those counterfactual input datasets) that could have come from the same frame. Adding invariants will weaken the privacy guarantees provided by DP (Abowd et al., 2022; Bailie et al., 2026c; Kifer et al., 2022). In general, the later the stage of the pipeline that is kept invariant, the greater the reduction in privacy. However, invariants may be justifiable when the output of the invariant steps can be considered as public knowledge (such as if the frame was sourced commercially rather than constructed from confidential information). Moreover, constraining some steps to be invariant has the advantage of reducing the sensitivity of weighted estimators and thereby decreasing the noise that must be added for privacy protection (Section 4).

3 DP with Complex Sampling Designs

Statistical agencies have been aware for decades that sampling can be a simple and effective strategy to reduce disclosure risks simply because an attacker can no longer be sure whether a specific target record is included in the sample or not. This is the main motivation why most statistical agencies only release samples from their censuses as public use micro datasets (they typically also apply additional measures to further increase the level of protection). This idea has been formalized in several papers in the context of DP (Balle et al., 2018; Bun et al., 2015; Kasiviswanathan et al., 2011; Wang et al., 2016, 2019). The authors show that the level of privacy is amplified through sampling, i.e., the actual privacy guarantees are higher than those implied by the chosen privacy loss parameters when protecting the sample output. Specifically, for small sampling rates r and

small privacy loss parameters ε , applying certain simple sampling designs (simple random sampling with and without replacement, and Poisson sampling) before running an ε -DP mechanism reduces the privacy loss to approximately $r\varepsilon$.

However, most surveys conducted by statistical agencies use complex multistage sampling designs, potentially with different sampling strategies at the different stages. These designs are primarily used to increase the accuracy of the survey outputs or to reduce the survey's operational costs. For example, the Current Population Survey (CPS), one of the flagship surveys of the U.S. Census Bureau, uses a two-stage sampling design in which stratified cluster sampling with probability proportional to size (PPS) is used to select clusters at the first stage and systematic sampling is used to sample households within clusters at the second stage (U.S. Bureau of Labor Statistics, 2018). There is no reason to believe that amplification effects for these complex designs are comparable to those obtainable for the simple designs discussed above. Bun et al. (2022) study the amplification effects for complex designs and find that amplification is small for most of the sampling designs used in practice. Their findings can be summarized as follows:

- Cluster sampling using simple random sampling without replacement to draw the clusters offers negligible amplification in practice except for small ε and very small cluster sizes.
- With minor adjustments, stratified sampling using proportional allocation can provide privacy amplification. For small ε , the amplification is still linear in the sampling rate up to a constant factor.
- Data dependent allocation functions such as Neyman allocation for stratified sampling will likely result in privacy degradation. (The effects will depend on the sensitivity of the allocation function.)
- With PPS sampling at the individual level, the privacy amplification will linearly depend on the maximum probability of inclusion (for small ε).
- Systematic sampling will only offer amplification if the ordering of the population is truly random. In all other cases, systematic sampling will suffer from the same effects as cluster sampling, leading to no amplification (assuming the ordering is known to the attacker).

In practice this implies that for many multistage sampling designs, which typically start with (multiple stages of) stratified cluster sampling, amplification effects can generally only be expected from those stages at which individual units or households are selected (typically the last stage of selection).

4 DP for Survey Weighted Estimates

As discussed in the introduction, the amount of noise that needs to be added to achieve a specific privacy loss ε directly depends on the sensitivity of the statistic of interest. Intuitively, this makes sense. If the statistic changes substantially when one record is changed in the data it will be easier to infer that record's value from observing the statistic and thus more noise will need to be added to sufficiently protect that record. From a utility perspective, this implies that more reliable (less noisy) DP outputs can be expected from statistics with low sensitivity. Thus, a common strategy with DP is to identify estimation strategies with low sensitivity and replace very sensitive estimates with less sensitive alternatives, for example by using robust statistics (Avella-Medina, 2021; Dwork & Lei, 2009).

