

# Differential Privacy for Government Agencies – Are We There Yet?

Jörg Drechsler

Institute for Employment Research (IAB), Germany

University of Maryland, College Park, USA

## Abstract

Government agencies always need to carefully consider potential risks of disclosure whenever they publish statistics based on their data or give external researchers access to the collected data. For this reason, research on disclosure avoiding techniques has a long tradition at statistical agencies. In this context, the promise of formal privacy guarantees offered by concepts such as differential privacy seem to be the panacea enabling the agencies to exactly quantify and control the privacy loss incurred by any data release. Still, despite the excitement in academia and industry, most agencies—with the prominent exception of the U.S. Census Bureau—have been reluctant to even consider the concept for their data release strategy.

This paper aims to shed some light on potential reasons for this. We argue that the requirements when implementing differential privacy approaches at government

agencies are often fundamentally different from the requirements in industry. This raises many challenging problems and open questions that still need to be addressed before the concept might be used as an overarching principle when sharing data with the public. The paper will not offer any solutions to these challenges. Instead, we hope to stimulate some collaborative research efforts, as we believe that many of the problems can only be addressed by inter-disciplinary collaborations.

## 1 Introduction

The concept of differential privacy has gained substantial attention in recent years as the only standing approach offering formal privacy guarantees, which hold irrespective of the assumptions regarding the background knowledge of a potential attacker. Since the seminal paper by Dwork *et al.* was published in 2006, thousands of papers mostly from the computer science community have addressed the topic from various perspectives. Given its roots in theoretical computer science, it is perhaps not surprising that most of these works approached the problem from a theoretical perspective. While new algorithms that satisfy the differential privacy requirements for specific analysis tasks are proposed almost every other day, their impacts on the accuracy of the obtained results is typically only evaluated analytically by looking at measures such as the maximum expected error under asymptotic regimes (that is, assuming  $n \rightarrow \infty$ ). These metrics offer only limited insights regarding the practical impacts for commonly applied analyses tasks. Besides, evaluations based on small sample sizes are sparse, although it is well understood that the relative advantages of different algorithms

crucially depend on the size and structure of the available data (see Alabi *et al.* (2020) for an illustration in the context of linear regression). Furthermore, problems that commonly arise in practice such as complex data structures, missing data, data cleaning, etc. tend to be ignored.

Although practical experience has been relatively limited so far, the concept of differential privacy has been embraced by the industry in recent years. Many companies, especially from the tech industry, such as Google (Erlingsson *et al.*, 2014), Apple (Apple’s Differential Privacy Team, 2017), Microsoft (Ding *et al.*, 2017), Facebook (Messing *et al.*, 2020) or Uber (Uber Security, 2017) have deployed the concept for some of their products or are currently conducting research with the aim of adopting the approach in the future.

Despite the excitement in academia and industry, the enthusiasm at government agencies and national statistical organizations (NSOs) has been limited so far. While some agencies explored the feasibility of the approach in limited settings (Soria-Comas and Drechsler, 2013; Bailie and Chien, 2019), the only large-scale deployment of the approach for many years was OntheMap, a graphical interface offered by the U.S. Census Bureau visualizing commuting patterns in the United States. The underlying data are protected using an algorithm which satisfies  $(\epsilon, \delta)$ -probabilistic differential privacy (Machanavajjhala *et al.*, 2008).

Recently, the U.S. Census Bureau also announced that it will adopt differential privacy for the decennial Census 2020 (Abowd, 2018). Compared to most other data products gathered at NSOs, which are based on surveys with limited sample sizes and hundreds of variables, protecting data from the decennial Census seems to be a straightforward task:

it contains more than one hundred million records and only asks seven questions. Still, the fact that a research team of computer scientists and statisticians has been working on this problem for several years now and the concerns regarding the accuracy of the results that were raised after results from a test run of the algorithm using 2010 Census data were released (Ruggles *et al.*, 2019; Wezerek and Van Ripper, 2020; Van Ripper *et al.*, 2020; Hallowell and Rector, 2020) illustrate the difficulties when trying to implement the ideas in practice. One of the key challenges which distinguishes the decennial Census from previous deployments in the tech industry is that (accurate) answers are required on a very detailed geographical level, which implies that despite its large size many of the released statistics are based on a very limited number of cases.

Many additional problems arise because the requirements when implementing differential privacy approaches at government agencies are fundamentally different from the requirements in industry: The amount of data is much more limited, the data should be available for many years, results should be reproducible, users of the data are typically interested in making inferences regarding a specific target population, agencies are not the final users of the data, incentives for sharing the data are virtually non-existent, etc. All these aspects need to be taken into account when considering whether the concept might be a viable approach for solving the ever existing dilemma between confidentiality protection and broad access to the data.

This paper is not meant to provide a road map for how to implement differential privacy at government agencies. Instead, it will highlight some important aspects that need to be on

everybody's radar and open questions that still need to be addressed when thinking about if and how the concept could be applied in the government context.

The remainder of the paper is structured as follows: before discussing the different aspects that need to be considered in the government context, Section 2 offers a short review of differential privacy and some of its properties. Section 3 discusses the challenges arising for the typical data products at government agencies with limited sample sizes but very detailed information. It also illustrates why neither the query response system, for which differential privacy was originally developed, nor restricted access to the unprotected data for accredited researchers seem to be an option when adopting differential privacy for government agencies. Section 4 addresses the difficulties in anticipating the impacts of differential privacy on the accuracy of the obtained results and highlights that understanding the level of protection provided through differential privacy is not trivial. Differential privacy in the survey context is the focus of Section 5, which illustrates that the approach might have negative effects on response rates and discusses the open questions regarding the interaction between differential privacy and common data processing steps such as weighting and imputation and the challenges when trying to account for the data protection procedures when making inference to the underlying population. Section 6 addresses the difficulties that government agencies face, when setting the privacy parameter  $\varepsilon$ . The paper concludes with some suggestions how to address the challenges and open questions raised, advocating for more interdisciplinary research to tackle these problems.

