Legacy Statistical Disclosure Limitation Techniques for Protecting 2020 Decennial US Census: Still a Viable Option

Krishnamurty Muralidhar[1] and Josep Domingo-Ferrer[2]

[1] University of Oklahoma, Price College of Business, Dept. of Marketing and Supply Chain Management, 307 West Brooks, Adams Hall Room 10, Norman, OK 73019, U.S.A.
krishm@ou.edu

[2] Universitat Rovira i Virgili, Dept. of Computer Engineering and Mathematics, CYBERCAT-Center for Cybersecurity Research of Catalonia, Av. Països Catalans 26, 43007 Tarragona, Catalonia
josep.domingo@urv.cat

Abstract

In our article "Database Reconstruction Is Not So Easy and Is Different from Reidentification", we show that reconstruction can be averted by properly using traditional statistical disclosure control (SDC) techniques, also sometimes called legacy statistical disclosure limitation (SDL) techniques. Furthermore, we also point out that, even if reconstruction can be performed, it does not imply reidentification. Hence, the risk of reconstruction does not seem to warrant replacing traditional SDC techniques with differential privacy (DP) based protection. In "Legacy Statistical Disclosure Limitation Techniques Were Not an Option for the 2020 US Census of Population and Housing", by Simson Garfinkel, the author insists that the 2020 Census move to DP was justified. In our view, this latter paper contains some misconceptions that we identify and discuss in some detail below. Consequently, we stand by the arguments given in "Database Reconstruction Is Not So Easy…".

Introduction

In this article, we address some of the issues that Garfinkel (2023) (hereafter referred to as G23) raises regarding our earlier article Muralidhar and Domingo-Ferrer (2023) (hereafter referred to as MDF23). Before we do however, we feel compelled to make the following preliminary remarks:

(1) G23 makes the following statements:

> Even if "more than half of the matches are incorrect, and an external attacker has no means of confirming them" (Jarmin 2019), a group of hoodlums out to terrorize mixed race families can easily confirm the putative matches by visiting the addresses with baseball bats.
>
> More accuracy means that demographers will likely publish better papers with improved *p* values using the published data. It also means that the

> hoodlums' bats will be more likely to find mixed race households when they pound on the door.
>
> There is no evidence to support the above claims. This is baseless fearmongering intended to inflame rather than inform.

(2) The author makes another similar statement relating to trans children in Texas, extracted from a non-refereed Opinion in the *Scientific American* (Keyes and Flaxman 2022), where the authors claim that they can identify children who changed gender between 2010 and 2020. It is important to note that the results are based on simulated data and not on real data. Regardless of the merits of Keyes and Flaxman (2022), there is no evidence that Census data have been (or can be) used to identify and target trans children.[1] Again, baseless fearmongering intended to inflame rather than inform.

We now proceed to address the main content of G23, much of which is only tangentially relevant to MDF23. Sections 1, 2, and 3 in G23 are a primer on differential privacy; Section 6 offers a defense of differential privacy. We never suggested, as G23 claims, that the USCB "should have used a combination of data suppression and swapping"; G23 uses this to engage in an extensive discussion of suppression. We ignore this discussion as well. We focus our attention primarily on Sections 4 and 5 in G23 and, where appropriate, on specific comments in the other sections.

To facilitate the readability of this paper, we first summarize the specific issues raised by MDF23:

(1) The relevance of Dinur and Nissim (2003) in Census database reconstruction.
(2) The Census reconstruction procedure.
(3) The Census reidentification procedure.
(4) The impact of geography and policy choices on the performance of swapping and the DP-based method.
(5) A discussion of the privacy offered by the DP-based method; and
(6) A discussion of the accuracy of the data generated using the DP-based method.

In the following sections, we address each issue above in greater detail in view of G23.

