Discrete Additive Noise Mechanisms for Differential Privacy:

Geometric and Gaussian Additive Noise

Jonathan Buttle

California Department of Finance

Demographic Research Unit

Noise Generation in Differential Privacy

- An Overview of Differential Privacy
- Noise Generation – Geometric and Gaussian Distributions
- Privacy Protected Data – Results
- Discussion
- References/Further Reading

An overview of differential privacy

- Differential privacy (DP) is a property of algorithms for answering queries. An algorithm is considered differentially-private for a given epsilon and delta (ε, δ) if, for two databases that differ by one record, it satisfies:

$$\Pr[A(D) \in T] \leq \exp(\varepsilon) \Pr[A(D') \in T] + \delta$$

- DP works by injecting statistically calibrated "noise" into a query or statistic derived from the underlying micro data.
- This "noise" is drawn from predetermined probability distributions that have characteristics consistent with DP.
- The literature identifies two fundamental probability distributions used to generate the "noise" – the Laplace distribution and the Gaussian distribution.
- Most mechanisms (including the Census Bureau's Disclosure Avoidance System (DAS) demonstration engine) employ the discrete version of these distributions (to generate integers) – the two-sided geometric distribution (Laplace) and the discrete Normal distribution (Gaussian). This is to address computational and security difficulties inherent in using continuous distributions.

Continued: An overview of differential privacy

- In general, each of the probability distributions are associated with one of two types of differential privacy – pure and approximate.
  - Pure differential privacy is the case when $(\varepsilon, \delta = 0)$. The mechanism is exclusively parameterized by $\varepsilon > 0$ (the privacy loss budget), which controls how much privacy loss an individual can suffer when a computation is performed using their data. The probability of a privacy loss of any individual exceeding ε is 0;
  - Approximate differential privacy $(\varepsilon, \delta > 0)$ is a relaxation of pure differential privacy that provides a less robust privacy guarantee. It guarantees that the probability of a privacy loss of any individual exceeding ε is bound by δ;
  - The parameter δ can be thought of as the probability that a catastrophic privacy breach/data release occurs in the presence of DP;
  - By contrast, 1 - δ is the probability that that the mechanism is ε-differentially private.
- The geometric distribution satisfies the requirements of pure DP, while the Gaussian distribution satisfies approximate DP.

Noise Generation – Geometric and Gaussian Distributions

- The geometric and Gaussian distributions differ in two ways: the DP condition satisfied, and the variance generated by each mechanism.
  - The geometric distribution satisfies pure DP (which is a stronger privacy condition than is approximate differential privacy), but with noise generated with a wider variance (its tails decay at a subexponential rate, $e^{-\varepsilon m}$);
  - The Gaussian distribution satisfies approximate DP and its tails decay at a subguassian rate ($e^{-m^2/2\sigma^2}$), which results in a smaller noise variance.
- Which distribution is better depends on whether privacy or utility (variance and accuracy) of the estimates is more important.
- Note that if pure differential privacy or approximate differential privacy for a very small $\delta > 0$ is needed, then the geometric distribution is preferred; while for the opposite situation the Gaussian distribution is preferred.

Continued: Noise Generation – Geometric and Discrete Gaussian Distributions

- The concept of concentrated DP was introduced to address some of the analytic and computation concerns raised by approximate DP [DR16].
- A random mechanism satisfies concentrated DP if the privacy loss has a small mean and is subgaussian.
- Concentrated DP was refined with the introduction of zero-concentrated DP (zCDP) by Bun and Steinke.
- They showed that there is a relationship between pure DP and zCDP: ε-DP implies $\left(\frac{1}{2}\varepsilon^2\right)$-zCDP [BS16].
- Figure 1 demonstrates the relationship between geometric and discrete Gaussian mechanisms for $\varepsilon = 0.5$.

