The Scalar Decomposition of Segregation Measures Under Differential Privacy

Lee Fiorio, *University of Washington*, fiorio@uw.edu
Neal Marquez, *University of Washington*, nmarquez@uw.edu

INTRODUCTION

Though the United States has grown steadily more diverse in terms of race over the last forty years, the population trends driving this shift have not occurred evenly across metropolitan regions or the municipalities and neighborhoods within them (Wright et al, 2014). As such, the geography of US racial residential segregation has become increasingly complex. Once primarily defined by stark neighborhood-level separation of Whites and Blacks in the inner city, the current residential geography of race reflects the interaction of multiple groups across a range of contexts – from urban to suburban and even rural, between and within municipal boundaries.

To understand the multifaceted and evolving trajectory of racial segregation in the US, social scientists have developed new methods for measuring segregation and neighborhood change. Much of this literature has emphasized the importance of geographic scale in explicating the causes and consequences of emergent residential patterns observed by race (Fowler, 2016; Farrell & Lee, 2011; Lee et al., 2008; Reardon et al. 2008; Fischer et al., 2004). The empirical findings in this body of research have prompted important new lines of scientific inquiry across the social and health sciences (Ard, 2016). Crucially, they have also generated new insights for public agencies and lawmakers, as well as community-based organizations and advocates seeking to understand and mitigate structural inequality, as it is instantiated through residential segregation.

The Census Bureau's new Disclosure Avoidance System (DAS) for the release of the 2020 decennial Census runs the risk of undercutting the growing attention to geographic scale in segregation research. For one, multiscalar measurements of segregation are often reliant on the accuracy and precision of data published at finer geographic scales, like the Census Block or Block Group -- the very scales which are under greatest danger of being compromised by noise injection under the new DAS. Moreover, though the Census Bureau has given careful consideration on how best to ensure the consistency of data published across the Census geographic hierarchy (from states and counties down to tracts and blocks), little if any consideration has been given to how best to maintain the accuracy of composite measures, like measures of segregation, and their interrelation across geographic levels. To put simply, although the Census Bureau has made it a priority to guarantee that a count published for an area is consistent with the counts published for the subdivisions that comprise it, we do not know how the distribution of counts across those subdivisions or in relation to other counts will be affected. The accuracy of distributions is important since the extent to which an area is racially integrated depends on the degree to which different race groups are spread evenly within it.

In this paper, we evaluate the potential impact of the 2020 DAS on the precision of multiscalar segregation analyses. We begin by using simulation to produce noise-injected counts of race, ethnicity and voting age at a range of geographic scales, changing the level of noise. This simulated data allows us to estimate the impact of the DAS on standard segregation measures and isolate the effect of noise injection from the effect of post-processing adjustments that are part of the DAS. Next we use the most recent release of the 2010 demonstration data to compare results from a scalar decomposition of Theil's H -- a multigroup entropy index -- with those observed in the 2010 summary file. Finally, we assess the impact of the DAS on our ability to measure the intensity of segregation over time. Results from these analyses provide an account of the potential impact of the 2020 DAS on multiscalar analyses of segregation and suggest that in addition to reducing precision, the DAS will potentially compromise the validity census geography at finer scales.

BACKGROUND

Under Title 13 of the US Code, the Census Bureau is mandated to ensure the privacy and confidentiality of census respondents. Throughout its history, the Census Bureau has utilized a variety of different methods for protecting privacy and confidentiality of published data, including top-coding, data swapping and data suppression (Boyd, 2019). Although each of these privacy methods comes with a cost to data accuracy, the Census Bureau has historically done its best to maximize the accuracy of its data products while also fulfilling its privacy mandate.

Differential Privacy

Since the beginning of the last decade, there has been growing concern that increasing computational power and advancements in database reconstruction techniques have rendered previous disclosure avoidance procedures insufficient, leaving the census vulnerable to a database reconstruction attack. The fear is that with state-of-the-art computational methods a potential attacker could reconstruct individual-level (i.e. confidential) response data by comparing and cross-referencing the many thousands of tables that the Census publishes with each decennial census. Linking reconstructed microdata to an outside data source, an attacker could then identify individuals in the data.

To thwart such an attack on the 2020 decennial census, the Census Bureau has adopted a new Disclosure Avoidance System (or DAS). This system rests on two major changes: (1) reducing the number of tables or "queries" published from decennial census microdata and (2) injecting noise into the cells of the queries that are published (Abowd et al., 2018; Garfinkel et al., 2018). Though both are important and worthy of debate, it is the second of these, the noise injection, that has garnered the most attention among the academic and policy communities.

The Census Bureau's approach toward noise injection is based on a framework from computer science literature known as differential privacy. Under the differential privacy framework, privacy can be defined mathematically as the condition in which a query performed on a database is not conditional on an individual person or entity's inclusion in that database. The logic is that if the values of a query are dependent on whether a specific individual appears in the data, then the difference between the query on the data containing the individual and the query on the data not containing the individual will reveal potentially private information about the individual.