## 2 A brief review of differential privacy

This section will only review the definition and properties of  $\epsilon$ -differential privacy as originally proposed by Dwork *et al.* (2006). Various relaxations of the initial concept—most notably  $(\epsilon, \delta)$ -differential privacy—have been proposed in the literature in the meantime. Since the subtleties of these variants are irrelevant for the discussions in the remaining paper, we refer interested readers to Dwork *et al.* (2014) and Vadhan (2017) for more detailed introductions to differential privacy, which also cover (some of the) relaxations of  $\epsilon$ -differential privacy.

The concept of differential privacy was originally developed for a query response system. Such a system accepts specific types of queries as input, for example a counting query or a query for the mean of a variable and then returns a noisy answer to the query, where the noise is calibrated to ensure that the requirements of differential privacy are met. The user of the system never accesses the underlying data directly, she will only see the noisy answer to the query. Under this setting, differential privacy guarantees that the influence that any record can have on the reported output is strictly limited. This ensures that the information that can be learned about any individual in the database is also limited. These guarantees are formalized by bounding the difference of the probability distribution of the query response when changing one record in the data.

**Definition 1** ( $\epsilon$ -differential privacy, Dwork *et al.* (2006)). A randomized mechanism  $\mathcal{M}$  gives  $\epsilon$ -differential privacy if, for all neighboring datasets  $D, D'$ , and all events  $S \subset \text{Range}(\mathcal{M})$

$$P(\mathcal{M}(D) \in S) \leq e^\epsilon P(\mathcal{M}(D') \in S). \quad (1)$$

Depending on the context, two datasets are called neighboring, if one could be obtained from the other by (a) adding or removing a single record (unbounded differential privacy) or by (b) changing the values of one record, while keeping the size of the database fixed (bounded differential privacy).

Differential privacy ensures that the probability of observing a specific output if database  $D$  is used as the input is never more than  $e^\epsilon$  times (and never less than  $e^{-\epsilon}$  times) the probability of observing the same output if database  $D'$  is used as the input, where the probability distribution is based on the randomness induced by the mechanism  $\mathcal{M}$ . The parameter  $\epsilon$  can be used to specify the level of protection. Larger values of  $\epsilon$  allow for larger differences in the output distribution between two neighboring databases, thus offering lower levels of protection. However, larger values of  $\epsilon$  will typically increase the level of accuracy of the reported output as the mechanism  $\mathcal{M}$  will need to introduce less noise to satisfy Equation (1). Thus,  $\epsilon$  can be seen as a tuning parameter that treats privacy for empirical accuracy of the obtained estimate.

Making this ever existing trade-off explicit is an important advantage of the concept compared to currently implemented data access regulations. Any data protection strategy, from aggregating geographic information, over top-coding, or swapping sensitive values to not releasing the data at all, makes implicit decisions on treating accuracy against privacy. However, the exact trade-off is more difficult to understand and to control under current regimes. With differential privacy it would be possible (at least in theory), to select the level of accuracy and privacy that is optimal for the society. Furthermore, since the protection

strategy is fully transparent, that is the selected value of  $\varepsilon$  and any information about the algorithm used to protect the data can be publicly released, it would in principle be possible to take the uncertainty induced by the algorithm into account when analyzing the protected data. Still, finding the optimal value for  $\varepsilon$  and integrating the extra uncertainty in downstream analyses are daunting tasks in practice (see also the discussions in Sections 6 and 5.2).

Differential privacy offers two additional attractive properties that will be relevant for the discussions in the remainder of this paper:

- **Postprocessing:** Differential privacy is immune to postprocessing meaning that any function of an output satisfying  $\varepsilon$ -differential privacy also satisfies differential privacy with the same level of  $\varepsilon$ .
- **Composition:** If the same database is used to answer  $K$  queries in a differentially private manner, each with its own privacy parameter  $\varepsilon_k$ ,  $k = 1, \dots, K$ , the overall privacy loss is bounded by  $\varepsilon = \sum_{k=1}^K \varepsilon_k$ .

These properties have several important implications in the government context: The first property guarantees that both the data producer as well as the user of the data can arbitrarily alter the output without increasing the privacy loss. This implies that the data disseminating agency can adjust the differentially private output to avoid releasing implausible values such as negative age values and to improve the utility. The agency also does not have to worry that a malicious user of the data might be able to learn sensitive information by manipulating the received information in a clever way. This property holds even if the output is combined



with information from other data sources.

The second property allows one to split the overall privacy budget  $\varepsilon$  across multiple queries. In principle this would allow one to assign individual privacy budgets to different users of the data. The users could then decide how to best spend their budget to obtain the information they are interested in (but see the caveats discussed in Section 3.2 below).

We refer the interested reader to Dwork *et al.* (2014) for an in-depth discussion of the properties of differential privacy including several examples of mechanisms that could be used to meet the requirements of differential privacy for various analysis tasks.

### 3 Data availability and access

Most companies that already ensure some form of formal privacy when collecting their data are from the tech industry. While these companies cannot rely on a trusted data curator, which means that the data need to be protected before they are transferred to the company (local differential privacy), they all collect massive amounts of data every day. This offers two advantages: First, for certain types of queries such as counting queries the amount of uncertainty that needs to be added to ensure differential privacy is only related to the sensitivity of the statistic of interest and not to the size of data. For example, the same amount of noise needs to be added to protect a frequency table irrespective of whether the table is based on 20 records or 10 million records. However, adding, say, a noise value of 3 to a cell count of 5 will obviously have a stronger impact on any findings based on this table than adding the same value to a cell count of 3.5 million. In other words, the signal clearly

dominates the noise in the latter case, but not in the first case. Second, the data become outdated and irrelevant quickly. With so much data being collected everyday, there is no need to still look at yesterday's data. This also implies that these companies can basically start with a new privacy budget regularly (there is a caveat as the information collected over time might be correlated, which will spill additional information, but it seems that this issue is currently ignored at least in some of the deployments (Tang *et al.*, 2017)).

### **3.1 Data collected at government agencies**

While large scale administrative data and other massive data sources such as satellite images or mobile phone data are being exploited more extensively in recent years, most of the collected information at government agencies is still based on surveys. Most of these surveys comprise less than 100,000 records and often collect very detailed information on every unit included in the database, making most records unique in terms of their value combinations across the variables. At the same time, interest in the data degrades at a much lower rate. The fact that the U.S. Census Bureau takes the effort of making Census data publicly available 72 years after the initial data collection indicates that there is still interest in these data even more than 70 years after they were collected. However, this implies that the privacy budget, that is, the level of  $\varepsilon$  spent for all information that is released, needs to protect the data over all these years.