Figure 1

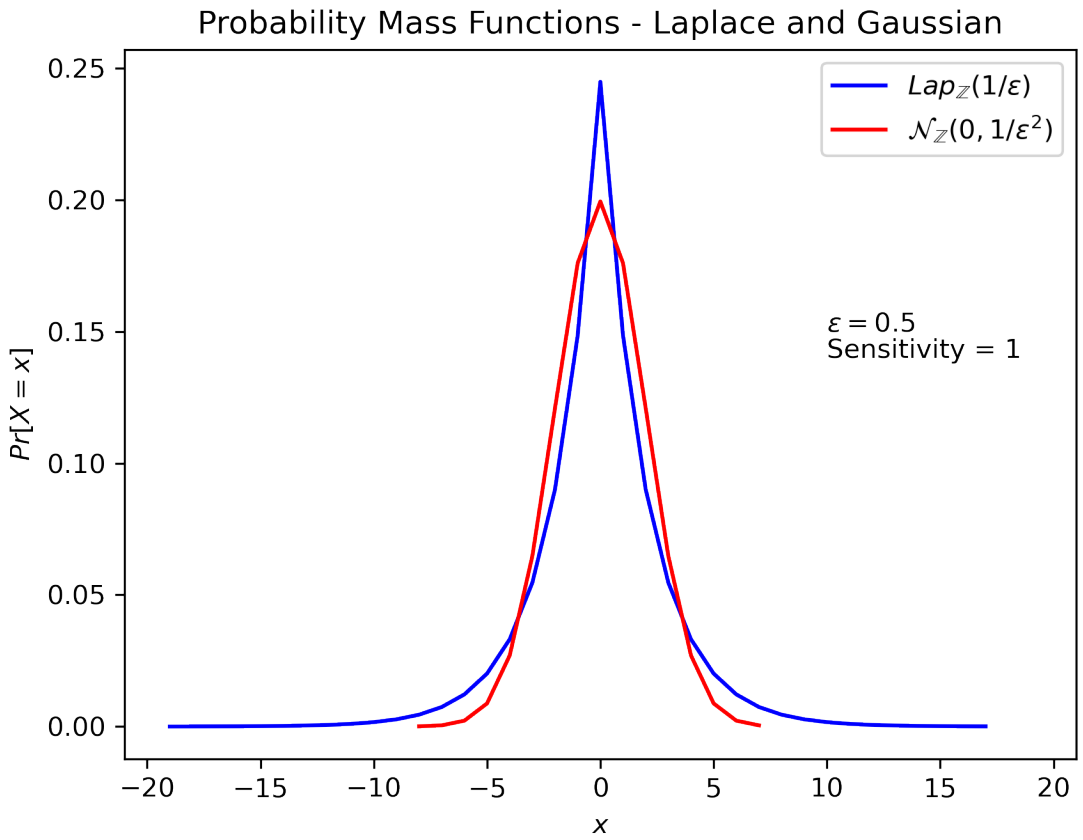

Continued: Noise Generation – Geometric and Discrete Gaussian Distributions

- Note that the noise mechanism for zCDP is defined by two parameters (ε, δ), whereas the noise mechanism for pure DP is solely defined by $\varepsilon > 0, (\varepsilon, \delta = 0)$.
- There is no consensus for determining the value of δ – to insure against catastrophic events, the literature recommends setting $\delta \ll n^{-1}$. For a modest-sized dataframe (694,675 records => Alaska), $\delta <$ 1.44e-6.
- Figure 2 illustrates how the choice of δ impacts the variance of the noise generated by the zCDP mechanism.

Figure 2

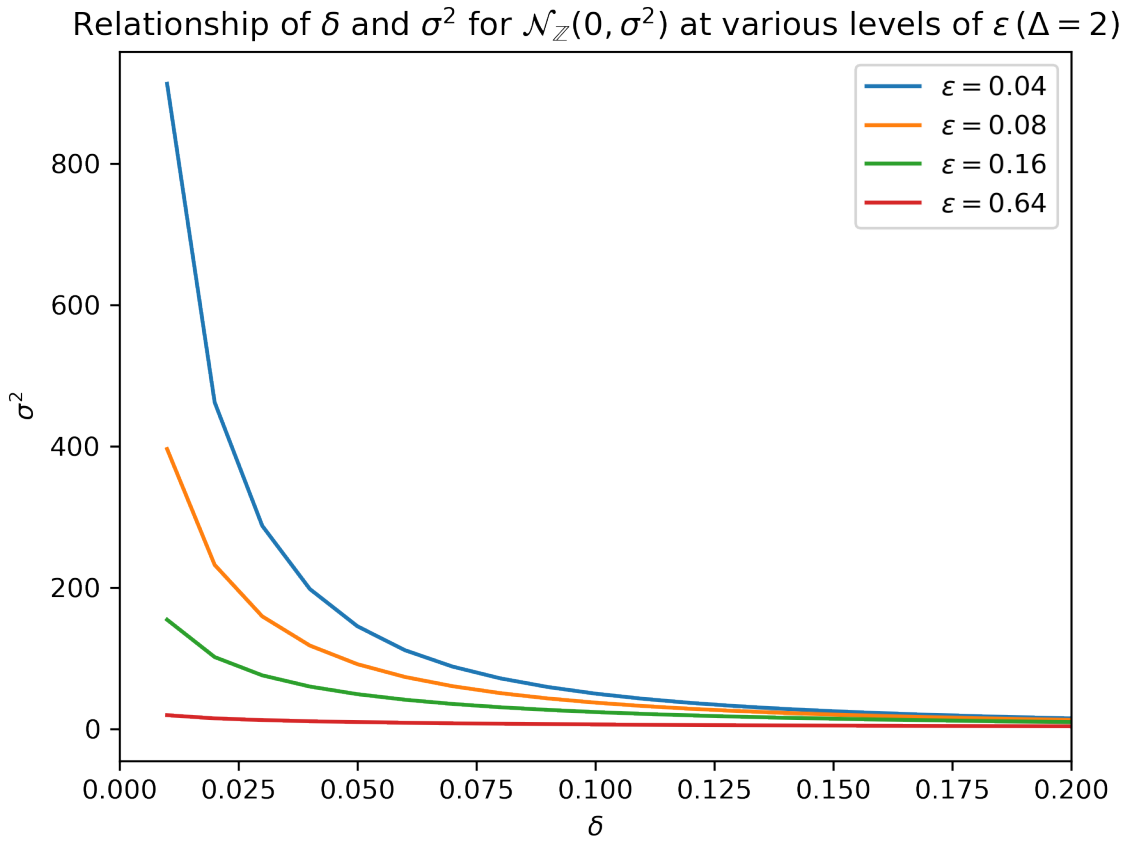

Continued: Noise Generation – Geometric and Discrete Gaussian Distributions

- The figures below further illustrate how the choice of delta impacts the magnitude of noise generated by the zCDP mechanism, compared to two-sided geometric mechanism.