The concept of epsilon-differential privacy elaborated by Dwork (2008) formalizes the mathematical definition of differential privacy and provides a method for guaranteeing it with the addition of Laplace noise. As part of the new 2020 DAS, the Census Bureau plans to inject Laplace noise into many of the published tables. The amount of noise injected will be based on a predetermined global privacy budget, epsilon, which is then divvied up and allocated to across queries and geographic levels.

Importantly, in relevant tables, noise will be added at the cell level rather than the table level. For example, imagine a table that gives two counts for a specific county: (1) the number of individuals of voting age and (2) the number of individuals not of voting age. In this table, noise will be injected into each count independently. Because it's very likely that the counts will receive non-reciprocal noise, this means that the Census Bureau will need to apply post-hoc adjustments to noise-injected data so that they are consistent across query (e.g. all tables for a county have the same total population) and across geographic level (e.g. all tract-level tables for a county add up to the total count for the county). Additional adjustments will be made to ensure that the published data are legible to the public, most importantly that counts are non-negative integers.

How best to implement these post-hoc adjustments has been the primary focus of the Census Bureau's last minute push to finalize the DAS before the release of the first noise infused tables in 2021. This has meant that most discussions over how to improve the DAS over the last year have centered around how to improve data accuracy though improvements to the adjustment algorithm. But as we discuss in the next section, this focus has resulted in a lack of clarity over what extent the inaccuracy introduced by the DAS is the result of the noise-injection step and what extent the inaccuracy is the result of the post-hoc adjustment.

Measuring Segregation in Noise-Injected Data

Accurately measuring racial residential segregation in noise-injected data is complicated by the fact that segregation measures are *composite* measures: combining two or more measures into a single index. On average, the count of a specific race group will be accurate (setting aside the potential effects of post-processing); however, since the noise added to one race group is independent from the noise added to other race groups, the impact of noise injection on measures of segregation is less straightforward.

At first blush, it might make sense that measures of segregation will be biased downwards in noise-injected data under the new DAS. Racial residential segregation is commonly theorized to be the outcome of a non-random individual-level (or family-level) residential mobility process. When faced with the decision of where to live, individuals and families are differentially (and relationally) constrained along the lines of race. Systematic differences in wealth, income or access to credit along racialized lines might make a desirable area a safe investment for one group and unattainable for another group. Members of one racialized group may feel impelled to move when a family of a different group moves in nearby. Individuals or families of another group might be steered away from certain areas or prefer to stay near the places with which they are most familiar.

From the logic of individual- or family-level residential mobility, it follows that racial residential *integration* will occur in places where these racial sorting mechanisms do not exist or where they have a diminished effect. This same logic suggests that noise-injection could reduce the levels of segregation observed. As the distribution of different race groups appears more random – either through noise injection or through the breakdown of social processes that differentiate groups on where they live – levels of segregation observed should be lower.

Scale and Segregation

But the degree to which noise-injected data will resemble a distribution of population expected under a random residential sorting process depends on the interaction between the noise-injection procedure and the geographic scale of the data. In a set of tables with large aggregate counts (e.g. counts of race at the county scale), adding Laplace noise might result in small net changes that downwardly bias segregation measures -- particularly if race groups are segregated by county in the true data. However, in a set of tables with small aggregate counts (e.g. counts of race at the block or block group scale), adding Laplace noise could result in very large net changes that resemble a redistribution of the population into segregated "lumps". Adding 100 Hispanics to an all White county of 15,000 will make that county appear slightly less segregated. Adding 100 Hispanics to a block of 20 Whites and 10 Hispanics will make that block look much more segregated. Because noise-injection under the new DAS will occur to each table cell independently, tables published at scales where the added noise can be expected to regularly approach or exceed the size of the cell may see an upward bias in measures of segregation.

To further illustrate these dynamics, it is helpful to consider the differences between Laplace noise-injection under the 2020 DAS and data swapping, a procedure that was a central part of the disclosure avoidance strategy in previous censuses. Under data swapping, responses are switched among a subset of individuals residing in nearby areas. The total population in each table and at each level is held fixed. As the amount of swapping increases, the data will become less accurate, but the effect on segregation will always be downward since swapping more or less approximates a random residential sorting mechanism. This is true whether swapping is done between coarse geographic units like counties or fine geographic units like blocks. In comparison, as the amount of noise added under the 2020 DAS increases, the data will become less accurate, but the effect on segregation is more ambiguous and will depend on the scale under consideration and the initial distribution. Unlike swapping, noise injection under the DAS does not approximate individual-level residential mobility. Instead, noise added to counts will resemble net population change caused by some unobserved population process.

RESEARCH QUESTIONS

Given the complexity of the processes that shape residential patterns by race and given the difficulty of anticipating the effect of the 2020 DAS on composite measures, like measures of segregation, the goals of this paper are descriptive. Fundamentally, we want to try to understand how the 2020 DAS will impact multi-scalar analyses of segregation and to what extent its impact varies by context (i.e. metropolitan areas versus micropolitan areas; segregated places versus integrated places). We also would like to parse the changes to counts brought about by noise injection and changes to counts brought about by post-processing, if possible. All together, this gives us the following three research questions:

1. Using simulated data, what is the effect of noise injection on measures of segregation between tracts and within tracts as epsilon varies?