## 3.2 Query response system not an option

As pointed out above, the concept of differential privacy was originally developed for a query response system. Such systems are not uncommon at statistical agencies, where they typically firm under the name of remote analysis servers (RAS). With RAS the user will never see the actual microdata. Instead she will define the analysis of interest by selecting the type of analysis to be performed and the variables to be included from various drop down menus. The RAS will compute the analysis on the underlying microdata but will only return the (potentially perturbed) results. Examples include the new dissemination platform, Microdata Access (<https://data.census.gov/mdat/#/>; the successor of the DataFerret), which allows user specified tabulations or the Cross-National Data Center in Luxembourg ([www.lisdatacenter.org](http://www.lisdatacenter.org)), which additionally allows specifying simple regression models.

In theory, addressing the fact that multiple queries are answered by a RAS over time is straightforward in the context of differential privacy due to the composition property described above. If the overall privacy loss deemed to be acceptable is  $\varepsilon$  and it is known that  $k$  queries should be answered, one option would be to assign  $\varepsilon/k$  of the privacy budget to each query. However, in practice a dynamic query system in which many queries are submitted by different users over a long period of time raises many challenging problems. For example, distributing the privacy budget equally across queries is not necessarily the optimal or even the fairest solution. Two queries spending the same level of  $\varepsilon$  might have completely different signal-to-noise ratios depending on the type of query and the sample size that is used for analysis. For example, to achieve the same level of  $\varepsilon$ , much more noise

needs to be added when estimating the mean compared to estimating frequency counts. Furthermore, it seems questionable to spend the same privacy budget on some intermediate analysis during data preparation as on estimating politically important measures such as the poverty rate, which will impact the allocation of billions of dollars. Thus, deciding how to split the privacy budget among multiple queries is a daunting task.

Still, these discussions are based on the assumption that all queries are known in advance. However, the main motivation for providing broad data access for the scientific community and the general public is the understanding that statistical agencies cannot anticipate all potential research questions for which the database might provide useful insights. If this would be the case, the agency could run all these analyses once the data have been collected and publish all the results (potentially under the constraints of differential privacy). Under this assumption no external access would be required. But once we accept that users will submit unforeseen queries, the difficult question to answer is: how much of the privacy budget should the agency hold back to be able to answer these questions? If the agency is too restrictive, answers to the submitted queries will be unnecessarily inaccurate, but if the agency is too generous, there will be a point at which all the privacy budget has been spent. At this point the system has to be shut down forever and no one can ever be allowed to access the data again. This first-come-first-served practice would be unacceptable for several reasons. Obviously, this strategy would setup an incentive to submit queries as early as possible. This will increase the risk of sloppy analyses as researchers might be tempted to avoid spending time and privacy budget on robustness checks and careful evaluations of the

modeling assumptions. Furthermore, important research questions might only emerge at a time when all the budget has already been spent. Besides, reproducibility research to validate earlier findings would be impossible. Especially, since careful reproducibility evaluations require to evaluate the results under various assumptions to check the robustness of the findings. This could imply that reproducibility checks might require more privacy budget than the initial research. And finally, given that most of the data hoisted by statistical agencies are collected spending taxpayer's money, it will be difficult to sell the idea that know one will ever be allowed to access the data again after the privacy budget has been spent.

Given all these arguments, it seems that the only sensible strategy for statistical agencies willing to adopt the approach will be to generate a differentially private copy of the microdata and disseminate these synthetic microdata to the public. Given that differential privacy is immune to postprocessing the users would be able to run as many analyses as they want on the released microdata without violating the differential privacy conditions. The only crucial requirement would be that the original data would never be touched again once the synthetic microdata have been generated (unless some of the privacy budget has been withhold for validation). We note that generating differentially private microdata is also the strategy followed by the U.S. Census Bureau for the decennial Census. Still, intuitively, protecting the entire microdata (input perturbation) is more difficult than protecting only specific statistics (output perturbation). Thus, the answers obtained from the differentially private microdata will necessarily be noisier than the answers that could be obtained if an

optimal differentially private algorithm could be used to answer all the queries of interest. Especially with the high dimensional survey data discussed in the introduction, the obtained answers from the synthetic data might be so noisy that they would no longer offer any useful insights. It would certainly be possible to develop synthesizers that provide very accurate answers for specific types of queries, but in order to achieve this, these queries would need to be known in advance leading to the same problems already discussed for the query response system.

### **3.3 Problems with tiered access for accredited researchers**

Some researchers argue that a potential solution to circumvent the difficulties of generating useful differentially private synthetic data could be that trusted researchers could still access the unprotected microdata at the premises of the statistical agency and only need to ensure that any result that they publish satisfies the requirements of differential privacy. This approach is in the spirit of current data access regulations at many statistical agencies. To facilitate access, many NSOs established research data centers (RDCs). Compared to the public use microdata samples that are disseminated to the public, more detailed and less protected datasets are typically available at the RDCs. Access is only granted to accredited researchers, the researchers cannot bring any own devices to the RDC, are often monitored while working at the RDC, and any research output can only be used outside the RDC after the output has been carefully scrutinized by employees of the RDC to ensure that it does not violate any confidentiality constraints.

Maintaining such a tiered access in the context of differential privacy would offer two important advantages: Fitting the final model of interest is typically only a very small part of an applied research project. Most of the work goes into data preparation and data cleaning, checking of model assumptions, etc. All these steps could be performed without spending privacy budget leaving more of the budget to get more accurate results for the final analysis of interest. Furthermore, since the final model is known, a tailor-made algorithm could be developed for producing the protected output, which optimizes the trade-off between accuracy and data protection.

However, there are some critical open questions that would need to be solved first, if the approach should be implemented in the future: First, the strategy requires that a privacy budget will be assigned to each researcher. This budget would be used to protect the results that should be released to the public. This leads to the dilemma outlined above that a decision needs to be made, how the privacy budget should be distributed across researchers. Second, in the current debate on the reproducibility crises, this approach seems to be a step in the wrong direction. If all the privacy budget that has been reserved for this research project is spent on protecting the final result, there is no way that the results could be verified in an independent reproducibility study. Of course, this problem could be mitigated by always reserving some privacy budget for reproducibility purposes, but this will negatively affect the accuracy of the publishable results and deciding how much of the budget needs to be assigned for reproducibility is a challenging question. Finally, the researcher needs to be trusted not to reveal any information about the data beyond the protected output.