When analyzing survey data, it is generally important to take the sampling design into account since the probabilities of selection typically vary between the units included in the sample. Unweighted estimates, especially those for descriptive statistics such as means and totals, will be biased whenever there are varying selection probabilities. To obtain unbiased estimates, each observation needs to be weighted by the inverse of its probability of selection. Hence, statistical agencies typically provide survey weights to enable researchers to take the survey design into account. In practice, these survey weights will also account for nonresponse and other data deficiencies such as undercoverage. (We will address this extra layer of complexity in the next section.)

Using survey weighted estimates raises the question: how (if at all) does the sensitivity of a statistic change when the survey design is taken into account? To illustrate the possible impacts, let us assume the analyst is interested in estimating the mean of some variable Y in the population using the sampled values y_i , $i = 1, \dots, n$, where n denotes the sample size. If the probabilities of selection were equal for all units, the sample mean would be an unbiased estimate for the population mean and its sensitivity would be R/n , where $R = \max(y_i) - \min(y_i)$ is the range of all possible values for y_i .⁵ When dealing with unequal probabilities of selection, a popular estimator for the population mean is the Horvitz-Thompson estimator (Horvitz & Thompson, 1952): $\hat{\mu}_Y^{HT} = \sum w_i y_i / N$, where w_i is the weight of unit i , for $i = 1, \dots, n$ and N is the size of the population. Note that we assume for simplicity that N is known and does not need to be protected and w_i is the design weight, i.e., it only accounts for the sampling design.

If we can treat the weights as fixed, the sensitivity of $\hat{\mu}_Y^{HT}$ is $\max(w_i)R/N$. Whether the maximum is over all units in the frame, over all units in the population, or over all possible counterfactual units, depends on which stages of the survey pipeline are treated as invariant as discussed in Section 2.2. Note that for equal-probability designs all $w_i = N/n$ and thus the sensitivity of the

⁵Throughout this section, we consider the bounded ε -DP setting. Similar arguments (with slightly different values for the sensitivity of a statistic) would apply for other settings.

Horvitz-Thompson estimator is the same as for the unweighted estimator. If $\max(w_i) > N/n$, the Horvitz-Thompson estimator will have larger sensitivity than the unweighted estimator.

However, these discussions assume that the weights can be treated as fixed, that is, they do not change if a record changes in the database. For most sampling designs used in practice, such an assumption is unrealistic. For example, with sampling proportional to size (PPS), the i th record's probability of inclusion is given by $\pi_i = (n \cdot x_i)/N \cdot \bar{x}$, where x_i is the value for unit i of the measure-of-size variable X that is used to improve the efficiency of the sampling design, and $\bar{x} = \sum_N x_i/N$ is the population mean of X . Changing the value of X for a single record will change the probabilities of inclusion and thus the survey weights for all other records in the sampling frame. Therefore, the sensitivity will be larger compared to the setting with fixed weights as we no longer only need to consider the maximum possible change in a single record's value for Y . We also need to consider the impact of the weight change for all the other records even if their values for Y don't change.

A recently-proposed strategy to mitigate this potentially-substantial increase in sensitivity is to regularize the weights, as explored by Seeman et al. (2024). (An extreme version of this strategy would set all weights to be equal; this could be justifiable if the increase in the privacy noise due to the weights dwarfs the bias introduced by ignoring the sampling design.) Another possible strategy is to treat the frame as invariant as discussed in Figure 2c. Frame invariance assumes any two neighboring datasets must always originate from the same frame and so can only differ at the sample level (or later). Thus, the probabilities of inclusion will be constant between neighboring datasets. However, treating the frame as invariant has two additional implications that need to be considered. First, fixing the frame implies that privacy amplification from sampling is no longer possible (we would need to have neighboring datasets at the frame level in order to achieve amplification). However, given the results of Bun et al. (2022), this amplification is likely small in practice and thus the positive effects of reducing the sensitivity will tend to outweigh the negative effects of losing the amplification effect. On the other hand, fixing the frame will restrict the possible counterfactual input datasets to those that are consistent with the realized frame. Because this restriction will fix the survey weights, it might introduce strong constraints on the possible neighboring datasets, depending on the sampling design. As a consequence, the actual privacy guarantees for a frame invariant setting could be significantly weaker than the guarantees under a non-frame-invariant setting even for the same privacy loss parameter. How problematic this reduction in privacy is in real settings is currently an open question for research.