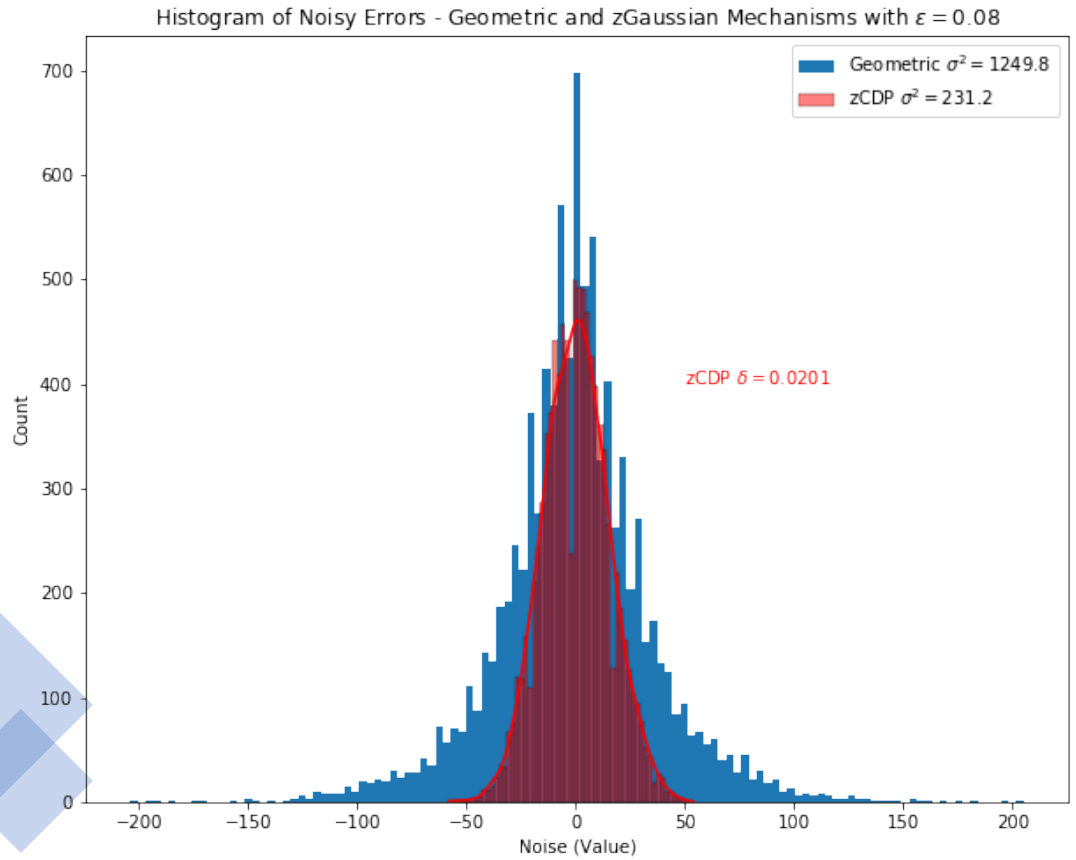

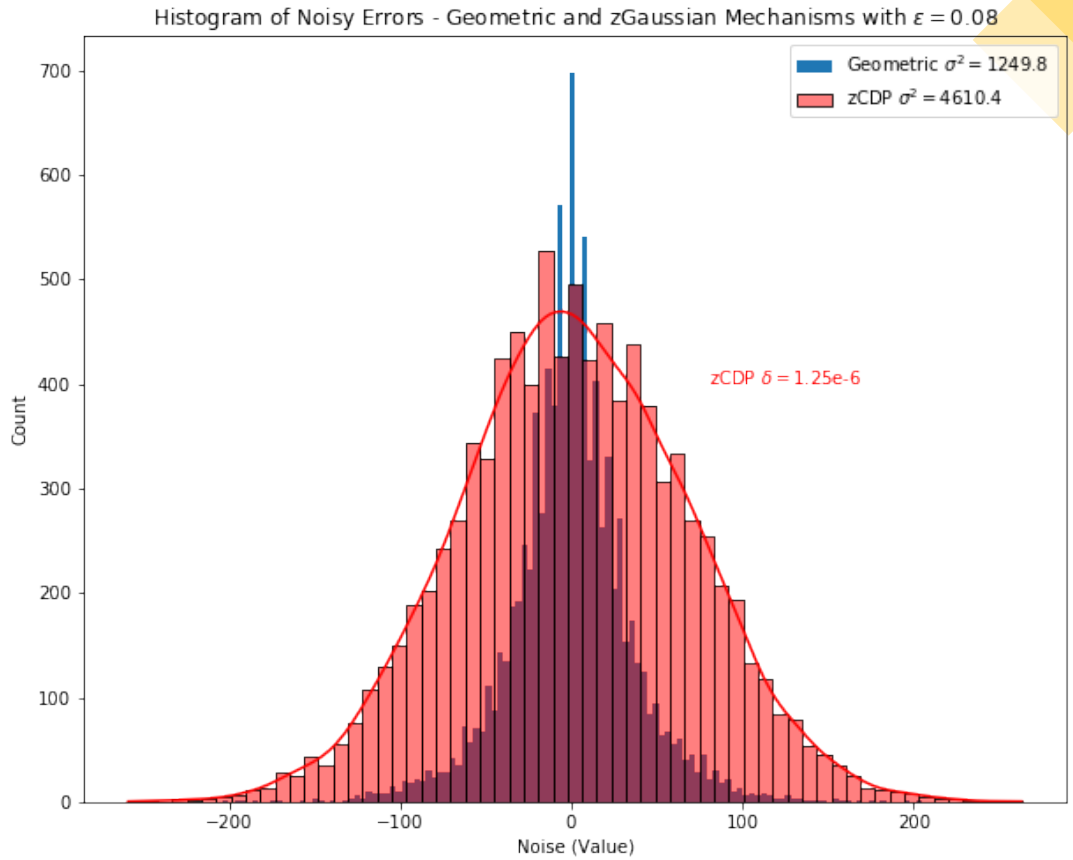

Continued: Noise Generation – Geometric and Discrete Gaussian Distributions

- The double-sided geometric distribution takes the form of:

$$\forall x \in Z,\ P[X = x] = \frac{e^{1/t}-1}{e^{1/t}+1} \cdot e^{-|x|/t}$$ and is symmetrical around 0.