This might not seem like a major obstacle in most situations. After all, researchers are typically not interested in revealing any information beyond their research results and the common practice of imposing heavy fines if personally identifiable information is purposefully released is a strong incentive to follow the protocol. However, there will be situations in which researchers will be tempted to violate the protocol. Imagine the following scenario: After spending several months cleaning and preparing the data, extensive data exploration, careful model evaluations and several robustness checks, the researcher is finally ready to run the final model of interest and is overjoyed to see that the results strongly support the research hypothesis. Given the novelty and importance of the results she is confident to publish the findings in one of the leading journals in the field. All that is needed now is to obtain a differentially private version of the final model (admittedly, this is already a simplifying assumption as major journals would also require proof for the robustness of the findings). Given the high stakes, the researcher decides to collaborate with some experts on differential privacy to come up with an algorithm specifically tailored to the final model to ensure that the error introduced is as small as possible. However, when run on the data, the differentially private results are so different from the results based on the original data that they would contradict the null hypothesis. As the concept of differential privacy necessarily requires some randomness when moving from the original output to the protected output, such an unlucky outcome could occur even for carefully designed algorithms and research output with low sensitivity. The question is whether the researcher would accept this outcome and silently bury all hopes of publishing the findings in a major journal. The problem arises



because the researchers only have one shot to obtain the final protected results. If they end up with an unlucky draw from the sanitization mechanism, they will have to accept the outcome. It seems questionable that researchers would accept such a lottery in practice.

Of course there is an analogy to sampling: even with carefully designed random samples, we can end up with a sample that is not representative of the population. However, the major difference is that we will never know whether this is the case or not. All we can do is to try to keep the sampling error as small as possible. It is a completely different story to publish results for which we already know that they are biased.

Establishing such a system despite these problems might actually violate the privacy guarantees in the long run. As results would likely only be published if they do not deviate too much from the results based on the original data, an attacker would know that the bounds of the noise that was introduced are tighter than the bounds that would actually be required for the selected level of the privacy parameter  $\epsilon$ .

One possible strategy to circumvent this problem would be to offer differentially private synthetic data to the accredited researcher for her exploratory analysis and to only compute the final model on the original data. Still, working only with synthetic data would likely not be an option for most researchers. Furthermore, parts of the privacy budget would need to be spent for generating the synthetic data, which would mean less budget for the accredited researchers.

## 4 Understanding the privacy guarantees and impacts on accuracy

Part of the attractiveness of the differential privacy approach lies in its intuitiveness. DP ensures that the influence that a single record can have on the reported output is strictly limited. This implies that the information that can be revealed about a single individual is also strictly limited. But understanding the impacts on the accuracy of the obtained results is more difficult and even the interpretation of the level of protection that is provided for a fixed level of  $\varepsilon$  is challenging.

### 4.1 Impacts on accuracy

Differential privacy has been criticised repeatedly for its strong negative impacts on the accuracy of the obtained results (Fienberg *et al.*, 2010; Bambauer *et al.*, 2013). Bounding the risks in the worst case scenario, that is, under the assumption that the attacker already knows the information for each record in the database except for one record, necessarily requires strong protection mechanisms. Still, there are many applications in which the approach can provide useful insights. However, assessing the impacts in advance is typically only possible for relatively simple algorithms such as the Laplace mechanism. With these algorithms the additional uncertainty that is introduced to protect the data can be quantified directly and thus the extra uncertainty can also be taken into account at the analysis stage (see also the discussion on statistical inference below). Other mechanisms introduce the randomness that

is necessary to fulfill the differential privacy requirements in a way that makes it difficult to assess the impacts on any subsequent analysis. For example the popular differential privacy stochastic gradient decent method (SGDM)(Abadi *et al.*, 2016)—a differentially private version of a Generative Adversarial Network (GAN)—adds noise to the gradients of the neural network and also truncates their range. The effects on downstream analysis are difficult to assess analytically.

The problem is magnified in the context of government agencies, because as outlined in Section 3.2, government agencies will typically have to provide differentially private synthetic datasets. Most of the algorithms that have been developed so far to generate synthetic data, such as the SGDM), fall into the second category of hard to quantify noise infusion techniques. Furthermore, the future analyses that will be performed on the synthetic data are typically unknown, making it impossible to anticipate the effects on accuracy and to optimize the algorithm to minimize these impacts. Finally, post-processing steps that are commonly employed to ensure that the protected data fulfill consistency requirements and do not contain implausible values such as negative age values also effect the accuracy of the results. The optimization procedures that are commonly applied to find solutions close to noisy result under consistency and non-negativity constraints again introduce biases in ways that are hard to anticipate and difficult to control for in any analysis using the protected data.

## 4.2 Privacy guarantees

While the general interpretation of the privacy guarantees offered by differential privacy is intuitive, the risk implications of specific values of  $\epsilon$  are much harder to understand. What are the actual risks for Bob, if the probability that the algorithm produces a specific noisy answer to the query is, say, ten times more likely if he joins the database? Should he be worried? To make it more concrete, consider the following simplified example. Assume a survey only asks about HIV status and plans to release the percentage of respondents that reported a positive status using an algorithm that ensures that the release satisfies the requirements of differential privacy. If  $\epsilon = 2.3$ , Bob knows that the probability that the released value equals any specific value  $k$  if he participates in the survey, is never more than 10 times (since  $e^\epsilon \approx 10$ ) the probability of releasing the same value if he would not participate. This holds for any value of  $k$  including the true value. But how can he link this information to the risk that his own HIV status could be revealed based on the released information? Understanding the impacts of changes in the value of  $\epsilon$  is even more difficult. How much of Bob's privacy is lost, if  $\epsilon$  is changed from 2.3 to 4.6? The probability ratio for observing a certain output with or without Bob increases from 10 to 100, but what does that mean for Bob? Yet, fully understanding the impacts of changes of  $\epsilon$  are vitally important if a value of  $\epsilon$  should be found that optimally addresses the trade-off between accuracy and privacy.

