Bayesian and Frequentist Intervals under Differential Privacy for Binomial Proportions

by

Hsuan-Chen Kao

Department of Statistical Science
Duke University

Defense Date: March 26^{th}, 2025

Approved:

Jerome P. Reiter, Supervisor

David L. Banks

Surya T. Tokdar

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in the Department of Statistical Science
in the Graduate School of Duke University
2025

ABSTRACT

Bayesian and Frequentist Intervals under Differential Privacy for Binomial Proportions

by

Hsuan-Chen Kao

Department of Statistical Science
Duke University

Defense Date: March 26th, 2025

Approved:

Jerome P. Reiter, Supervisor

David L. Banks

Surya T. Tokdar

An abstract of a thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in the Department of Statistical Science in the Graduate School of Duke University
2025

Abstract

This paper compares and proposes interval inference methods for binomial proportions, true p, under differential privacy (DP) with the Laplace mechanism (ε-DP) and the discrete Gaussian mechanism (Rényi DP). We first assess the frequentist approaches, including adjusted plug-in Wald and Wilson intervals. Notably, the Wilson interval traditionally served as a more robust alternative to Wald in terms of the out-of-bound problem, which is no longer a viable substitute in this setting after adding in the Laplace noise (or discrete Gaussian noise and others) due to persistent out-of-bound issues. Additionally, we propose three alternatives: First, ε-DP Bayesian credible intervals with uniform prior and Jeffrey prior—derived from the posterior distribution of noisy observations $f(p|\hat{p}^*)$. Second is an ε-DP sampling-based interval, which is a practical alternative to the Bayesian method without MCMC. It is less complex and achieves high coverage, though the intervals can be slightly longer and somewhat conservative. Third is the ε-DP exact interval, mainly motivated by Clopper-Pearson's method, which is straightforward and easy to interpret. Lastly, for the Rényi DP mechanism, we only demonstrate the Bayesian mechanism in this thesis, as it provides a better balance between achieving the nominal coverage rate and avoiding overly conservative interval lengths, based on our evaluation of the ε-DP with the Laplace mechanism.

To bring the informative evaluation, we discuss the Laplace noise and discrete Gaussian noise controlled by the privacy parameter ε. We examine the impact of the specific pairing of varying noise levels ε and the binomial proportions p on the accuracy and coverage of these intervals. We aim to emphasize the trade-offs between privacy and statistical inference precision in differentially private data dissemination.

Dedication

To Chung, Kimberly, Tina, Judy, and Joe.

Contents

List of Tables

List of Figures

Acknowledgements

Over the past two years of my academic journey in the MSS program at Duke, I have deeply appreciated the challenges and opportunities that have pushed my research forward. The experiences allow me to know what I have a huge passion for, and I am eager to construct a model that is novel, cool, and accessible to the world, along with cultivating my ability to appreciate the beauty of statistics. With the exceptional faculty at Duke, I am closer than ever to realizing my dream.

First, I am incredibly grateful to my master's thesis advisor, Dr. Jerry Reiter. His zestful interaction in our regular meetings has been the most meaningful part of my time at Duke. His flexibility, trust, and encouragement allowed me to find a topic I was truly interested in, tolerating my weird ideas when my logic was jumping around with lost direction and guiding me when needed while giving me the space to explore on my own. His precise and strategic vision allows me to learn the mechanism of synthetic data, multiple imputations, and differential privacy under the Bayesian perspective, where many are new to me. His sincere and clear guidance has reinforced my commitment to contributing to this field, whether in industry or academia. I have a similar actuarial background to Jerry's, which allows us to have many things in common and reinforces our shared motivation to contribute to statistics after deviating from an actuarial career. Our similar dedication to advancing statistical methodology has made working with him an enriching experience, and I am deeply grateful to have him as my advisor.

Second, I thank Dr. Surya Tokdar, the first statistics professor I've met(in front of the chapel) at Duke, who is also my program advisor. His Theory of Statistical Inference course reshaped my understanding and gave me a new perspective on the relationship between Bayesian and Frequentist approaches—their complementarities as well as their conflicts. With his encouragement, I took graduate-level real analysis, which later gave me the chance to contribute, in a small way, to his research on heavy-tailed density estimation. He explained the trade-offs between comparing his Bayesian semiparametric method to the peak-over-threshold approach and taught me the reasoning behind defending a model—what

properties his team prioritized and what they chose to forgo. His sharp attention to detail and insightful ideas have constantly challenged me to think more deeply. I want to echo his words here: "Don't take any detail as given, and don't skip even the smallest part of the mathematics." I will carry this lesson with me as I take the next step as a Ph.D. student in Statistics.