2. Comparing the 2010 Summary File (SF) to the noise-injected 2010 demonstration data, what is the effect of the DAS (noise injection and post-processing) on the scalar decomposition of Theil's H?
3. Incorporating data from the 2000 SF, what is the effect of the DAS (noise injection and post-processing) on the ability to accurately measure change to levels of segregation over time and across scale?

DATA AND METHODS

Data

We use race and ethnicity counts from the 2000 and 2010 Census data for block groups. We use data prepared by IPUMS which harmonizes 2000 census block group data to the geography definitions used in the 2010 census so that geographies are comparable over time and segregation metrics are not influenced by different geographic boundaries. In addition, we also use the 2010 Census end to end test of the Differential Privacy DAS as our comparison data set. For each block group and year we create mutually exclusive counts of the number of individuals who are non-Hispanic White, non-Hispanic Black, non-Hispanic Asian, and Hispanic. We remove any block group which has a zero total population count in both the 2010 true census release data and the DAS.

Measuring Segregation

The literature on racial residential segregation has produced a variety of approaches to quantifying segregation. Different segregation indices measure distinct dimensions segregation including: evenness, or the degree to which two groups are similarly distributed across a set of areas; isolation, or the degree to which a group only shares space with members of its own group; and clustering, or the degree to which a group is clustered in space (Massey & Denton, 1988). Differences between measures of segregation are also a reflection of changing data availability particularly as access to finer grained data, like block data, and auxiliary spatial data, like data on geographic adjacency, have become the norm. In the mid-20th century, when the most detailed information available were tables of race groups by tract -- with information about which tracts were adjacent to one another being difficult to obtain-- social scientists made use of simple but intuitive measures like the Index of Dissimilarity and the Isolation Index. Though these measures are still widely used in segregation research, research in the intervening decades have identified and attempted to rectify their limitations (Manley et al., 2019; Napierala & Denton, 2017; Wong, 1993 ; Morrill, 1991).

For the purposes of our analysis, we use Theil's H index, a segregation measure which increased in use in recent decades (Fischer et al., 2004). Like the Index of Dissimilarity, Theil's H measures evenness, or the degree to which race groups exhibit the same distribution across a set of areas. But Theil's H has two distinct advantages. First, unlike the Index of Dissimilarity which is limited to comparing the distribution of two groups (e.g. Whites and Hispanics) or one group and its residual (e.g. White and non-White), Theil's H allows the comparison of multiple groups. Second, Theil's H has mathematical properties that allow it to be decomposed into its component parts.

It is this second property that has made Theil's H a useful tool in multiscalar segregation analyses as Theil's H can be decomposed along a nested geographic hierarchy. The concept is simple and the results are generally informative. First, using counts of race groups for a particular metropolitan area, their distributions are assessed across the smallest geographic sub-unit in the data, like census blocks. This gives an overall Theil's H index for the metropolitan region. Then, because blocks are nested within tracts and tracts are nested within counties, it is then possible to determine how much each intermediate scale contributes to overall segregation.

To illustrate this concept and why it is informative, we provide two figures to demonstrate how groups may be distributed differently across different scales. Figure 1 shows two hypothetical metropolitan areas subdivided into 100 blocks each. Spread across these blocks are two groups, group A and group B, our populations of interest. In both metropolitan areas, both groups are completely segregated at the block-level. No member of group A or B shares a small square with a member of the other group. Figure 1 also shows that each metropolitan area is also made up of four intermediate subdivisions seen in bold. Let's say these are tracts or neighborhoods. In Metro Area 1 (left), groups A and B are completely segregated at the tract level. In Metro Area 2 (right), groups A and B are completely integrated at the tract level but remain segregated by block. Analyzing these hypothetical data using Theil's H index would allow us to estimate the portion of segregation occurring between tracts and the portion that is occurring within tracts (i.e. between blocks). In Metro Area 1, the uneven distribution of groups A and B are 100% explained by their segregation into different tracts. In Metro Area 2, tracts explain 0% of the uneven distribution of groups A and B, meaning that the overall level of segregation observed is 100% attributable to the distribution of groups A and B across blocks.