When discussing the concept of differential privacy with the general public it is also important to emphasize that guaranteeing differential privacy will not be enough to ensure

that the data are protected. The level of protection will obviously depend on the selected value of  $\varepsilon$  but also on the selected mechanism to achieve differential privacy.

A simplified example will again help to illustrate this point. Assume a database contains five categorical variables, two of them are binary and three of them have three categories. The agency decides to use the geometric mechanism (Ghosh *et al.*, 2009), which is also the workhorse in the current version of the TopDown mechanism to be used for the decennial Census 2020 (we ignore any hierarchical data structures or postprocessing steps for simplicity). The algorithm consists of three steps. In the first step, all variables are fully cross-classified resulting in a contingency table with up to  $2^2 \cdot 3^3 = 108$  cells (structural zeros, that is, implausible value combinations can be dropped). In the second step, random noise from a two-sided geometric distribution is added independently to each of the cells of the contingency tables. In the final step, the noisy table is turned back into microdata to be released (potentially using some postprocessing to deal with negative counts in the noisy table). The agency decides to set  $\varepsilon = 8.6$  as this is the value used in the only existing differentially private data release – the OntheMap visualization tool of the U.S. Census Bureau. Assuming unbounded differential privacy, this implies that the probability that any particular cell remains unchanged, that is, that a value of zero is added to the true count, is 99.963%. Since noise is added independently to each cell, this implies that with a probability of more than 96.1% the released data would match the original data in every single record (subject to random shuffling of the records). Thus, a potential attacker who recognizes a record based on some of the attribute values can be very confident that the

remaining, potentially sensitive attributes will still contain the true values for this record. This holds irrespective of how many records the database contains.

To be fair, these probabilities decrease quickly with the number of variables contained in the database, postprocessing steps which are common in practice will further reduce these probabilities, and such large values of  $\varepsilon$  are generally not recommended. The point is that it is important to consider the context and to be aware that differential privacy does not automatically guarantee an acceptable level of protection unless a sufficiently small value of  $\varepsilon$  is selected.

A final point that is important for statistical agencies is the fact that ensuring differential privacy even with very low values of  $\varepsilon$  will not guarantee that the data release does not cause any harm. Consider the scenario, in which the released noisy data reveal that the risk of getting a specific type of cancer is substantially increased in certain geographic areas (even when accounting for the extra uncertainty from noise infusion). As a consequence, some health insurance companies may increase the premium or decide not to offer any insurance for customers living in this area. Obviously, the concept of differential privacy cannot be blamed for this adverse event, since the approach was never meant to protect against such negative outcome. In fact, one of the fundamental underpinnings that motivated the development of differential privacy was the understanding that some information could be revealed from a database about Terry Gross, no matter, if Terry Gross decides to join the database or not and that there is no reasonable way of preventing this from happening (Dwork, 2006). However, unless there are regulations that prevent such events, statistical agencies still need

to consider the potential harm they could inflict on parts of the society by publicly releasing aggregate statistics from their collected data, even if these statistics fulfill the requirements of differential privacy.

## 5 Differential Privacy in the Survey Context

Most information obtained by statistical agencies is collected through surveys which typically comprise only a small sample from the underlying population of interest. Thus, it will be crucial to understand the impacts of differential privacy for survey data. Some open questions for this context will be discussed in this section.

### 5.1 Incentives to share the data

In all successful deployments of differential privacy in the industry context, the data providers have strong incentives to share their data, as they will benefit from the services offered by these companies in exchange for the data. Furthermore, the well known privacy paradox that has been confirmed in several independent studies shows that although users express concerns about their privacy, their behaviour does not reflect these concerns (Taddicken, 2014). The desire to use the products seems to outweigh any privacy concerns.

The situation is different in the survey context. Although there are obvious benefits for society from the insights obtained from the collected data, the direct benefits for the survey respondent are less obvious. Although many surveys offer a small monetary compensation for respondents, the offered amounts are low to avoid introducing bias in the collected data

because the incentive might work better for some subgroups. As a consequence, statistical agencies are facing constantly increasing nonresponse rates. While response rates above sixty percent were still achievable with carefully mounted surveys in 1997, these rates dropped to 22% by 2012 according to a study conducted by the PEW Research Center (Kohut *et al.*, 2012).

At first sight, offering strong formal privacy guarantees to the respondents will improve the quality of the collected data as more respondents might be willing to participate and to respond truthfully, if they know that their confidentiality is guaranteed. However, this only holds if the respondents fully understand and trust the protection mechanisms. Previous experience shows that this is not always the case. Randomized response proposed by Warner (1965), which coincidentally happens to be the earliest protection mechanism satisfying the requirements of differential privacy, protects answers to sensitive questions by introducing randomness into the response process. In its basic form, the respondent flips a coin and depending on the result of the coin flip either answers the sensitive question or provides an answer to a completely different non-sensitive question. Since only the respondent knows whether the reported answer is the answer to the sensitive or non-sensitive question her sensitive information is protected. However, several studies have shown that introducing the concept of randomized response in a survey does not increase response rates or the likelihood that respondents answer truthfully (Edgell *et al.*, 1982; Landsheer *et al.*, 1999; Coutts and Jann, 2011; Kirchner, 2015). It seems that respondents do not trust the concept enough to reveal the true answers. See also Oberski and Kreuter (2020) for a similar



argument.

On the other hand, guaranteeing differential privacy can undermine the willingness to participate. Survey institutes often try to motivate potential participants by emphasizing the relevance of the study and the important insights that might be obtained from the collected data. But differential privacy sends a different signal, basically telling the respondent: no matter which information you provide, we will make sure it will be irrelevant for the final findings. The results will be more or less the same even if you do not participate. This message might destroy one of the few remaining incentives for survey participation: the feeling that the time spent answering a long and boring questionnaire is well spent as the provided answers will be an important contribution to help researchers to better understand and potentially improve our society.

## 5.2 Valid inference

With survey data, the goal is often to make inferences regarding an underlying (finite) population, that is, the results based on the survey data are treated as estimates for the true values in the population. Furthermore, major interest often lies in identifying causal effects. To be able to achieve these goals, it is essential to quantify the uncertainty in the estimates obtained from the survey data and sophisticated methods have been developed in the survey statistics literature to quantify the uncertainty in the obtained estimates accounting for complex sampling designs and nonresponse adjustments (see for example Särndal *et al.* (2003)).