The Dinur-Nissim Reconstruction

One of the primary motivations of our study was to highlight the fact that, while the Dinur and Nissim (2003) study is an important contribution, it has limited relevance in the context of the Census. Specifically, we wanted to bring attention to the "CD Model" described in Dinur and Nissim (2003), which states that a data curator may be able to modify the original data set to produce a "modified" data set that could possibly be securely shared. It is easy to see that this is

[1] Interested readers are invited to view the video of the presentation of this paper at the 2020 Census Data Products – Public Workshop on the Demographic and Housing Characteristics Files held June 21-22, 2022, in Washington DC. This paper is found in the video presentation under Use Cases – Part II. (https://www.nationalacademies.org/event/06-21-2022/2020-census-data-products-workshop-on-the-demographic-and-housing-characteristics-files)

exactly the procedure that the USCB employs. Statistical disclosure limitation is applied to the Census Edited File to create the Hundred-percent Detail File. All tabular data are then generated from the Hundred-percent Detail File (and no access is provided to the Census Edited File) (Zayatz et al 2009). Furthermore, unlike the Dinur and Nissim (2003) reconstruction procedure where the adversary can issue any query, in the Census context only the results of a pre-determined set of queries (the tabular data results) are allowed. The result of this is that the adversary can only reconstruct the Hundred-percent Detail File and never the Census Edited File. Another possible way to state this is as follows … *you can only reconstruct what the Census allows you to reconstruct and nothing else*.

G23 contends that there are two types of protection models, local and trusted curator. According to him, the local model involves adding noise to microdata while the trusted curator model involves adding noise to the output. But this is a misunderstanding of the methods. Indeed, the local and trusted curator models represent *where* the protection is applied, *not what protection is applied*. In a local model, the respondent modifies the response value prior to sharing with the curator. In a trusted curator model, the respondent shares the true value with the curator. In fact, it is possible to implement noise addition to microdata in the trusted curator model (which is the case with the American Community Survey (McKenna 2019)). Note that, by definition, a local model precludes adding noise to the output.

G23 also questions our statements regarding the level of noise required to prevent reconstruction. Our statements are based directly on the Dinur and Nissim (2003) conclusions (our emphasis): "a database algorithm that is within $O(\sqrt{n})$ perturbation and is private against polynomial adversaries *in the strongest possible sense*." Thus, our original statement is accurate. That there are other bounds is irrelevant.

G23 also makes another dubious comment that adding noise at the microdata level is less efficient than adding noise at the query level: "With counting queries, the amount of noise added to each record in local model is the same as the total amount of noise added to the final result in the trusted curator model." This claim is made without any substantiation or support. We already addressed *this exact question* in our paper MDF23 (Section 3, paragraph 5) and showed that when many queries are issued, the noise added by a DP-based method will be higher than with an alternate procedure such as randomized response. Bach (2022)'s statement is consistent with our analysis.

Finally, the 2010 Decennial Census tabular data release *did not involve any noise addition*. Instead, data swapping was used as the disclosure limitation tool (Zayatz et al 2009). One of the biggest benefits of data swapping is that it *does not modify the original values of the attributes*. In addition, data swapping was implemented with the explicit requirement that voting age population must be preserved exactly at the block level. As a result, the question "how many children are in this block" *will be answered exactly, without any error*. The same cannot be said for the 2020 Decennial Census with the DP-based method as many demographers have attested to.

The Census Reconstruction Procedure

The entire justification for adopting differential privacy for the 2020 Census was based on the results of the reconstruction and reidentification experiments of the 2010 tabular data release. The USCB has also made the argument explicit in Abowd (2021):

> …. only over the last few years have computing power and the sophisticated numerical optimization software necessary to perform these types of reconstructions advanced enough to permit reconstruction attacks at any significant scale.

From the very beginning, Professor Ruggles at the University of Minnesota Population Center has argued that the reconstruction amounts to little more than creating a list of individuals from the tabular data: "any tabular data can be expressed as microdata." (Ruggles 2018) This was later proved by Muralidhar (2022), who was able to perform the reconstruction using only the data released to the public and an Excel spreadsheet on a laptop, without any numerical optimization software. G23 makes the following statement:

> When the 2010 data tables were published, Census Bureau leadership did not think reconstruction possible.