Figure 1: A Stylized Example of Tract- and Block-level Segregation

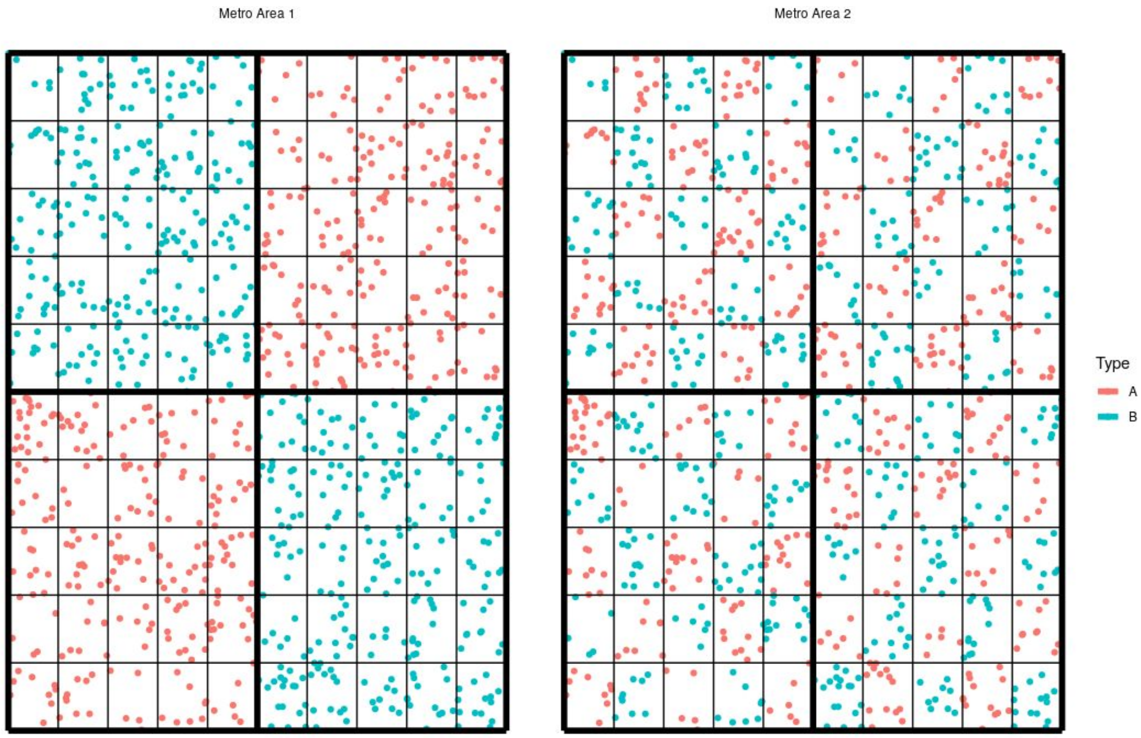

Figure 2 illustrates how Theil's H can be decomposed, not just by scale, but also by group. In addition to asking how much does each level of geography contribute to the overall segregation observed, we can ask how much does the distribution of each group contribute to the overall segregation observed. In Figure 2, we again show two hypothetical metropolitan areas, but this time with three different subpopulations, A, B and C. Again, the overall measure of Theil's H index in each metropolitan area is the same, despite the obvious differences. In Metro Area 1 (left), all three groups are sometimes observed sharing a block with another group. In Metro Area 2, group A is completely segregated from B and C. Decomposing Theil's H by group, we can quantify the different relative distributions of A, B and C in Metro Area 1 and Metro Area 2 by estimating what portion of the overall measure is the result of segregation between A versus B and C and and what portion is the result of segregation between the remaining groups, B versus C. Using Theil's H, we would observe that 66% of the overall segregation of Metro Area 1 is attributed to the relative isolation of group A, while in Metro Area 2, the distribution group A explains 100% of the segregation observed.