Third, I am deeply grateful to Dr. David Banks. He gave me my first opportunity to engage in research through the Bayesian Graph Navigation project. When I first approached him, I had no prior research experience, but under his guidance, I learned how to structure lab meetings to share new insights and gained hands-on experience in advancing the statistical aspects of our work. Beyond research, his confidence in me led to my role as the head teaching assistant for STA 521 Predictive Inference in Fall 2024 and STA 611 Mathematical Statistics in Spring 2025. We had many meaningful discussions about course design, and he encouraged me to take bold steps in structuring the class and leading the TA team.

Lastly, I express my deepest thanks to Dr. Edric Tam. He is the best Ph.D. mentor I have ever met. I really enjoyed the time with him, especially the thoughtful conversations about the frontiers of statistics, application strategies, and school choices. I will never forget the first day I attended his boot camp at Duke—it was the starting point of our relationship that will continue to evolve, with many more meaningful moments ahead.

1. Introduction

The primary mission of federal statistical agencies is to share data responsibly with the public (Reiter, 2019) to support research across various domains. The binomial proportion is one kind of common data structure; for example, the federal agency may provide data to enable educational researchers to examine trends, such as the proportion of individuals engaging in specific activities like illegal substance use over the past year, or to analyze the customer's default rate in the banking and financial system. These data releases require careful consideration of their privacy and confidentiality before the government makes them accessible to the public; nonetheless, the disclosure risks have become increasingly intricate as the flourishing of machine learning and deep learning models empower the linkage of traditionally anonymized data back to individuals. The conventional anonymization (Abowd & Schmutte, 2015; Ohm, 2010) techniques often fail to prevent re-identification, as illustrated by high-profile cases like the Netflix privacy breach (Narayanan & Shmatikov, 2008). Consequently, many studies (Awan & Slavković, 2020; Dwork, 2006) recommend using differential privacy (DP) mechanisms, which introduce noise into the data before release to protect individual privacy.

Although this excitement around DP has already prompted companies like Apple (Differential Privacy Team, Apple, 2017), Uber (Near, 2018), Microsoft, Meta, Google, and many other tech leaders to form teams aiming to adopt DP in their data, and has spurred academia to delve into DP research with great enthusiasm, government agencies and other data curators remain limited (Drechsler, 2023). Even if they adopt these recommendations, most federal data products do not yet fully align with standard implementations of ε-DP and Rényi DP techniques, and substantial progress is still required for agencies to adopt even partial implementations of ε-DP mechanisms (Reiter, 2019). Although intermediate alternatives, such as synthetic data, provide some level of privacy protection (Barrientos et al., 2018; Raghunathan, 2021) and are grounded in rigorous mathematical frameworks, synthetic data can sometimes degrade data quality and may not be suitable for certain research purposes.

In this thesis, we construct the binomial proportion interval inference methodology and comparison under the ε-DP (Dwork, 2006) with the Laplace noise mechanism, and Rényi DP (Mironov, 2017) with the discrete Gaussian noise mechanism. Throughout this thesis, we assume practically that the ε-DP parameter, ε, which controls the level of privacy protection, is disclosed alongside the noise-added binomial proportion $\hat{p}^*$. The ε parameter indicates the data's accuracy or the blurring level, allowing the public to understand the trade-off between privacy and precision for the government-disclosed data.

The rest of the thesis is organized as follows: Section 2 reviews the classical Wald and Wilson intervas without DP, emphasizing the advantages of the Wilson interval over Wald, and the definitions and properties of ε-DP and Rényi DP we used. Section 3 highlights the lack of reasonable bounds when applying adjusted plug-in methods to the Wald and Wilson intervals under the DP mechanism. Section 4 includes the ε-DP Bayesian credible interval with uniform and Jeffreys prior, the ε-DP sampling-based interval, the ε-DP exact binomial interval, and the Rényi DP Bayesian credible interval for the discrete Gaussian mechanism. Section 5 discusses the simulation studies' findings and compares the interval methods' repeated sampling properties. Finally, Section 6 discusses the conclusion and topics that merit further investigation.