It speaks very poorly of the leadership at the USCB both in 2010 *and later* that they did not realize that tabular data releases enable reconstruction. It speaks even more poorly of the post-2010 Census Bureau leadership that they chose to attribute their failure to advances in "computing power and the sophisticated numerical software." Note that G23 does not acknowledge Muralidhar (2022) and the detailed explanation the latter provides as to how the reconstruction is simple and straightforward and that the technology needed to perform reconstruction has been available for 50+ years.

An even more important contribution of Muralidhar (2022) is that he showed that, even for small geography (at the Census tract level), there are billions of feasible reconstructions that are equally likely. This has significant implications for the claims of reidentification as we now discuss.

The Census Reidentification Procedure

Of all the claims made by the Census Bureau, the most dramatic is the following (Abowd 2021, Appendix B):

> If the external data on name, address, sex, and age are comparable to the 2010 Census, then the attacker will putatively re-identify 238 million persons (77% of the 2010 Census resident U.S. population). Confirmed re-identifications will be 179 million (58% of the same population). This means that with the best quality external data, relative to the 2010 Census, as many as 179 million persons could be correctly re-identified using the attack strategy outlined here.

Abowd (2021) leaves no doubt as to what he means by reidentification:

> … link the reconstructed microdata to a real name and address associated with the block, sex, age, race, and ethnicity on the reconstructed microdata.

These are dramatic statements because they directly contradict the legal requirements in Title 13 of the United States Code, described in Abowd et al (2020) as follows:

> Additionally, the Department of Commerce (2017), in which the Census Bureau is housed, has issued directives regarding the protection of personally identifiable information (PII) and business identifiable information (BII). These directives largely mirror those issued by other government agencies and prohibit release of information that can be used "to distinguish or trace an individual's identity, such as their name, social security number, biometric records, etc., alone or when combined with other personal or identifying information which is linked or linkable to a specific individual, such as date and place of birth, mother's maiden name, etc."

Abowd (2021) states explicitly that this was the primary reason for adopting differential privacy for the 2020 Decennial Census was due to fact that the tabular data release from the 2010 Decennial Census violated Title 13 requirements:

> To defend against this known vulnerability, the Census Bureau explored different confidentiality methods that explicitly defend against database reconstruction attacks and concluded that the best tool to protect against this modern attack while also preserving the accuracy and usability of data products comes from the body of scientific work called "differential privacy."

MDF23 pointed out that the claims of reidentification were vastly overstated. These results are supported by the empirical analysis of 2010 Census tabular data by Muralidhar (2022). The procedure described in Abowd (2021) was shown to directly contradict the reidentification procedure suggested by the Census Bureau's own Research and Methodology Directorate (McKenna 2019). The above statements are also in direct contradiction to the statement by Acting Director of the Census Bureau Dr. Ron Jarmin, who said in 2019: "more than half the matches are incorrect." Based on all this information, we were pointing out the fact that what Abowd (2021) claims as reidentification was *in fact the accuracy of reconstruction*, as illustrated in Section 5.2 of MDF23. What is lacking is the ability to link the respondents in the reconstructed data to their identity.

G23 seems to acknowledge that our conclusion is correct when he says, "there are many ways that even unidentified census microdata with block-level geography can cause harm." This may be true, but this is a far cry from claiming that the reconstructed data can be linked to individual names and addresses. G23 now claims that even if the reidentification was incorrect, it can be verified by the adversary by visiting the address. If we take this to its logical conclusion, for the adversary to make the same claim as Abowd (2021) regarding reidentification, the adversary *will have to conduct a new Census*.

In summary, G23 offers nothing to contradict our conclusion that the reidentification claims regarding the 2010 Census tabular data release have been vastly overstated by the USCB reconstruction attacks.