- The discrete Gaussian distribution takes the form of:

$$\forall x \in Z,\ \ P[X = x] = \frac{e^{-(x-\mu)^2/2\sigma^2}}{\sum_{y \in Z} e^{-(y-\mu)^2/2\sigma^2}}$$, where μ is assumed to be 0 for purposes of DP.

- The discrete Gaussian has properties similar to those of the continuous Gaussian [CKS20]:
  - The privacy guarantee is almost equal to the one offered by the continuous distribution;
  - The discrete distribution offers the same or slightly better accuracy than does the continuous distribution; and
  - Functionally, it is practical to sample the discrete Gaussian on a finite computer.

Continued: Noise Generation – Geometric and Discrete Gaussian Distributions

- The zCDP algorithm used in this simulation was developed by Thomas Steinke at IBM [Dga].
- Here's how the algorithm works:
    1. The first step is to choose the values of $\varepsilon \; and \; \delta$;
    2. The algorithm then computes a value for ρ such that ρ-CDP implies (ε, δ)-differential privacy;
    3. The parameter ρ is used to determine the value of σ^2 (Adding samples derived from either a continuous or discrete Gaussian distribution with parameter σ^2 provides ρ-CDP for $\rho = \Delta/2\sigma^2$ (where Δ represents the sensitivity of the mechanism), the value of σ^2 is determined by $\Delta/2\rho$).
- The Python code for the discrete Gaussian mechanism was adapted to the Census Bureau's DAS 2020 program.
- The simulation was run on a single-node stand-alone computer.
- The synthetic data microdata is based on the U.S. Synthetic Population 2010 (version 1) developed by RTI International [RTI]. It includes population and housing records for Alaska (694,675 records) and eight counties in California (Alpine, Fresno, Humboldt, Los Angeles, Riverside, Sacramento, San Francisco, and Santa Barbara – 15,568,335 records).

Privacy Protected Data - Results

- The charts show DAS errors for tract data for three major race groups (White, Black, and American Indian Native Alaskan).
- All three runs set $\varepsilon = 0.04$.
- The ”noisy” estimates were generated using three mechanisms – Geometric, Gaussian $\delta = 0.0201$, and Gaussian $\delta = 5.44\text{e-}8$.
- The $\delta = 0.0201$ Gaussian had the smallest variance followed by Geometric, and then the $\delta = 5.44\text{e-}8$ Gaussian.