A key challenge, which needs to be addressed when dealing with differential privacy for survey data, is how to obtain statistically valid inferences from the protected data, that is, how to take the extra uncertainty from the protection mechanism into account. Research in this area is still limited. While it is straightforward to quantify the extra uncertainty for simple algorithms such as the Laplace mechanism which simply adds noise to the generated output, measuring the uncertainty for more complex algorithms is challenging. Furthermore, some algorithms introduce systematic bias, which is difficult to control for, especially for multivariate statistics. Finally, postprocessing steps, which are typically performed to ensure consistency and to avoid implausible values such as negative cell counts in a frequency table further affect the final estimates in complex ways invalidating potential adjustments even for simple algorithms.

### 5.3 No Secrecy of the Sample?

Common intuition suggests that sampling offers additional protection for certain types of risks, as a potential attacker does not know whether the target is part of the sample or not. This intuition has been formalized in the differential privacy literature for certain sampling types such as simple random sampling with or without replacement or Poisson sampling (Balle *et al.*, 2018), showing that these sampling designs can indeed lead to privacy amplification. For example, in a blog post, Adam Smith (2009) illustrated that given a mechanism offering a privacy guarantee of  $\varepsilon_1$  if run on the entire population and a sampling design based on simple random sampling with replacement with sampling rate  $r$ , the privacy

guarantee provided by running the mechanism on the sample instead of the population is  $\varepsilon_2 = r\varepsilon_1$ .

Quantifying the privacy amplification offered through sampling would allow agencies to either report more accurate statistics for a given privacy budget or offer a higher level of data protection to the respondents. However, the sampling designs employed by government agencies are typically far more complex than the ones studied so far. Stratified, cluster, and PPS sampling are commonly applied to improve the efficiency and/or reduce the costs of data collection. Furthermore, the final sample is typically drawn in multiple stages combining different sampling strategies at the various levels.

Initial research indicates that complex sampling designs can actually lead to privacy degradation instead of amplification. Bun *et al.* (2020) find that stratified sampling designs can negatively affect the privacy guarantees and amplification will generally be negligible for cluster sampling. However, they also find that privacy amplification can be retained for stratified sampling using proportional allocation, if random rounding is used instead of conventional rounding when determining the final sample sizes in each stratum. Better understanding the impacts of various (multistage) sampling designs on the final privacy guarantees that can be offered is an important area of future research.

## 5.4 Understanding the impacts of preprocessing

Statistical agencies aim at collecting samples that are representative of the target population and employ various methods to ensure that valid inferences can be obtained regarding the

population of interest. Before any analyses are run on the collected data, the raw data are also cleaned and updated in several pre-processing steps. Unit nonresponse is usually addressed by adjusting the initial survey weights, item nonresponse is often addressed by imputing missing values, and implausible values are corrected using various editing and imputation procedures. Currently there is limited understanding how the typical data production process affects the privacy guarantees or how the process could be adjusted to satisfy differential privacy.

Assuming the goal is to release differentially private microdata, one key question is whether these pre-processing steps should be carried out before or after generating the protected data. There are some arguments for both approaches. For example, data editing mostly deals with finding and adjusting implausible outliers by establishing so-called editing rules that raise a flag whenever a record violates a rule. Since editing can be seen as a form of clamping or truncation of the raw data, which typically helps to reduce the sensitivity of potential queries to be run on the data, it might actually improve the privacy guarantees provided by privacy-respecting analyses if editing rules would be enforced for the protected data.

Intuitively, having imputed values in the data could also increase the level of privacy, as these values are only estimates for the true values and thus already introduce an extra level of uncertainty. However, all imputation strategies at their core rely on models which are estimated using the observed data, and thus they can increase privacy leakage. The extent to which sensitive information is leaked depends on the imputation strategy used.

For example, early research (Clifton *et al.*, 2019) found that hot-deck imputation, popular amongst statistical agencies, is problematic since it imputes missing values by transferring the values of similar records that are fully observed (and is thus highly sensitive to changes in the data). This would be an argument to start by generating differentially private data first and then impute based on the generated data. As differential privacy is immune to post-processing, the imputation routines would not require any extra privacy budget as long as they only rely on the differentially private data. However, this rises the question how to deal with missing values in the data when enforcing differential privacy.

Similar questions arise for the weighting adjustments that are commonly applied to control for unit nonresponse and other deficiencies such as undercoverage in the sampling frame. If unit-nonresponse adjustments and calibration techniques, which ensure that the weighted survey estimates exactly match certain values that are known for the population, are used prior to protecting the data, it will be important to understand, how these procedures affect the sensitivity of the data. If the data are protected first, this would require additionally protecting all the information beyond the survey data which is used for the weighting adjustments. For nonresponse weighting this would include the design variables from the sampling frame, for calibration this would imply that only noisy benchmarks from the population could be used. Adopting this strategy raises challenging questions how to account for these extra measures when trying to make inferences regarding the population of interest.

## 6 Setting the value of $\varepsilon$

Deciding which value to use for the privacy parameter  $\varepsilon$  is arguably the most difficult decision in any deployment of differential privacy in practice. Partly because of the difficulties of understanding the implications of different values of  $\varepsilon$  for privacy and accuracy as discussed in Section 4, but also because the optimal value is inherently a social choice: how much privacy am I willing to give up in change for more accurate results? But there are several aspects, which make finding the optimal value for  $\varepsilon$  specifically challenging for government agencies.

### 6.1 Government agencies often act as intermediates

From an economic perspective, the industry setting involves two parties with relatively simple utility functions. The data provider will aim at minimizing the amount of privacy leakage while still getting the service offered by the company. The data recipient will try to maximize the privacy leakage to get more accurate data, while ensuring the data provider is still willing to participate. Leaving aside all the challenges in understanding  $\varepsilon$ , information asymmetries, and the privacy paradox, economic theory tells us that this should lead to a market equilibrium at the optimal value of  $\varepsilon$ .