5 DP and Weighting Adjustments

In practice, two adjustment steps are commonly applied to the design weights to correct for unit nonresponse and other data deficiencies such as over- or undercoverage in the sampling frame:

nonresponse adjustments and calibration. Nonresponse is typically taken into account by modeling each survey unit's probability to respond and then multiplying the design weights with the inverse of the estimated response propensities. Calibration techniques rely on benchmarks known from other sources such as census data or large scale surveys such as the American Community Survey (U.S. Census Bureau, 2022a). These techniques can be used to adjust the survey weights in such a way that the survey weighted estimates will match the known benchmarks exactly. How these adjustment steps interfere with differential privacy has not been studied so far. (We are currently at an early stage of trying to address this problem.) However, both steps are data dependent, that is, they use information from the survey units for the adjustments. This implies that these steps cannot be ignored from a privacy perspective as the adjusted weights leak some personal information. Looking at the impacts on the sensitivity of the final statistic of interest (which uses the adjusted weights), similar problems as those discussed in the previous section will arise: changing one record in the database can potentially change the weight-adjustment factors for all other units in the survey. Thus, it seems imperative not to only account for these adjustment steps at the analysis stage. Better results in terms of the privacy-accuracy trade-off might be achieved if the weight-adjustment steps would be carried out in a differentially private way. More research is needed to better understand this trade-off. For example, it seems beneficial to identify robust adjustment strategies as less noise would be required to satisfy DP for these strategies.

In the particular case of post-stratification (which is a simple type of calibration), one such robust adjustment strategy has been proposed by Clifton et al. (2023). Another strategy would be to regularize the nonresponse and calibration weight adjustments. (This would be similar to the survey weight regularization strategy of Seeman et al., 2024, discussed in the previous section.)

6 DP and Imputation

All survey data are plagued by item nonresponse as survey respondents are often unwilling or unable to respond to all survey questions especially if they request sensitive information. A common strategy to deal with this problem is to impute the missing values before analyzing the data. Imputation is especially helpful if the response process is selective, that is if it is not missing completely at random as defined by Rubin (1976). In this case, using only the fully observed cases for the analysis would give biased results. However, imputations are always data dependent as they typically build a model based on the observed data and use this model to impute the missing values. As a consequence, the implications of imputation on the DP guarantees need to be considered regardless of whether or not the imputation procedure is included inside the data release mechanism. Some preliminary results for this problem are discussed in Das et al. (2022).

Similar to the problem of weighting adjustments, there are two possible strategies to account for

imputation under DP. The first strategy only considers the effects when analyzing the imputed data. The second strategy modifies the imputation routines to ensure that the imputations already satisfy DP. As Das et al. (2022) have shown, the first strategy implies that in the worst case the sensitivity increases linearly with the number of imputed observations. This substantial increase of the sensitivity arises because changing one record in the database can potentially impact all of the imputed values. Whether the worst case applies depends on the analysis of interest and on the selected imputation procedure. Still, for statistical agencies offering pre-imputed datasets for accredited researchers, this strategy is not an option since they cannot anticipate which analyses might be performed on the imputed data.

The second strategy can break the dependence of the sensitivity on the number of imputed records – at least for certain imputation strategies. The key requirement for breaking this dependence is that the imputation model m can be written as $D_{imp}^{(i)} \sim m(D_{obs}^{(i)}, \hat{\theta})$, where $D_{imp}^{(i)}$ and $D_{obs}^{(i)}$ contain the imputed and observed variables for record i and $\hat{\theta}$ denotes the model parameters estimated on the complete data. The model implies that, given $\hat{\theta}$, the imputed values of record i only depend on the observed values of that record and not on any other record. If these requirements are met and the parameters θ of the imputation model are estimated using any suitable differentially private mechanism with privacy loss parameter ε_1 , then, given any ε_2 -differentially private mechanism used for analyzing the data, the overall privacy loss is given by $\varepsilon_1 + \varepsilon_2$.