Figure 2: A Stylized Example for Variations of Multi-group Segregation

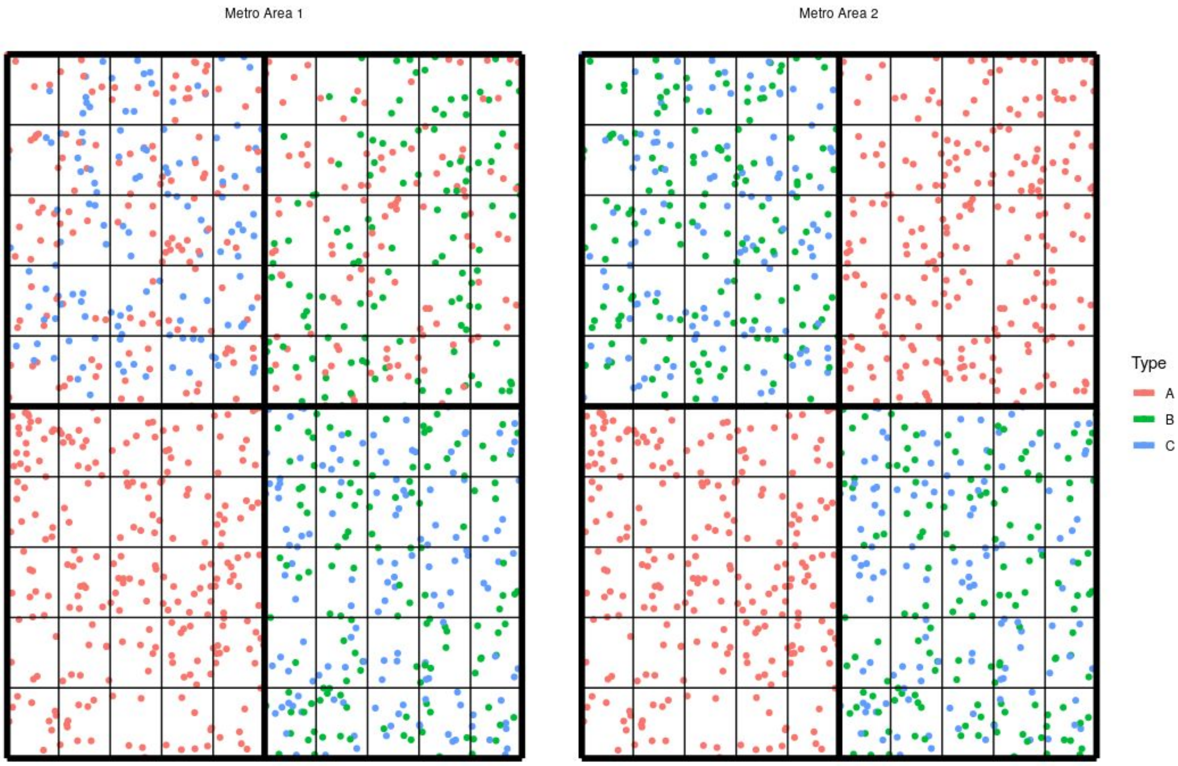

In this paper, we will examine the effect of noise-injection and the 2020 DAS on the decomposition of Theil's H for Census Core-Based Statistical Areas (CBSAs). For each CBSA, we measure segregation at the block group level, using the Theil's H-index for four racial and ethnic groups, Non-Hispanic Asians, Non-Hispanic Blacks, Non-Hispanic Whites, and Hispanics. We decompose Theil's H in two ways. First we calculate the proportion of segregation that is attributable to the distribution of groups between census tracts and the proportion attributed to the distribution of groups within census tracts (i.e. between block groups). Second, we calculate the percent of segregation that is attributable to separation of Non-Hispanic Whites from all other groups in the analysis. In addition, we also report results aggregated by CBSA type, either metropolitan or micropolitan region, as we anticipate locations with smaller populations will be more susceptible to bias from the noise injection process.

Simulating Noise-Injection

To answer our first research question about the extent to which noise-injection impacts levels of between- and within-tract segregation as epsilon varies, we conduct a simulation experiment, adding noise to published census tables in a manner that replicates the initial noise-injection step of the 2020 DAS. In this simulation we create cross tabulation for the racial and ethnic groups considered in our analysis for each block group in the United States. For each cell in a table we inject noise from a double sided exponential function equivalent to a specified epsilon value and a set of tables across the United States is created for each epsilon noise injection. The epsilon values range from .005-.125. It should be noted that while the DAS does inject noise in this manner, there are also a series of post processing steps undertaken as well to ensure geographic harmony of population estimates and that no value below zero is created. The post processing algorithm is made publically available by the US Census Bureau, however, because of the computational costs of running the DAS, it is infeasible to rerun this process across the US many times for different levels of epsilon. As we are unable to replicate this post-processing, we instead simply round all values less than zero to zero.

We calculate Theil's H index, segregation proportion attributable to Non-Hispanic White and all other groups, as well as segregation proportion attributable to census tract differences for all CBSA's. We repeat this process 100 times to create process uncertainty such that for every CBSA and value of epsilon we have 100 simulated data sets from which we may calculate a noise injected measure of segregation and examine how different the results are from the true measure of segregation for that area. Using this process uncertainty we calculate the 2.5% and 97.5% percentiles of the simulation for each metric, CBSA, epsilon value from which we calculate the proportion of CBSA's which have a true segregation metric which overlaps with the interval. We also calculate bias percentage points from the mean value of Theil's H index from simulations, $\bar{x}$, and the true value of Theil's H, μ, in the following manner.

$$Bias\ \% = \frac{\bar{x} - \mu}{\mu}$$

Analyzing the End-to-End Demonstration Data

To answer our second and third research questions, we make use of the 2010 end-to-end demonstration data released by the Census Bureau in May. These data allow comparison of the "true" 2010 Census tables with a version of the 2010 tables that have undergone noise-injection and postprocessing in the DAS. Our second research question asks how the demonstration data differ from the true 2010 data when we perform a decomposition of Theil's H. Again we calculate Theil's H index and our proportions of interest. We assess systematic bias by assessing the difference between the two measures for each CBSA in a two sample paired T-test for each segregation metric.

Lastly, we are interested in examining how the DAS may alter our understanding of segregation trends over time both in total measured segregation, the geographic level at which segregation is occuring, and how racial and ethnic groups may be differentially segregated. We again calculate our measures of segregation for all CBSA's this time using the 2000 census data. Once calculated we take the difference between the 2010 census segregation measures and the 2000 census segregation measures for each CBSA to get a measure of segregation change over time. We repeat the process this time using the 2010 DAS end to end test and the 2000 census such that we have two sets of segregation metric differences. We compare the two sets of differences using a two sample paired T-test for each segregation metric.

RESULTS

We consistently find that the process of noise injection, both in the simulation and the DAS implemented by the Census, leads to biased estimates of both total segregation as well as the proportion of segregation that is attributable to small vs large scale segregation and Non-Hispanic White segregation.