The situation is more complex in the government context, since the data collecting agencies are often not the end-users of the collected data. For example, data collected by national statistical agencies such as the U.S. Census Bureau are used by various stakeholders, including politicians, journalists, and social scientists. The agencies not only face the dilemma that

they need to anticipate the utility functions of all these stakeholders, the utility functions will likely also vary considerably between the stakeholders. Furthermore, as discussed above, the agency cannot anticipate all the future analyses that users will be interested in, making it impossible to evaluate the impacts on accuracy. In a worst case scenario, the agency will settle on a value of  $\varepsilon$  that data providers consider insufficient to protect the data while the potential users of the data are no longer willing to work with the data as they feel that too much accuracy has been compromised.

## **6.2 Low values of $\varepsilon$ potentially harmful for the data provider**

Another aspect makes finding the optimal value for  $\varepsilon$  more challenging in the government context. Once the data have been shared, the only party benefitting from the data in the industry context, is the company. Further usage of the data will typically only have negative effects for the data provider, for example if the company sells the data to a data broker. Thus, the primary goal of the provider will always be to keep the privacy leakage as small as possible.

This could be different in the government context where data with low accuracy might be at least as harmful as data with low levels of privacy. To illustrate, we can look at the results of the test run of the TopDown algorithm of the U.S. Census Bureau on the Decennial Census 2010. As illustrated in a New York Times article by Wezerek and Van Ripper (2020), the version of the algorithm used in the test run introduced a systematic downward bias in the counts of Native Americans on reservations. Since according to the Census, the counts

from the next decennial Census will be used “to inform the allocation of hundreds of billions in federal funding” (U.S. Census Bureau, 2020), it is obvious that any underestimation in the counts will have direct negative impacts for the subjects involved. The problem has been fixed in the TopDown algorithm in the meantime, but this example illustrates that the utility function of the data provider is far more complex as low levels of accuracy will also affect their utility. Arguably, considering the scenario described above, there will likely be data providers that would prefer giving up most of their privacy, if they would otherwise risk receiving less government funding. Obviously, similar situations would arise in other scenarios, for example, if government money is spent on health care measures based on differentially private findings from a survey or poverty rates are computed using differentially private data.

### **6.3 The optimal value for the society**

Finding the optimal value of  $\epsilon$  is difficult for various reasons. Obviously, different respondents will have different views on privacy. For example, healthy respondents will typically be much less concerned regarding questions about their health status than respondents suffering from a rare disease. But how do we aggregate across these individual preferences to find the optimal value for the society? Would it be ethically defensible, if the released data would still be accurate enough to allow health care insurances to better target their premiums? There might be a benefit for the majority of respondents as their premiums would decrease, but some respondents might have to pay the price of losing their insurance. Thus, simply



maximizing the utility across respondents would not be an option. The problem is magnified by the fact that deciding not to participate in the survey will not save the respondent from these types of negative consequences.

## 7 The Road Forward

Despite its critical tone, this article should not be read as a general critique of differential privacy nor should it imply by any means that the concept is generally not suitable for the government context. Being able to offer formal privacy guarantees would be a major leap forward regarding how statistical agencies value the privacy of their respondents and the reconstruction attacks conducted by the U.S. Census Bureau on some of their own previously released data products (Abowd, 2019) clearly illustrated that the protection methods that are still popular among statistical agencies around the world are no longer sufficient to adequately protect the data.

The aim of the article is to raise the awareness that adopting the concept of differential privacy at government agencies raises many challenges and questions that did not have to be addressed in previous deployments. The hope is to stimulate interdisciplinary research to address these challenges. To be fair, some challenges cannot be solved easily. It will always be true, for example, that the size of the database will be relatively small in the survey context. However, the increasing popularity of administrative records and found data for statistical purposes might mitigate some of these problems. Besides, little attention has been given so far to the performance of algorithms if run on small samples. Substantial gains in

terms of accuracy might still be possible in practice if the focus will shift to developing optimal algorithms for fixed sample sizes.

Other challenges can be addressed by increased collaborations between disciplines. Up to date the vast majority of papers published on differential privacy have been authored by Computer Scientists. Since topics such as statistical inference, complex sampling designs, or nonresponse adjustment typically only play a minor role in this field, it is not surprising that research in this area has been limited so far. While some papers have appeared in recent years that also look at the problem from a statistical perspective (see, for example, Wasserman and Zhou (2010); Awan and Slavković (2020); Karwa *et al.* (2016); Bowen *et al.* (2020)), more collaborations between Computer Scientists and Statisticians—especially Survey Statisticians—would certainly help address many of the questions raised in Section 5 of the paper.

Finally, collaborations need to be expanded further when trying to identify the value of  $\varepsilon$  that is optimal for the society. Economists, Social Scientists, Psychologists, experts on Data Ethics, and Survey Methodologists could all make invaluable contributions to better understand individual feelings about privacy, identify models that accurately describe the social welfare function as discussed in Abowd (2019), understand social behaviour in response to threats to privacy, etc. The successful collaborations between Computer Scientists and Legal Experts to study the legal implications of differential privacy (Altman *et al.*, 2015; Nissim *et al.*, 2017; Nissim and Wood, 2018) already illustrate the knowledge gains possible from such collaborations.

Differential privacy is currently at a critical transition stage, moving from a very attractive but purely theoretical concept to becoming the quasi-standard that serves as a benchmark for any research on data privacy (for the good or the bad it is almost impossible these days to publish any research that does not meet the requirements of differential privacy without detailed explanations, why differentially private methods were not considered to address the problem). Whether the concept will also be adopted in practice for the data dissemination strategies of government agencies will depend on finding answers and solutions to the questions and challenges raised in this article. The theoretical properties of differential privacy have been studied extensively over the last decade. Now is the time to better understand the implications of adopting the idea in practice and this will require joint efforts from various disciplines as many of the problems arising in practice are clearly outside the realm of theoretical computer science.

## References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
- Abowd, J. M. (2018). The US Census Bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2867–2867.