We note that the conditional independence assumption of the imputation model holds for many imputation methods, for example, parametric imputation models based on linear regression. However, it does not hold for hot-deck imputation, an imputation method commonly applied at statistical agencies.

7 Discussion

DP is theoretically intuitive and elegant. It provides quantifiable and composable guarantees of privacy protection (although these guarantees have been subject to some confusion and misinterpretation, see Tschantz et al., 2020). By putting data privacy on a mathematical basis, it has supplied a calculus for reasoning about the protection offered by sophisticated data release mechanisms.

Yet implementing DP mechanisms in practice often entails unforeseen complexities. In this paper we have focused on some of the complexities that arise in the context of survey data. Many of the same complexities can also emerge in settings with data preprocessing steps or with multistage data collection (such as national censuses). The goals of this paper are to draw attention to these complexities, review the current progress on addressing them, and spur renewed research activity to resolve those that remain outstanding.

We should emphasize at this point that we have explicitly focused on the complexities arising when using classical survey estimators, such as the Horvitz-Thompson estimator. These estimators all make use of the survey weights. However, there is an ongoing debate in economics about whether to include weights in econometric models (Faiella, 2010; Gelman, 2007; Magee et al., 1998; Solon et al., 2015). Many researchers tend to ignore survey weights and instead account for the sampling design and nonresponse by adding control variables to their models. Most of the complexities discussed in this paper (excluding the discussions on the increased sensitivity due to weighting in Sections 4 and 5) are still relevant in this context. Some of these complexities are further exacerbated when using complex econometric models, because computing the sensitivity for such models tends to be more difficult (as compared to simple statistics such as a Horvitz-Thompson estimator of a univariate mean). If we also want to account for earlier steps of the survey pipeline (e.g., nonresponse imputation or setting the sample size), calculating these models' sensitivity can become a major obstacle in practice.

We should also emphasize that in this paper we have only investigated the problem of releasing a single survey statistic, whereas in reality survey publications typically contain many outputs. For example, the American Community Survey asks respondents more than seventy questions, generates a dataset with over 64,000 variables, and produces more than 11 billion statistics annually (Jacobsen et al., 2023; U.S. Census Bureau, 2022b, 2025). This introduces an additional complexity in adopting DP for survey data. When a record contributes to multiple statistics, the overall privacy loss ε is typically equal to the sum of these statistics' individual privacy losses. This can quickly lead to either 1) very high levels of noise being added to all of the statistics; or 2) DP providing very weak – essentially meaningless – guarantees of protection. For example, given an overall privacy loss of $\varepsilon = 10$ (an upper limit for what can still be considered meaningful, see Near and Abuah, 2025), releasing 64,000 (unweighted) count statistics would require infusing each of these counts with noise that has variance 81 million. An alternative would be to add noise that has variance 100,000, but this would result in an overall $\varepsilon \approx 286$ – a value that is basically vacuous. (Setting the variance to be smaller would result in even larger values of the privacy loss parameter ε). While there are methods that achieve better tradeoffs between ε and the noise's variance (e.g., Abowd et al., 2022; Li et al., 2015; McKenna et al., 2018; Xiao et al., 2023), they cannot currently handle a data product of the scale and complexity of the American Community Survey, and adjusting these methods to fit with the complications of survey data is a nontrivial, unresolved problem.

Having identified a multiplicity of challenges in obtaining DP – some of which may be unduly constraining – we suggest that future research investigates pragmatic modifications to “completely-by-the-books” implementations of DP. The goals of such modifications should be: to provide a solution that is feasible to implement; to retain the essence of DP even while not strictly satisfying DP; and to not unduly sacrifice the accuracy of the released data, nor the privacy of the data

subjects, nor the resources of the statistical agency (in implementing the solution). Of course, any such modifications should be principled, in the sense that the associated risks to privacy are properly quantified and are outweighed by gains in data utility or implementability. Assessing the privacy risks of these modifications will likely involve a combination of theoretical and empirical analyses, and require measures of data privacy that lie outside the framework of DP.