In our simulation analysis we find that measures of total segregation, as measured by the Theil's H index, are biased upwards in the noise injection process. A majority, 93.4% of micropolitan and 88.9% of metropolitan, of the CBSA's saw an upward bias in their total segregation metric, as measured by the Theil's H index, from the noise injection process across all simulations. The bias a CBSA was measured to have is strongly correlated to that location's true measure of Theil's H, with lower values of Theil's H being more susceptible to upward bias, figure 3. As expected, bias was greater for lower values of epsilons, averaging at 2.06 and 1.01 for micropolitan and metropolitan areas respectively at the minimum value of epsilon (.005) and .24 and .06 for micropolitan and metropolitan areas respectively at the maximum value of epsilon (.125).

Repeating a noise injection process for each pair of CBSA and value of epsilon allowed us to calculate process uncertainty from which we can see for 100 iterations of a noise injection process how many times the true value fell between the 2.5 and the 97.5 percentiles of the simulations. Unsurprisingly, coverage was greater for metropolitan areas than micropolitan areas and coverage increased when epsilon was greater. Even for the greatest value of epsilon, however, coverage was quite low 33.2% for both micropolitan and metropolitan areas. For the lowest values of epsilon coverage was 4.41% and 2.89% for micropolitan and metropolitan areas respectively.

We find consistent bias of both measures of segregation attributable to between census tract and non-hispanic White to other groups. We find both measures of segregation to be downward biased in our simulation results for micropolitan and metropolitan areas. Percent of segregation attributable to between census tracts was found to be X and Y biased while non-hispanic White to other groups bias was found to be X and Y, for micropolitan and metropolitan areas respectively.