The Impact of Policy and Geography Choices

MDF23 pointed out that the performance of data swapping and DP-based method were evaluated using different benchmarks. Perhaps the most important of this is reflected in the *invariants*, that is, counts that are preserved and the geographic level at which these counts are preserved. The 2010 Decennial Census tabular data release preserved *exact counts at the block level for the total population, population age 18 and above, and population below age 18*. Regarding the 2020 Decennial Census, a 2017 Census publication (Dajani et al 2017) states that:

> By agreement with the Department of Justice (2000), the Census Bureau will provide exact counts at the Census block level for the following variables:
> - Number of people: total, age 18+ (voting age), and less than age 18,

The updated version of the same document in 2020 (Abowd et al 2020) states:

> The Census Bureau currently plans the following invariants for the 2020 Census data publications:
> - Total number of people by state;

This is a remarkable change in policy (and one that has had a significant detrimental impact on data accuracy as we will discuss later). While the 2017 invariants would have been consistent with the 2010 invariants, the 2020 invariants have remarkably relaxed the requirement, from the block level to the state level. From the Census demonstration data products released to the public, we know that the DP-based method for the demonstration products only preserved totals at the state level, while the original 2010 tabular data release using swapping preserved totals at the block level. Considering that the number of blocks varies from 6012 (Washington DC) to 668,757 (Texas), this represents a massive change in policy. MDF23 point out that such changes in policy have a direct impact on the protection afforded by the statistical disclosure limitation procedure.

It is true that the USCB has recently released comparisons of applying data swapping with different parameters (the proportion of swapped records) (Hawes 2022). However, it is still not clear whether these new parameters were applied with the old policy (invariants at the block level) or the new policy (invariants at the state level). Just changing the proportion of records swapped without changing the invariant requirement is unlikely to change the performance of the swapping algorithm. It would have been relatively easy for the Census Bureau to release data where both data swapping and the DP-based method were held to the same standard. By not doing so, the Census Bureau has made any meaningful comparison impossible. Unfortunately, G23 does not address this issue at all.

Instead, G23 chooses to make new policy by stating “the primary risks that DP was designed to protect: the release of confidential information about race.” This is quite a remarkable statement

and, in some ways, a curious one. G23 claims that information regarding race and ethnicity data are "in general, not for sale." While this claim is arguable, let us assume it is true. If this is the case, then race and ethnicity are the confidential variables that the adversary *does not know*. Remember that reconstruction was performed by matching age and sex from the reconstructed data to the externally obtained data. This is a common phenomenon encountered in statistical disclosure limitation where the attributes used to perform the matching are referred to as quasi-identifiers. To prevent the adversary from using the quasi-identifiers (age and sex) to learn the confidential attributes (race and ethnicity), privacy models such as k-anonymity (Samarati 2001) would modify the quasi-identifiers and not the confidential attributes. Thus, if the objective is to prevent reconstruction, it would make more sense to protect the age and sex attributes rather than the race and ethnicity attributes. The rationale is simple – if you cannot reconstruct the data accurately, you cannot infer the confidential attributes.

The Privacy Provided by DP

We must admit that we were astounded by the statement in G23 that the privacy loss parameter ϵ and the corresponding randomized response probability is "not relevant to the use of differential privacy for the Census." G23 contends that "there is no single query over the 2020 Census confidential data for which this privacy loss value is used. Instead, the privacy loss is spread over tens of millions of individual queries to produce the so-called noisy measurements." This represents a profound misinterpretation of the global privacy parameter.

According to Dwork et al (2019)

> Differential privacy allows us to quantify cumulative privacy loss as data are analyzed and re-analyzed, shared, and linked. These differentially private systems will, for the first time, allow us to measure and compare the total privacy loss due to these personal data-intensive activities.

The privacy parameter ϵ represents this total (cumulative) privacy loss. It would be a serious error to pretend that, since the privacy loss is distributed over tens of millions of individual queries, the cumulative privacy loss does not matter.

MDF23 does not say that "there are objectively correct values for data utility and privacy loss," as GF23 asserts. In our opinion, the statements "Taking a very small ϵ entails unaffordable utility loss, but taking very large ϵ entails very little noise addition and offers little to no protection against reconstruction" (MDF23) and "choosing a small ϵ means that the data owner has chosen to highly value privacy at a cost of data utility, while large ϵ values mean the data owner has chosen to highly value data utility at the cost of privacy" (G23), say the same thing albeit in different words.