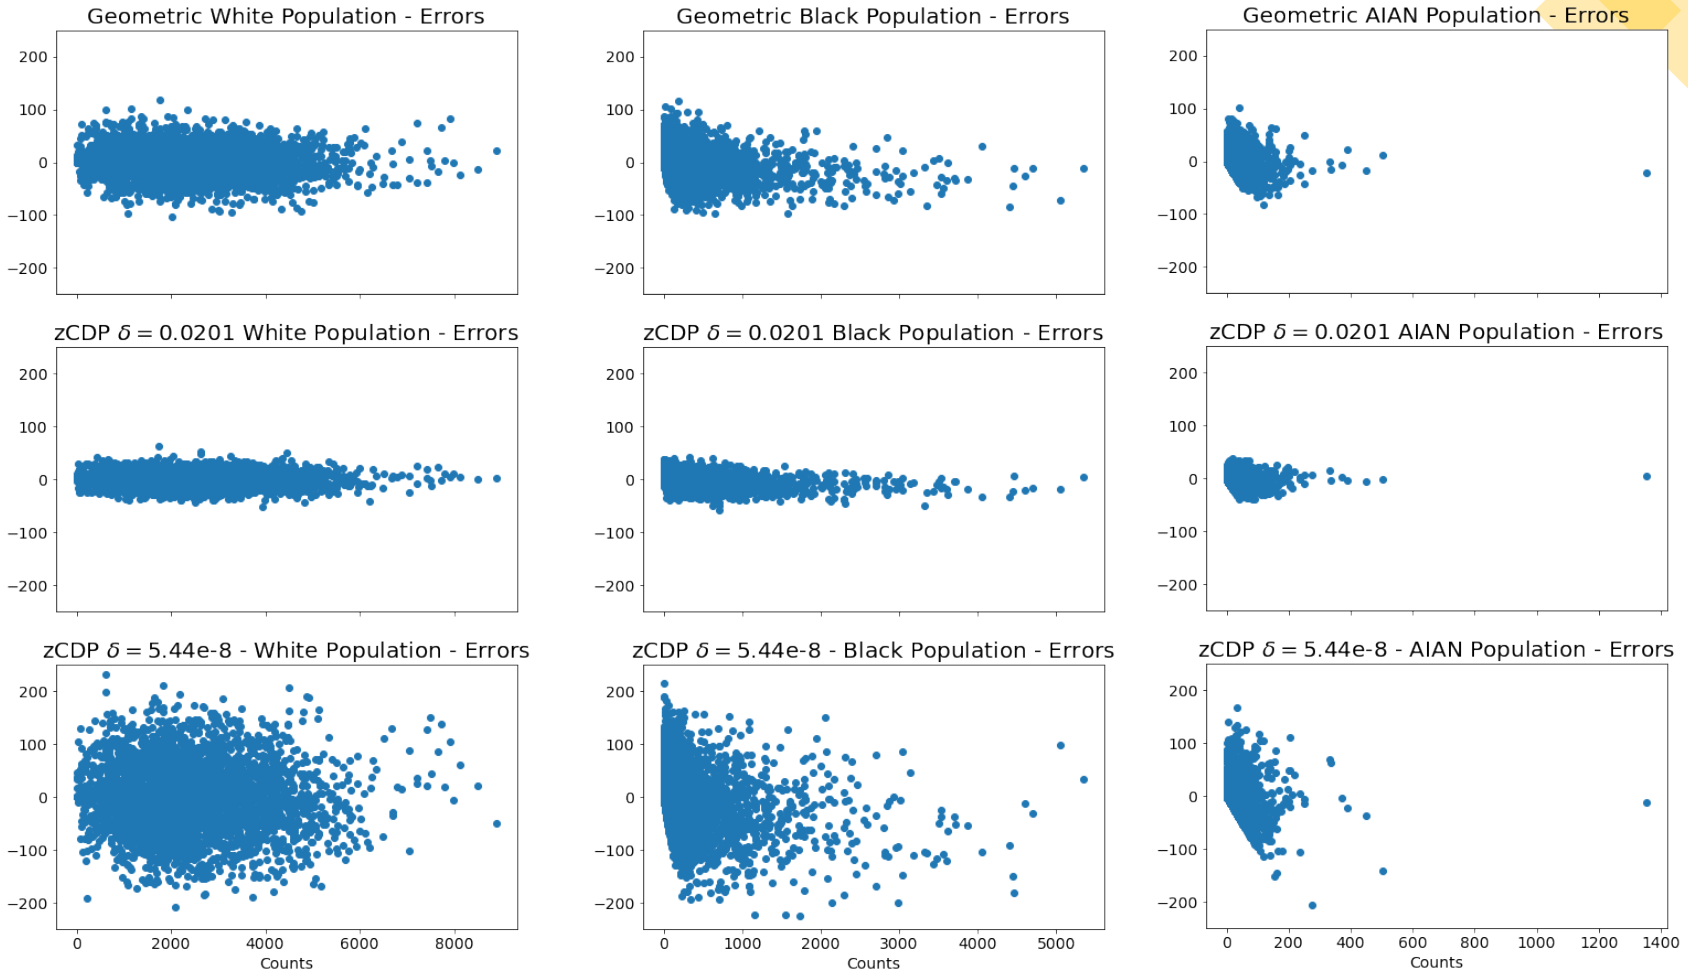

Continued: Privacy Protected Data - Results

- The charts show DAS errors for tract data for four major race groups (Asian, Native Hawaiian and Other Pacific Islander, Other, and Two or More).