Figure 3: Simulation Bias as a Function of CBSA's true Segregation Measure

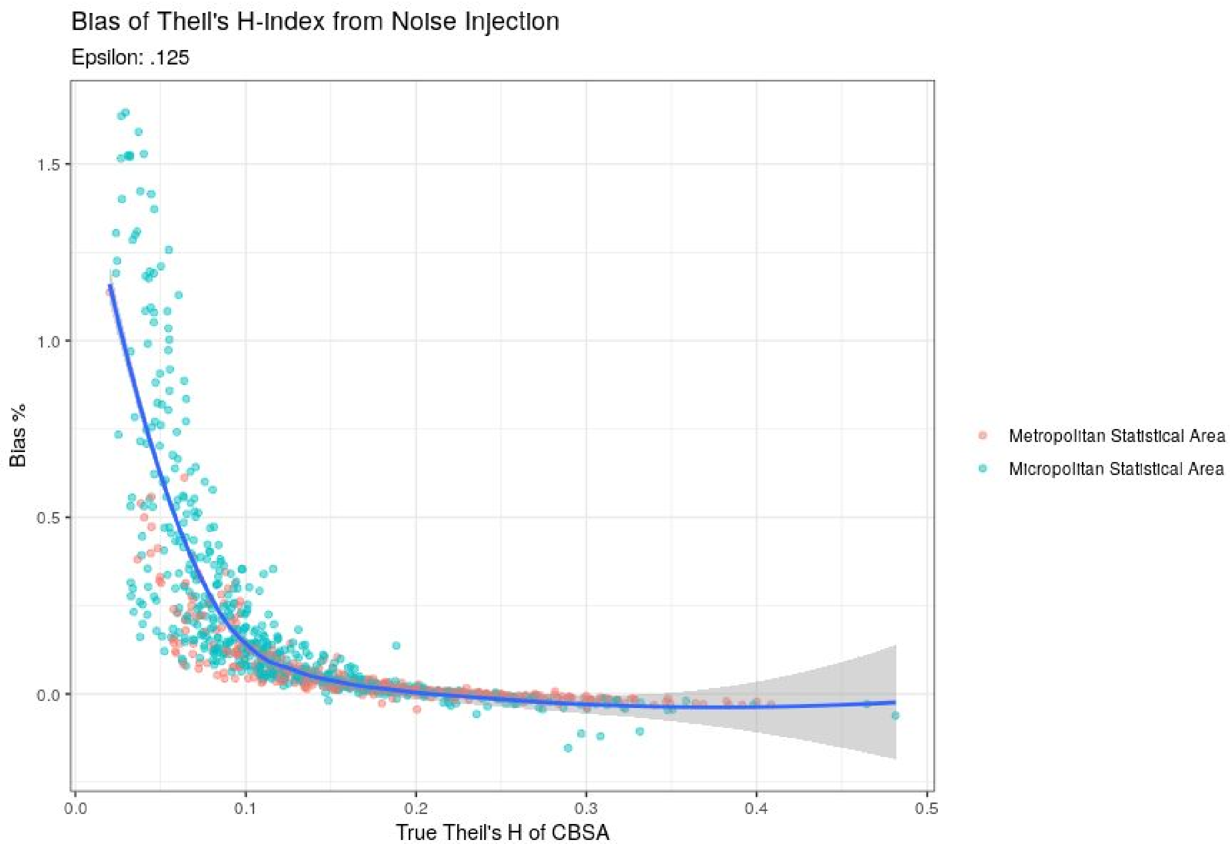

When we analyze the 2010 differential privacy end to end test we find similar results. Total segregation was found to be upward biased, 7.0% and 1.8%, while the percent attributable to between census tract, -2.6% and -2.1%, and non-hispanic White to other groups, -5.6% and -2.4%, were both downward biased, with each pair of values being for micropolitan and metropolitan areas respectively. Scatter plot measures for all areas for each metric shown in Figure 4. For all sets of t-tests run, total segregation, percent attributable to between census tract, and percent attributable to non-Hispanic White from other groups, we find a significant difference between the end to end test and the actual 2010 Census release data.

Figure 4: Scatter plot of CBSA Segregation Measures

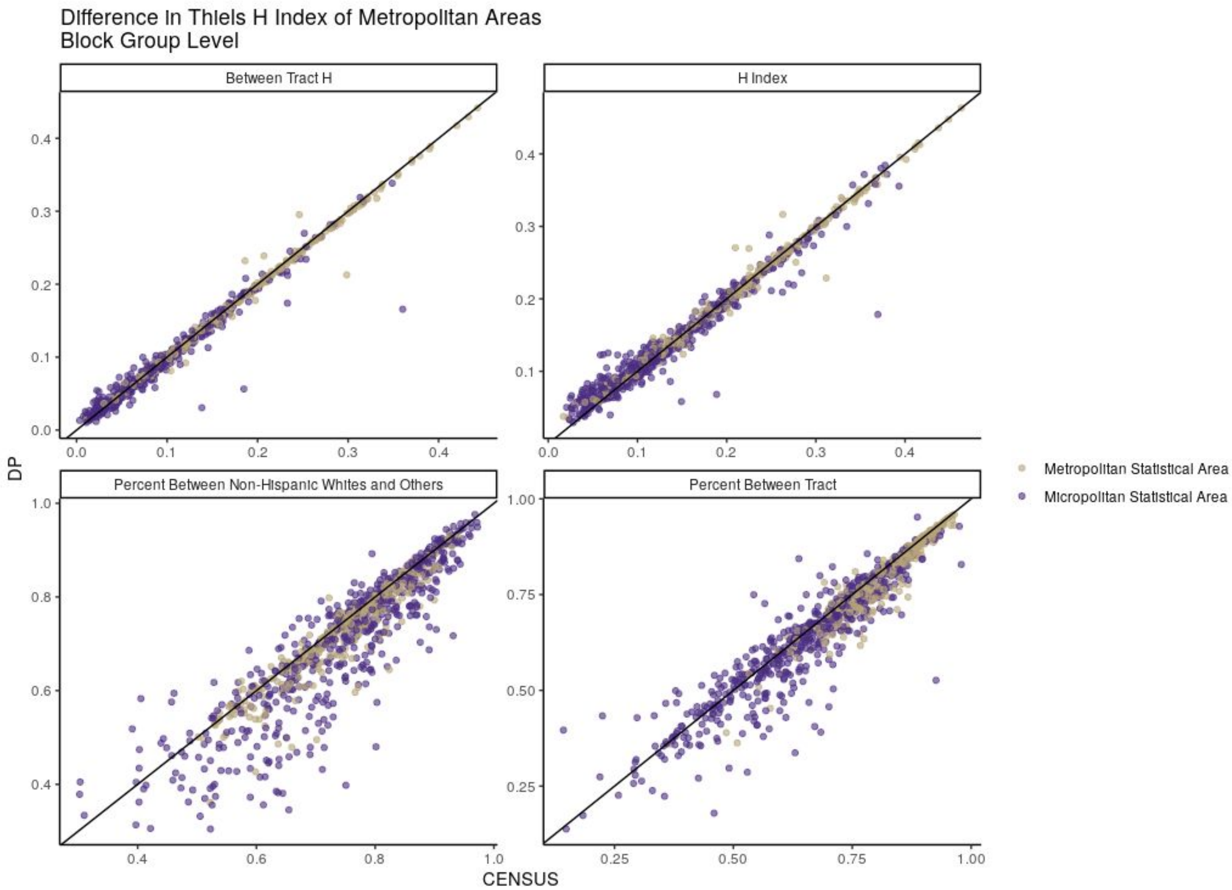

When comparing the change in segregation over time between 2000 and 2010 using the true 2010 and and the 2010 DP end to end test we find that their is again a significant difference for measures of change in total segregation ($p <= .05$), percent attributable to between geographies ($p <= .05$), and percent attributable to non-Hispanic White other segregation ($p <= .05$). Most notably for micropolitan areas we find that the while in the true release data, the majority of areas are increasing in their proportion of segregation attributable to between census tracts(57%) while in the end to end test, the pattern is masked (48%) such that we see a decline in the importance of census tract level segregation.