An example of one possible modification is the non-DP publication of a data-dependent sampling design. A description of a survey's sampling design is crucial information for data users. Yet, as outlined in Section 2.2, if the sampling design was chosen with reference to the frame (as is often the case), then DP requires noise to be added to it before it can be published. Moreover, designing a DP mechanism to publish a sampling design will likely be difficult. On the other hand, it defies intuition that a simple description of a survey's sampling design should be disclosive of private information. This suggests it may be reasonable to modify the DP data release mechanism, allowing the sampling design to be released exactly (i.e. without noise) even while the other outputs are protected in line with the exact requirements of DP. But to justify this pragmatic violation of DP, the statistical agency should first address the questions: Can the risks associated with publishing a sampling design be quantified (without resorting to DP)? And when is it principled (in the sense given in the previous paragraph) to publish a sampling design as is, without privacy protection?

Acknowledgements

JB gratefully acknowledges partial financial support from the Australian-American Fulbright Commission and the Kinghorn Foundation. JD gratefully acknowledges support from Cooperative Agreement CB20ADR0160001 with the U.S. Census Bureau. The views expressed in this paper are those of the authors and not those of the Census Bureau or any other sponsor.

References

Abowd, J., Ashmead, R., Cumings-Menon, R., Garfinkel, S., Heineck, M., Heiss, C., Johns, R., Kifer, D., Leclerc, P., Machanavajjhala, A., Moran, B., Sexton, W., Spence, M., & Zhuravlev, P. (2022). The 2020 Census disclosure avoidance system TopDown Algorithm. *Harvard Data Science Review*, (Special Issue 2). <https://doi.org/10.1162/99608f92.529e3cb9>

Abowd, J. M., & Schmutte, I. M. (2019). An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, 109(1), 171–202. <https://doi.org/10.1257/aer.20170627>

Avella-Medina, M. (2021). Privacy-preserving parametric inference: A case for robust statistics. *Journal of the American Statistical Association*, 116(534), 969–983. <https://doi.org/10.1080/01621459.2019.1700130>

Bailie, J., & Drechsler, J. (2026). Whose data is it anyway? A formal treatment of differential privacy for surveys. In preparation. A working paper version was presented at the 2024 NBER conference *Data Privacy Protection and the Conduct of Applied Research: Methods, Approaches and Their Consequences*. https://conference.nber.org/conf_papers/f194306.pdf.

Bailie, J., Gong, R., & Meng, X.-L. (2026a). Differential privacy meets invariant statistics: Some conundrums in quantifying trade-offs. In R. Gong, V. J. Hotz, & I. M. Schmutte (Eds.), *Data privacy protection and the conduct of applied research: Methods, approaches and their consequences*. National Bureau of Economic Research. <https://doi.org/10.48550/arXiv.2504.15246>

Bailie, J., Gong, R., & Meng, X.-L. (2026b). *Privacy differentials in differential privacy* [Article in preparation]. Department of Statistics, Harvard University.

Bailie, J., Gong, R., & Meng, X.-L. (2026c). A refreshment stirred, not shaken: Invariant-preserving deployments of differential privacy for the U.S. Decennial Census. *Harvard Data Science Review*, (Special Issue 6). <https://doi.org/10.1162/99608f92.dab78690>

Balle, B., Barthe, G., & Gaboardi, M. (2018). Privacy amplification by subsampling: Tight analyses via couplings and divergences. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* (Vol. 31). Curran Associates, Inc. <https://papers.nips.cc/paper/paper/2018/hash/3b5020bb891119b9f5130f1fea9bd773-Abstract.html>

Bun, M., Drechsler, J., Gaboardi, M., McMillan, A., & Sarathy, J. (2022). Controlling privacy loss in sampling schemes: An analysis of stratified and cluster sampling. *Foundations of Responsible Computing (FORC 2022)*, 1:1–1:24.

Bun, M., Nissim, K., Stemmer, U., & Vadhan, S. (2015). Differentially private release and learning of threshold functions. *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, 634–649.