- Abowd, J. M. (2019). Staring down the database reconstruction theorem.
- Alabi, D., McMillan, A., Sarathy, J., Smith, A., and Vadhan, S. (2020). Differentially private simple linear regression. *arXiv preprint arXiv:2007.05157* .
- Altman, M., Wood, A., O’Brien, D. R., Vadhan, S., and Gasser, U. (2015). Towards a modern approach to privacy-aware government data releases. *Berkeley Technology Law Journal* **30**, 3, 1967–2072.
- Apple’s Differential Privacy Team (2017). Learning with privacy at scale. *Apple Machine Learning Journal* **1**, 8.
- Awan, J. and Slavković, A. (2020). Structure and sensitivity in differential privacy: Comparing k-norm mechanisms. *Journal of the American Statistical Association* , (online first), 1–56.
- Bailie, J. and Chien, C.-H. (2019). Abs perturbation methodology through the lens of differential privacy.
- Balle, B., Barthe, G., and Gaboardi, M. (2018). Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, 6280–6290.
- Bambauer, J., Muralidhar, K., and Sarathy, R. (2013). Fool’s gold: an illustrated critique of differential privacy. *Vand. J. Ent. & Tech. L.* **16**, 701.

- Bowen, C. M., Liu, F., *et al.* (2020). Comparative study of differentially private data synthesis methods. *Statistical Science* **35**, 2, 280–307.
- Bun, M., Drechsler, J., Gaboardi, M., and McMillan, A. (2020). Controlling privacy loss in survey sampling. *arXiv preprint arXiv:2007.12674* .
- Clifton, C., Hanson, Eric, J., Merrill, K., and Merrill, S. (2019). Smooth sensitivity for k-nearest neighbor imputation.
- Coutts, E. and Jann, B. (2011). Sensitive questions in online surveys: Experimental results for the randomized response technique (rrt) and the unmatched count technique (uct). *Sociological Methods & Research* **40**, 1, 169–193.
- Ding, B., Kulkarni, J., and Yekhanin, S. (2017). Collecting telemetry data privately. In *Advances in Neural Information Processing Systems*, 3571–3580.
- Dwork, C. (2006). Differential privacy. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, 1–12. Springer, Berlin.
- Dwork, C., Mcsherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*, 265–284.
- Dwork, C., Roth, A., *et al.* (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* **9**, 3-4, 211–407.

- Edgell, S. E., Himmelfarb, S., and Duchan, K. L. (1982). Validity of forced responses in a randomized response model. *Sociological Methods & Research* **11**, 1, 89–100.
- Erlingsson, Ú., Pihur, V., and Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, 1054–1067.
- Fienberg, S. E., Rinaldo, A., and Yang, X. (2010). Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *International Conference on Privacy in Statistical Databases*, 187–199. Springer.
- Ghosh, A., Roughgarden, T., and Sundararajan, M. (2009). Universally utility-maximizing privacy mechanisms. In *Proceedings of the 41st annual ACM symposium on Theory of computing, STOC '09*, 351–360, New York, NY, USA. ACM.
- Hallowell, A. and Rector, A. (2020). Maine state economist letter to census on differential privacy.
- Karwa, V., Slavković, A., *et al.* (2016). Inference using noisy degrees: Differentially private  $\beta$ -model and synthetic graphs. *The Annals of Statistics* **44**, 1, 87–112.
- Kirchner, A. (2015). Validating sensitive questions: A comparison of survey and register data. *Journal of Official Statistics* **31**, 1, 31–59.
- Kohut, A., Keeter, S., Doherty, C., Dimock, M., and Christian, L. (2012). Assessing the representativeness of public opinion surveys. *Washington, DC: Pew Research Center* .

- Landsheer, J. A., Van Der Heijden, P., and Van Gils, G. (1999). Trust and understanding, two psychological aspects of randomized response. *Quality and Quantity* **33**, 1, 1–12.
- Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J., and Vilhuber, L. (2008). Privacy: Theory meets practice on the map. In *2008 IEEE 24th international conference on data engineering*, 277–286. IEEE.
- Messing, S., DeGregorio, C., Hillenbrand, B., King, G., Mahanti, S., Mukerjee, Z., Nayak, C., Persily, N., State, B., and Wilkins, A. (2020). Facebook Privacy-Protected Full URLs Data Set.
- Nissim, K., Bembenek, A., Wood, A., Bun, M., Gaboardi, M., Gasser, U., O’Brien, D. R., Steinke, T., and Vadhan, S. (2017). Bridging the gap between computer science and legal approaches to privacy. *Harv. JL & Tech.* **31**, 687.
- Nissim, K. and Wood, A. (2018). Is privacy privacy? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **376**, 2128, 20170358.
- Oberski, D. L. and Kreuter, F. (2020). Differential privacy and social science: An urgent puzzle. *Harvard Data Science Review* **2**, 1.
- Ruggles, S., Fitch, C., Magnuson, D., and Schroeder, J. (2019). Differential privacy and census data: Implications for social and economic research. In *AEA papers and proceedings*, vol. 109, 403–08.

- Särndal, C.-E., Swensson, B., and Wretman, J. (2003). *Model assisted survey sampling*. Springer Science & Business Media.
- Smith, Adam (2009). Differential privacy and the secrecy of the sample.
- Soria-Comas, J. and Drechsler, J. (2013). Evaluating the potential of differential privacy mechanisms for census data. In *UNECE work session on statistical data confidentiality*.
- Taddicken, M. (2014). The ‘privacy paradox’ in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure. *Journal of Computer-Mediated Communication* **19**, 2, 248–273.
- Tang, J., Korolova, A., Bai, X., Wang, X., and Wang, X. (2017). Privacy loss in apple’s implementation of differential privacy on macos 10.12. *arXiv preprint arXiv:1709.02753* .
- Uber Security (2017). Uber releases open source project for differential privacy.
- U.S. Census Bureau (2020). Why your answers matter.
- Vadhan, S. (2017). The complexity of differential privacy. In *Tutorials on the Foundations of Cryptography*, 347–450. Springer.
- Van Riper, D., Kugler, T., and Ruggles, S. (2020). Disclosure avoidance in the census bureau’s 2010 demonstration data product. In *International Conference on Privacy in Statistical Databases*, 353–368. Springer.
- Warner, S. L. (1965). Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association* **60**, 309, 63–69.



Wasserman, L. and Zhou, S. (2010). A statistical framework for differential privacy. *Journal of the American Statistical Association* **105**, 489, 375–389.

Wezerek, X. and Van Ripper, D. (2020). Changes to the census could make small towns disappear.