DISCUSSION AND CONCLUSION

Our analysis shows that noise injection similar to the DP process systematically biases results of measures of total segregation, as well as the proportion of segregation attributable to geographic and compositional factors. Two points are of great concern to segregation scholars. First, in both DP noise injection simulations and the 2010 end-to-end demonstration data, the segregation betweennon-Hispanic Whites and non-Whites is consistently downward biased.

Given that considerable focus of segregation research is centered on the geographic isolation of non-Hispanic Whites (Ellis et al., 2018), the dampening of this large portion of segregation could have strong implications for our ability to accurately observe and understand segregation trends. Second, noise injection appears to increase total segregation by increasing the amount of segregation that occurs at small geographic scales. Given that in recent years segregation at larger geographic levels has increased in its relative importance, DP may mask how the geographic pattern of segregation is unfolding.

In addition to the risks that the 2020 DAS poses to the ability of researchers to conduct substantive research on segregation in the United States and its trajectory over time, the 2020 DAS could have negative implications for the validity of census geography. Census tracts are heavily used in social science research, and while data published at the tract level will be less affected by the DAS than data published at finer geographic levels, these sub-tract data play a key role in establishing the validity of tract data. Census tracts are fundamentally arbitrary units of aggregation. They serve a statistical purpose, not an administrative or political purpose . If the 2020 DAS makes residential patterns by race appear clustered at the sub-tract level, then tracts -- as they are drawn -- will seem less useful and data published at the tract level, less meaningful.

As such, the findings presented in this paper call for deeper engagement with noise-injection mechanisms as they relate to the real-world population processes that produce the patterns observed in the data. When large amounts of noise is injected into block- or block group-level counts of race and ethnicity, the resulting distributions no longer resemble those which are plausible under normal patterns of residential mobility: whole clusters of population are added and subtracted in a way that would require great shifts in population and housing stock. Work is needed to provide guidance for how to interpret noisy data published at fine geographic scales.

REFERENCES

Abowd, J. M. (2018, July). The US Census Bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 2867-2867).

Ard, K. (2016). By all measures: An examination of the relationship between segregation and health risk from air pollution. *Population and Environment*, *38*(1), 1-20.

boyd, d. (2019). *Balancing data utility and confidentiality in the 2020 US census*. Technical report, Data and Society, New York, NY.

Dwork, C. (2008, April). Differential privacy: A survey of results. In *International conference on theory and applications of models of computation* (pp. 1-19). Springer, Berlin, Heidelberg.

Ellis, M., Wright, R., Holloway, S., & Fiorio, L. (2018). Remaking white residential segregation: metropolitan diversity and neighborhood change in the United States. *Urban geography*, *39*(4), 519-545.

Farrell, C. R., & Lee, B. A. (2011). Racial diversity and change in metropolitan neighborhoods. *Social Science Research*, *40*(4), 1108-1123.

Fischer, C. S., Stockmayer, G., Stiles, J., & Hout, M. (2004). Distinguishing the geographic levels and social dimensions of US metropolitan segregation, 1960–2000. *Demography*, *41*(1), 37-59.

Fowler, C. (2016). Segregation as a multiscalar phenomenon and its implications for neighborhood-scale research: The case of South Seattle 1990–2010. *Urban geography*, *37*(1), 1-25.

Garfinkel, S. L., Abowd, J. M., & Powazek, S. (2018, January). Issues encountered deploying differential privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society* (pp. 133-137).

Lee, B. A., Reardon, S. F., Firebaugh, G., Farrell, C. R., Matthews, S. A., & O'Sullivan, D. (2008). Beyond the census tract: Patterns and determinants of racial segregation at multiple geographic scales. *American Sociological Review*, *73*(5), 766-791.

Manley, D., Jones, K., & Johnston, R. (2019). Multiscale Segregation: Multilevel Modeling of Dissimilarity—Challenging the Stylized Fact That Segregation Is Greater the Finer the Spatial Scale. *The Professional Geographer*, *71*(3), 566-578.

Massey, D. S., & Denton, N. A. (1988). The dimensions of residential segregation. *Social forces*, *67*(2), 281-315.

Morrill, R. L. (1991, February). On the measure of geographic segregation. In *Geography research forum* (Vol. 11, No. 1, pp. 25-36).

Napierala, J., & Denton, N. (2017). Measuring residential segregation with the ACS: How the margin of error affects the dissimilarity index. *Demography*, *54*(1), 285-309.

Reardon, S. F., Matthews, S. A., O'sullivan, D., Lee, B. A., Firebaugh, G., Farrell, C. R., & Bischoff, K. (2008). The geographic scale of metropolitan racial segregation. *Demography*, *45*(3), 489-514.

Wright, R., Ellis, M., Holloway, S. R., & Wong, S. (2014). Patterns of racial diversity and segregation in the United States: 1990–2010. *The Professional Geographer*, *66*(2), 173-182.

Wong, D. W. (1993). Spatial indices of segregation. *Urban studies*, *30*(3), 559-572.

Zhang, W., & Logan, J. R. (2017). The emerging spatial organization of the metropolis: Zones of diversity and minority enclaves in Chicago. *Spatial demography*, *5*(2), 99-122.