Casacuberta, S., Shoemate, M., Vadhan, S., & Wagaman, C. (2022). Widespread underestimation of sensitivity in differentially private libraries and how to fix it. *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 471–484. <https://doi.org/10.1145/3548606.3560708>

Clifton, C., Dajani, A. N., Hanson, E. J., Clark, S., Merrill, K., Merrill, S., & Rodriguez, R. (2023). *Preliminary report on differentially private post-stratification* (Working Paper No. ced-wp-2023-004) (<https://www.census.gov/library/working-papers/2023/adrm/ced-wp-2023-004.html>). U.S. Census Bureau.

Cummings, R., Desfontaines, D., Evans, D., Geambasu, R., Huang, Y., Jagielski, M., Kairouz, P., Kamath, G., Oh, S., Ohrimenko, O., Papernot, N., Rogers, R., Shen, M., Song, S., Su, W., Terzis, A., Thakurta, A., Vassilvitskii, S., Wang, Y.-X., ... Zhang, W. (2024). Advancing differential privacy: Where we are now and future directions for real-world deployment. *Harvard Data Science Review*, 6(1). <https://doi.org/10.1162/99608f92.d3197524>

Das, S., Drechsler, J., Merrill, K., & Merrill, S. (2022). *Imputation under differential privacy*. ArXiv. <https://doi.org/10.48550/arXiv.2206.15063>

Drechsler, J. (2023). Differential privacy for government agencies—Are we there yet? *Journal of the American Statistical Association*, 118(541), 761–773. <https://doi.org/10.1080/01621459.2022.2161385>

Dwork, C., Kohli, N., & Mulligan, D. (2019). Differential privacy in practice: Expose your epsilons! *Journal of Privacy and Confidentiality*, 9(2). <https://doi.org/10.29012/jpc.689>

Dwork, C., & Lei, J. (2009). Differential privacy and robust statistics. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing - STOC '09*, 371–380. <https://doi.org/10.1145/1536414.1536466>

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In S. Halevi & T. Rabin (Eds.), *Proceedings of the Third Theory of Cryptography Conference, TCC 2006* (pp. 265–284, Vol. 3876). Springer. https://doi.org/10.1007/11681878_14

Faiella, I. (2010). *The use of survey weights in regression analysis* (Temi di discussione (Working papers) No. 739). Bank of Italy. <https://www.bancaditalia.it/pubblicazioni/temi-di-discussione/2010/2010-0739/index.html>

Gelman, A. (2007). Struggles with survey weighting and regression modeling. *Statistical Science*, 22(2), 153–164. <https://doi.org/10.1214/088342306000000691>

Horvitz, D. G., & Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260), 663–685.

Hsu, J., Gaboardi, M., Haeberlen, A., Khanna, S., Narayan, A., Pierce, B. C., & Roth, A. (2014). Differential privacy: An economic method for choosing epsilon. *Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium*, 398–410. <https://doi.org/10.1109/CSF.2014.35>

Jacobsen, L. A., Mather, M., & Reamer, A. (2023). *America's essential economic and social data at risk: A vision to preserve and enhance the American Community Survey*. The Census Project. <https://thecensusproject.org/wp-content/uploads/2023/05/census-acrs-report-2023-v4.pdf>

Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., & Smith, A. (2011). What can we learn privately? *SIAM Journal on Computing*, 40(3), 793–826. <https://doi.org/10.1137/090756090>

Kifer, D., Abowd, J. M., Ashmead, R., Cumings-Menon, R., Leclerc, P., Machanavajjhala, A., Sexton, W., & Zhuravlev, P. (2022). *Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 Census*. ArXiv. <https://doi.org/10.48550/arXiv.2209.03310>

Li, C., Miklau, G., Hay, M., McGregor, A., & Rastogi, V. (2015). The Matrix Mechanism: Optimizing linear counting queries under differential privacy. *The VLDB Journal*, 24(6), 757–781. <https://doi.org/10.1007/s00778-015-0398-x>

Magee, L., Robb, A. L., & Burbidge, J. B. (1998). On the use of sampling weights when estimating regression models with survey data. *Journal of Econometrics*, 84(2), 251–271. [https://doi.org/10.1016/S0304-4076\(97\)00086-9](https://doi.org/10.1016/S0304-4076(97)00086-9)