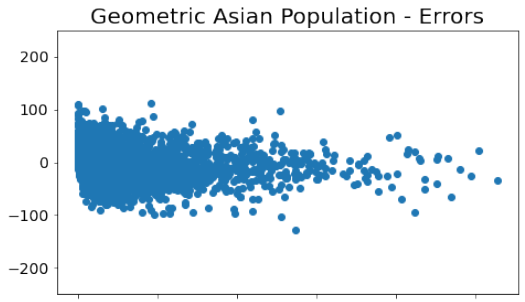

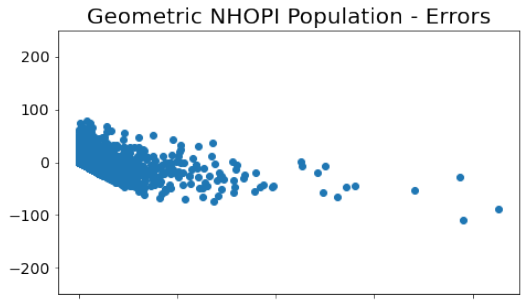

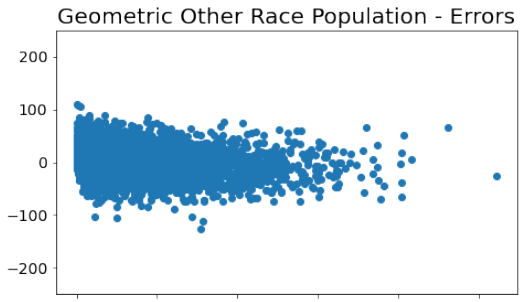

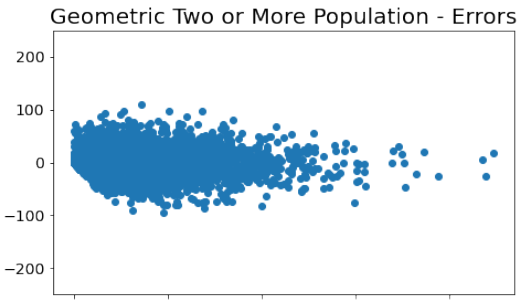

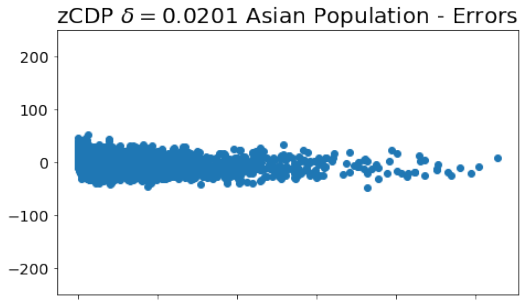

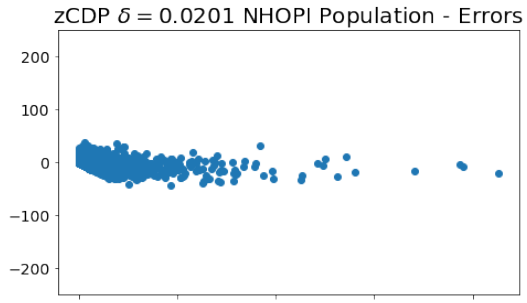

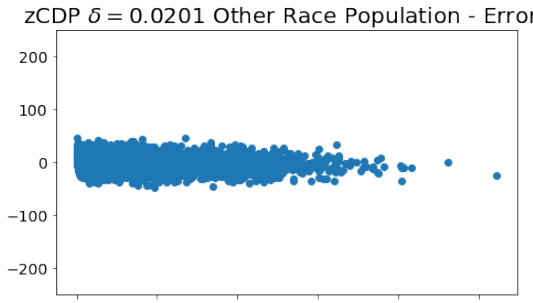

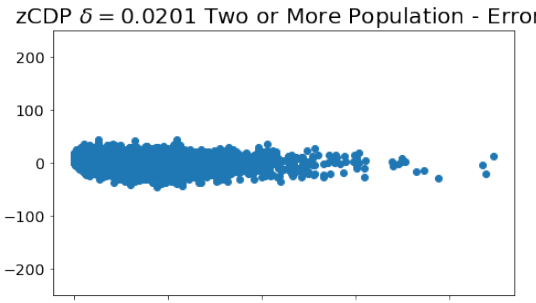

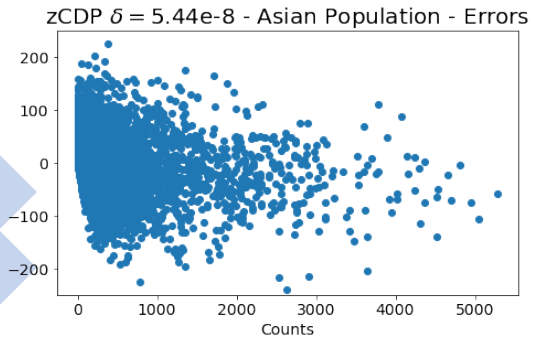

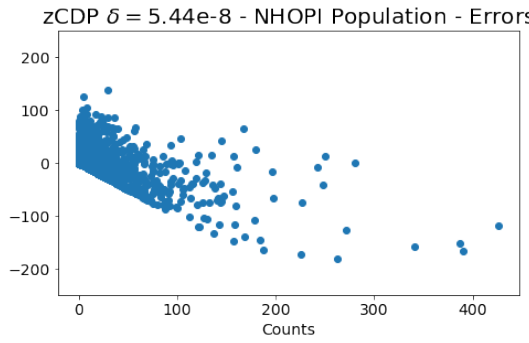

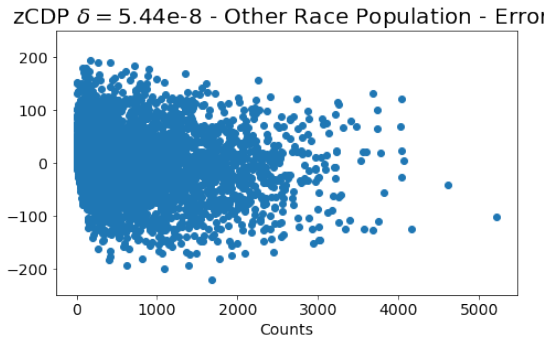

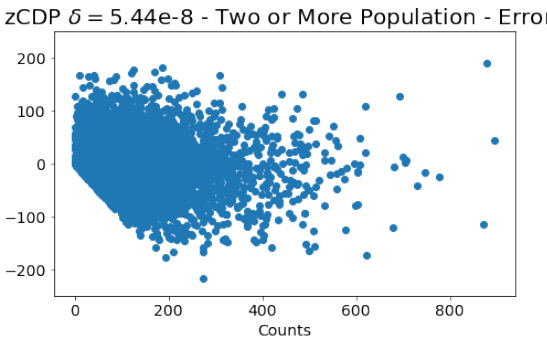

Continued: Privacy Protected Data - Results

- The charts show DAS errors for block data for three major race groups (White, Black, and American Indian Alaskan Native);
- All three runs set $\varepsilon = 0.04$;
- The "noisy" estimates were generated using three mechanisms – Geometric, Gaussian $\delta = 0.0201$, and Gaussian $\delta =$ 5.44e-8;
- The $\delta = 0.0201$ Gaussian had the smallest variance followed by Geometric, and then the $\delta =$ 5.44e-8 Gaussian;

Continued: Privacy Protected Data - Results

- The charts show DAS errors for block data for four major race groups (Asian, Native Hawaiian and Other Pacific Islander, Other, and Two or More).