Data Accuracy of the DP-based Method

The lack of accuracy of the DP-based method is well documented, as shown in MDF23. Unfortunately, G23 does not address any of the multiple *research* studies cited in MDF23. He only affirms that "the Census Bureau was largely successful, as evidenced by numerous data

analyses published since that make use of the data." According to Menger (2021), as many as 11% of the blocks have "impossibilities" (more households than people, zero housing units with non-zero population, zero population but non-zero occupied housing units, children but no adults, and very large households). Whether this constitutes a large successful implementation is up to question. As in many other instances, G23 makes the statement of numerous data analyses, but provides no citation.

G23 has chosen to address one humorous example highlighted in the NY Times of people living under water. G23 categorically states that this error was not due to DP. But if there were the case, one would have expected the USCB to issue some type of formal explanation for this. G23 further contends that this error could not have happened since, according to Abowd et al (2022), there is a constraint that "the number of householders (person one on the questionnaire) cannot be greater than the number of housing units". But there have been documented instances where "4.83% [of blocks] have zero occupied housing units which have people living in them (that is, there are no occupied dwellings, but there are people living there)" (Menger 2021). Clearly, the constraint described by Abowd et al (2022) is being violated.

Finally, we address the rather serious claim by G23: "MDF fails to address a very real harm to the statistics of vulnerable populations done by suppression and swapping. That harm is *erasure*. By design, these techniques hide the statistical contribution of minorities." As we have emphasized before, we absolutely deny that we recommend suppression. Removing suppression from the discussion and focusing only on swapping, *we can categorically state that data swapping does not result in erasure*. Data swapping simply exchanges the attribute values between records. As such, there is no addition or subtraction to the counts. Hence, we reiterate that *data swapping does not result in erasure*.

The real irony is that the DP-based method *does result* in erasure. With the DP-based method, the frequency counts are modified by adding or subtracting a random noise term. As a result, *every time negative noise is applied to a count of minorities, there is erasure*. And there is no way to prevent this from happening. Thus, contrary to the claim made by G23, it is the DP-based method that is likely to erase the contribution of the minorities, not data swapping.

Conclusions

We have carefully gone through the G23 paper. The paper offers no direct substantive, justifiable criticism of MDF23. We have gone to great lengths to address every issue that we found relevant.

G23 calls for finding a constructive way forward. The first step in moving forward is to perform a comprehensive comparison of data swapping and the DP-based method using the same criteria as in the 2010 tabular data release. The second step is to re-evaluate the reidentification claims. Moving forward also requires acknowledging the work of researchers on the other side and not ignoring all their arguments.

References

Abowd, J. M. 2021. Declaration of John M. Abowd. Case no. 3:21-CV-211-RAH-ECM-KCN, U. S. District Court for the Middle District of Alabama. Apr. 13.

Abowd, J., Ashmead, R., Cumings-Menon, R., Garfinkel, S., Heineck, M., Heiss, C., Johns, R. 2022. "The 2020 Census Disclosure Avoidance System TopDown Algorithm." *Harvard Data Science Review*, no. Special Issue 2 (June). https://hdsr.mitpress.mit.edu/pub/7evz361i (accessed May 14, 2023).

Abowd, J. M., Benedetto, G. L., Garfinkel, S. L., Dahl, S. A., Dajani, A. N., Graham, M., Hawes, M. B. 2020. "The Modernization of Statistical Disclosure Limitation at the U.S. Census Bureau," August. https://www.census.gov/library/working-papers/2020/adrm/CED-WP-2020-009.html (accessed May 14, 2023).

Bach, F. 2022. "Differential Privacy and Noisy Confidentiality Concepts for European Population Statistics." *Journal of Survey Statistics and Methodology*, 10:642-687. DOI: 10.1093/jssam/smab044

Dajani, A. N., Lauger, A.D., Singer, P. E., Kifer, D., Reiter, J. P., Machanavajjhala, A., Garfinkel, S. L., Dahl, S. A., Graham, M., Karwa, V., Kim, H., Leclerc, P., Schmutte, I. M., Sexton, W. N., Vilhuber, L., Abowd, J. M. 2017. "The modernization of statistical disclosure limitation at the U. S. Census Bureau." In: *Census Scientific Advisory Committee Meeting*, Sep. 14-15, Suitland MD, USA. https: //www.census.gov/library/video/2017/2017-09-sac.html (accessed May 14, 2023).