McKenna, R., Miklau, G., Hay, M., & Machanavajjhala, A. (2018). Optimizing error of high-dimensional statistical queries under differential privacy. *Proceedings of the VLDB Endowment*, 11(10), 1206–1219. <https://doi.org/10.14778/3231751.3231769>

Nanayakkara, P., Smart, M. A., Cummings, R., Kaptchuk, G., & Redmiles, E. M. (2023). What are the chances? Explaining the epsilon parameter in differential privacy. *Proceedings of the 32nd USENIX Conference on Security Symposium*, 1613–1630. <https://www.usenix.org/system/files/usenixsecurity23-nanayakkara.pdf>

National Academies of Sciences, Engineering, and Medicine. (2024). *A roadmap for disclosure avoidance in the Survey of Income and Program Participation*. The National Academies Press. <https://doi.org/10.17226/27169>

Near, J. P., & Abuah, C. (2025). *Programming differential privacy*. <https://programming-dp.com/>

Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3), 581–592.

Seeman, J., Si, Y., & Reiter, J. P. (2024). *Differentially private finite population estimation via survey weight regularization*. ArXiv. <https://doi.org/10.48550/arXiv.2411.04236>

Seeman, J., & Susser, D. (2024). Between privacy and utility: On differential privacy in theory and practice. *ACM Journal on Responsible Computing*, 1(1), 3:1–18. <https://doi.org/10.1145/3626494>

Solon, G., Haider, S. J., & Wooldridge, J. M. (2015). What are we weighting for? *The Journal of Human Resources*, 50(2), 301–316. <https://www.jstor.org/stable/24735988>

Tschantz, M. C., Sen, S., & Datta, A. (2020). SoK: Differential privacy as a causal property. *2020 IEEE Symposium on Security and Privacy (SP)*, 354–371. <https://doi.org/10.1109/SP40000.2020.00012>

U.S. Bureau of Labor Statistics. (2018). Current Population Survey: Design. <https://www.bls.gov/opub/hom/cps/design.htm>

U.S. Census Bureau. (2018). Protecting the confidentiality of America's statistics: Adopting modern disclosure avoidance methods at the Census Bureau. https://www.census.gov/newsroom/blogs/research-matters/2018/08/protecting_the_confi.html.

U.S. Census Bureau. (2022a). *American Community Survey and Puerto Rico Community Survey design and methodology* (Version 3.0). https://www2.census.gov/programs-surveys/acs/methodology/design_and_methodology/2022/acs_design_methodology_report_2022.pdf

U.S. Census Bureau. (2022b). *Census API: Datasets in /data/2022 and its descendants*. <https://api.census.gov/data/2022.html>

U.S. Census Bureau. (2022c). Disclosure avoidance protections for the American Community Survey. <https://www.census.gov/newsroom/blogs/random-samplings/2022/12/disclosure-avoidance-protections-acs.html>

US Census Bureau. (2025). *The American Community Survey questionnaire - informational copy (2025)*. <https://www2.census.gov/programs-surveys/acs/methodology/questionnaires/2025/quest25.pdf>

Wang, Y.-X., Balle, B., & Kasiviswanathan, S. P. (2019). Subsampled Rényi differential privacy and analytical moments accountant. In K. Chaudhuri & M. Sugiyama (Eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics* (pp. 1226–1235, Vol. 89). PMLR. <https://proceedings.mlr.press/v89/wang19b.html>

Wang, Y.-X., Lei, J., & Fienberg, S. E. (2016). Learning with differential privacy: Stability, learnability and the sufficiency and necessity of ERM principle. *Journal of Machine Learning Research*, 17(1), 6353–6392. <https://dl.acm.org/doi/10.5555/2946645.3053465>

Xiao, Y., He, G., Zhang, D., & Kifer, D. (2023). An optimal and scalable Matrix Mechanism for noisy marginals under convex loss functions. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, & S. Levine (Eds.), *Advances in Neural Information Processing Systems* (pp. 20495–20539, Vol. 36). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2023/file/414f4c9fe9653e5de98fad6964d50315-Paper-Conference.pdf