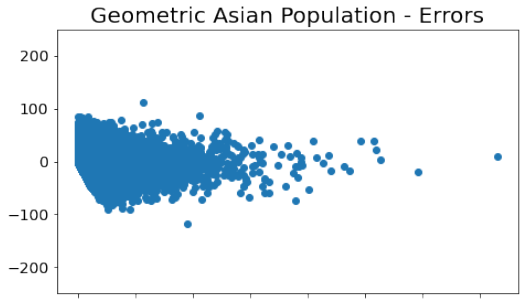

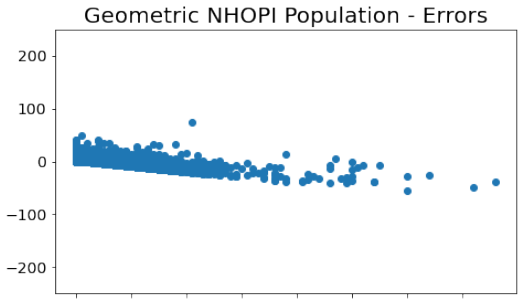

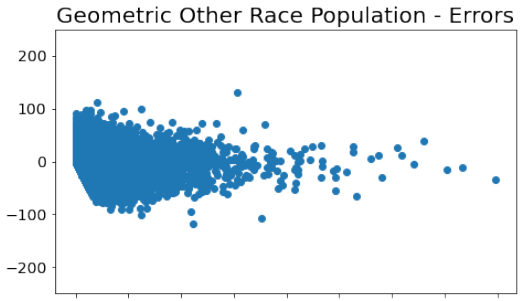

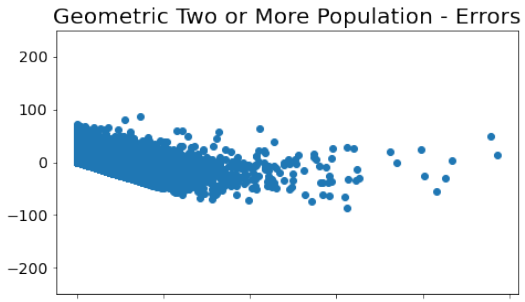

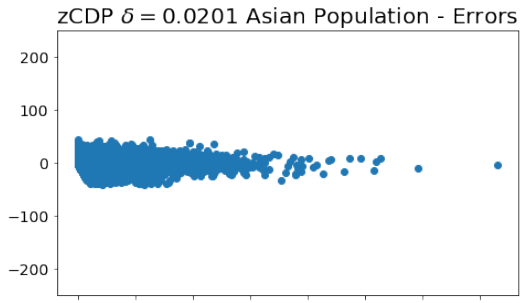

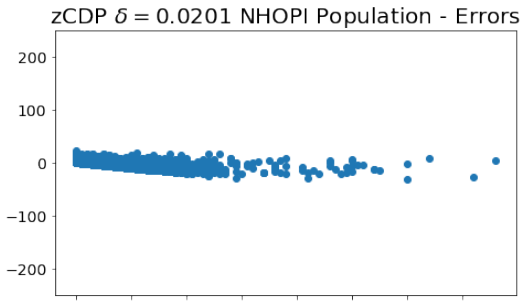

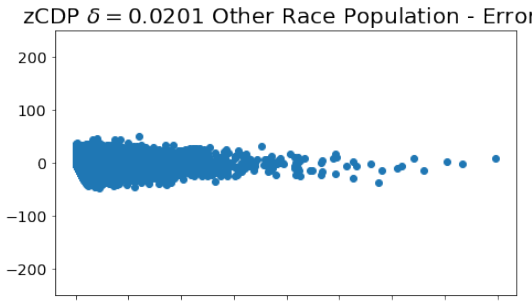

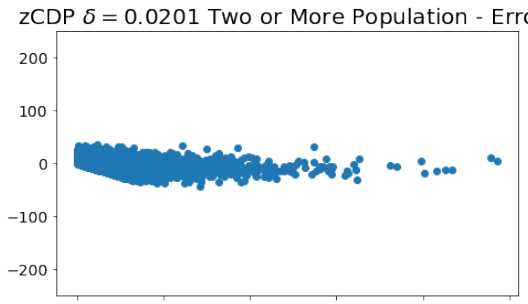

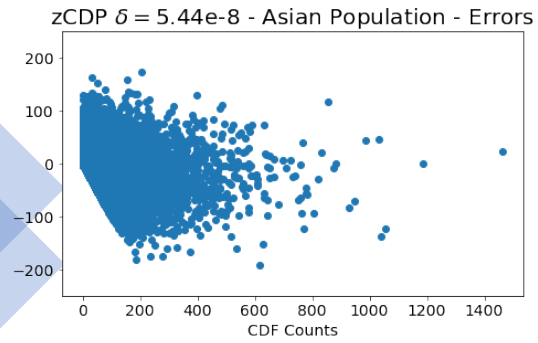

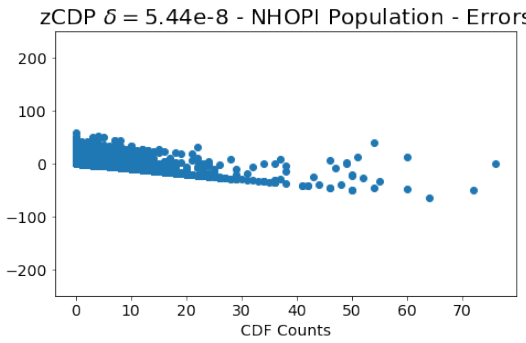

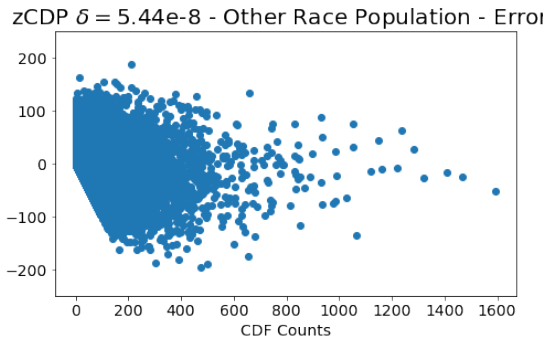

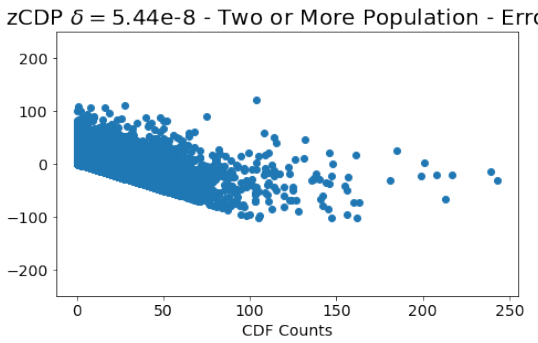

Continued: Privacy Protected Data - Results

- These charts show the mean absolute error for Census blocks and tracts for the eight California counties.
- All three runs set $\varepsilon = 0.1$ for blocks and for tracts.

	Blocks	Mean Absolute Error		
		Geometric	Gaussian	
			δ = 0.0201	δ = 5.44e-6
All Blocks	167,400	17.61	8.31	31.41
Blocks with total population less than 100	121,538	15.15	7.55	25.01
Blocks with total population 100 to 499	41,894	23.02	10.00	45.63
Blocks with total population 500 to 999	3,289	34.64	13.21	74.42
Blocks with total population 1,000 or more	679	41.27	16.05	92.97

	Tracts	Mean Absolute Error		
		Geometric	Gaussian	
			δ = 0.0201	δ = 5.44e-6
All Tracts	3,606	23.74	10.62	53.96
Tracts with total population less than 100	7	30.86	9.86	57.29
Tracts with total population 100 and 999	21	26.09	10.74	61.50
Tracts with total population 1,000 and 9,999	3,544	23.70	10.61	53.91
Tracts with total population 10,000 or more	34	24.76	11.90	48.95

Continued: Privacy Protected Data - Results

- These charts show the number of blocks and tracts that switched from majority (50%+) White to majority all other races and vice versa for the eight California counties.
- All three runs set $\varepsilon = 0.04$ for blocks and tracts.

	Blocks	Geometric		Gaussian δ = 0.0201		Gaussian δ = 5.44e-6	
		White to Minority	Minority to White	White to Minority	Minority to White	White to Minority	Minority to White
All Blocks	167,400	29,301	15,177	21,577	12,008	34,384	18,295
Blocks with total population less than 100	121,538	26,872	12,157	20,314	10,287	29,812	13,966
Blocks with total population 100 to 499	41,894	2,398	2,878	1,247	1,647	4,484	4,062
Blocks with total population 500 to 999	3,289	26	129	16	67	80	235
Blocks with total population 1,000 or more	679	5	13	0	7	8	32

	Tracts	Geometric		Gaussian δ = 0.0201		Gaussian δ = 5.44e-6	
		White to Minority	Minority to White	White to Minority	Minority to White	White to Minority	Minority to White