Dinur, I., Nissim, J. 2003. "Revealing Information While Preserving Privacy." In: Proc. of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, pp. 202-210, Jun. 9-12, 2003, San Diego CA, USA. DOI: 10.1145/773153.773173 https://dl.acm.org/doi/10.1145/773153.773173

Dwork, C., Kohli, N., Mulligan, D. 2019. "Differential Privacy in Practice: Expose Your Epsilons!" *Journal of Privacy and Confidentiality*, 9(2). DOI: 10.29012/jpc.689

Garfinkel, S. 2023. "Legacy Statistical Disclosure Limitation Techniques Were Not An Option for the 2020 US Census of Population and Housing". *Journal of Official Statistics*, this issue.

Hawes, M. 2022. "Reconstruction and Reidentification of the Demographic and Housing Characteristics File (DHC)", Presentation to the Census Scientific Advisory Committee, September 29-30, 2022, Washington DC. https://www2.census.gov/about/partners/cac/sac/meetings/2022-09/presentation-reconstruction-and-re-dentification-of-dhc-file.pdf (accessed March 14, 2023)

Jarmin, R. 2019. "Census Bureau Adopts Cutting Edge Privacy Protection for 2020 Census," Feb. 15, 2019. https://www.census.gov/newsroom/blogs/random-samplings/2019/02/census_bureau_adopts.html (accessed May 14, 2023).

Keyes, O., Flaxman, A. D. 2022. “How Census Data Put Trans Children at Risk.” *Scientific American,* September 21. https://www.scientificamerican.com/article/how-census-data-put-trans-children-at-risk/ (accessed May 14, 2023).

McKenna, L. 2019. “Disclosure Avoidance Techniques Used for the 1960 Through 2010 Decennial Censuses of Population and Housing Public Use Microdata Samples,” Research and Methodology Directorate, U.S. Census Bureau. https://www.census.gov/library/working-papers/2019/adrm/CED-WP-2019-007.html (accessed May 14, 2023)

Menger, G. 2021. “Using 2020 Census Data.”, Applied Geographic Solutions, Sep. 16. https://appliedgeographic.com/2021/09/using-2020-census-data/ (accessed May 14, 2023).

Muralidhar, K., Domingo-Ferrer, J. 2023. “Database Reconstruction Is Not So Easy and Is Different from Reidentification.” *Journal of Official Statistics,* this issue.

McKenna, L. 2019. “U.S. Census Bureau Reidentification Techniques.” Research and Methodology Directorate, U.S. Census Bureau. April. https://www2.census.gov/adrm/CED/Papers/CY19/2019-04-Reidentification%20studies-20210331FinRed.pdf (accessed May 14, 2023).

Muralidhar, K. 2022. “A Re-Examination of the Census Bureau Reconstruction and Reidentification Attack.” In: *Privacy in Statistical Databases – PSD 2022*, Sep. 21-23, Paris, France. Springer, pp. 312-323. DOI: 10.1007/978-3-031-13945-1 https://link.springer.com/chapter/10.1007/978-3-031-13945-1_22

Ruggles, S. 2018. “Implications of Differential Privacy for Census Bureau Data Dissemination,” Institute for Social Research and Data Innovation, University of Minnesota, December 2018. https://apps.bea.gov/fesac/meetings/Ruggles%20Presentation%20Revised.pdf (accessed May 8, 2023).

Samarati, P. 2001. “Protecting Respondents Identities in Microdata Release.” *IEEE Transactions on Knowledge and Data Engineering* 13(6):1010-1027.

Zayatz, L., J. Lucero, P. Massell, A. Ramanayake. 2009. “Disclosure Avoidance for Census 2010 and American Community Survey Five-year Tabular Data Products,” Census Bureau Research Report RRS2009-10. https://www.census.gov/content/dam/Census/library/working-papers/2009/adrm/rrs2009-10.pdf (accessed May 14, 2023).