All Tracts	3,606	13	15	10	6	40	39
Tracts with total population less than 100	7	1	1	1	0	2	0
Tracts with populations between 100 and 999	21	2	0	1	0	5	0
Tracts with populations between 1,000 and 9,999	3,544	10	14	8	6	33	39
Tracts with populations of 10,000 or more	34	0	0	0	0	0	0

Discussion

- These simulations demonstrate that the discrete Gaussian (zCDP) mechanism has better utility and accuracy than does the geometric mechanism.

- However, the gains in accuracy and utility come at the expense of less privacy. The privacy loss can be mitigated by reducing the value of δ.

- Even with the accuracy gains provided by the discrete Gaussian mechanism, small-population areas still are disproportionality impacted by additive noise compared with areas that have populations greater than the magnitude of the tails.

Contact Information

Jonathan Buttle – jonathan.buttle@dof.ca.gov

California Department of Finance
Demographic Research Unit
dof.ca.gov/forecasting/demographics/

(916) 323-4086

References/Further Reading

References –

- [BS16] Mark Bun, and Thomas Steinke. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In: Hirt M., Smith A. (eds) Theory of Cryptography. TCC 2016. Lecture Notes in Computer Science, vol 9985. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-53641-4_24. https://arxiv.org/abs/1605.02065.
- [CKS20] Clement Canonne, Gautam Kamath, and Thomas Steinke. The Discrete Gaussian for Differential Privacy. In: arXiv: 2004.00010. https://arxiv.org/abs/2004.00010.
- [DR16] Cynthia Dwork, and Guy Rothblum. Concentrated Differential Privacy. In: arXiv: 1603.01887. https://arxiv.org/abs/1603.01887.
- [Dga] https://github.com/IBM/discrete-gaussian-differential-privacy.
- [RTI] 2010 RTI U.S. Synthetic Population Ver. 1.0 , RTI International. May 2014. https://fred.publichealth.pitt.edu/syn_pops.

Further Reading –

- Yu-Hsuan Kuo, Cho-Chun Chiu, Daniel Kifer, Michael Hay, and Ashwin Machanavajjhala. “Differentially private hierarchical count-of-counts histograms”. In: Proceedings of the VLDB Endowment 11.11 (2018), pp. 1509–1521. https://arxiv.org/abs/1804.00370.
- Damien Desfontaines. (2019, June 27). Almost differential privacy. [Blog post]. Retrieved from https://desfontain.es/privacy/almost-differential-privacy.html.
- Damien Desfontaines. (2019, June 27). The privacy loss random variable. [Blog post]. Retrieved from https://desfontain.es/privacy/privacy-loss-random-variable.html.