2. Background

We first discuss the Wald interval and Wilson interval. Then we discuss the mechanism of ε-DP with Laplace noise and Rényi DP with discrete Gaussian noise.

2.1 Wald and Wilson Interval with no DP

Under the binomial proportion, $X \sim \mathsf{Binomial}(n, p)$, $p \in (0, 1)$, when we observe $x_1 \ldots x_n \sim \mathsf{Bernoulli}(p)$. It is well known that, for large n, $\hat{p} \sim N\left(p, p(1-p)/n\right)$. Hence, considering $\hat{p}$ as a random variable, we have

$$\Pr\left[-z_{\alpha/2} < \frac{\hat{p} - p}{\sqrt{p(1-p)/n}} < z_{\alpha/2}\right] = 1 - \alpha. \tag{2.1}$$

Two well-known methods to construct interval estimates for the unknown p include

- Wald interval:

$$\hat{p} \pm z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}} \tag{2.2}$$

- Wilson interval:

$$\underbrace{\left(\frac{1}{1 + \frac{z_{\alpha/2}^2}{n}}\right)}_{①} \underbrace{\left(\hat{p} + \frac{z_{\alpha/2}^2}{2n}\right)}_{②} \pm z_{\alpha/2} \underbrace{\frac{\sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z_{\alpha/2}^2}{4n^2}}}{1 + \frac{z_{\alpha/2}^2}{n}}}_{③} \tag{2.3}$$

The overshooting problem, where the upper bound exceeds 1 or the lower bound falls below 0, is one of the main criticisms of the Wald interval. Many studies (Brown et al., 2001; Newcombe, 1998; Wallis, 2013) recommend using alternative methods to address this and other issues associated with the Wald interval. One such alternative is the Wilson interval, long regarded as an intuitive and straightforward remedy for the Wald interval in terms of addressing the out-of-bounds problem.

To see this, note that the binomial proportion is bounded, i.e., $0 \leqslant \hat{p} \leqslant 1$. We consider three cases: $\hat{p} = 1$, $\hat{p} = 0$, and $0 < \hat{p} < 1$. If $\hat{p} = 1$, the upper bound of (2.3) simplifies to $\frac{1 + \frac{z_{\alpha/2}^2}{2n}}{1 + \frac{z_{\alpha/2}^2}{n}} + \frac{\frac{z_{\alpha/2}^2}{2n}}{1 + \frac{z_{\alpha/2}^2}{n}} = 1$, and the lower bound simplifies to $\frac{1}{1 + \frac{z_{\alpha/2}^2}{n}} > 0$. Similarly, if $\hat{p} = 0$, the (2.3) simplifies to $\left[0, \frac{\frac{z^2_{\alpha/2}}{n}}{1+\frac{z^2_{\alpha/2}}{n}}\right]$. In the last case, $0 < \hat{p} < 1$, the upper bound satisfies: $(①②) + ③ < 1$, while the lower bound satisfies: $(①②) - ③ > 0$. Thus, all three cases satisfy that the interval remains within a reasonable bounded $[0,1]$.

However, in Section 3, we will demonstrate that the overshooting issue persists even with the Wilson interval once the DP mechanism is applied.

2.2 Laplace Mechanism

We will use the definition of ε-DP in Dwork (2006). First, define two databases as neighboring databases if differing only one observation in the dataset $x' \in \mathbb{N}^{|\mathcal{X}|}$ from the original dataset $x \in \mathbb{N}^{|\mathcal{X}|}$, we have $\|x - x'\|_1 \leqslant 1$. Second, we define Δf as the ℓ_1-sensitivity to quantify the maximum deviation of the function change with only one observation change from x to x':

$$\Delta f = \max_{\substack{x, x' \in \mathbb{N}^{|\mathcal{X}|} \\ \|x - x'\|_1 = 1}} \|f(x) - f(x')\|_1, \tag{2.4}$$

and we say the DP mechanism preserves ε level of privacy if

$$\frac{P[\mathcal{M}(x, f(\cdot), \varepsilon)]}{P\left[\mathcal{M}\left(x', f(\cdot), \varepsilon\right)\right]} = e^{\varepsilon}. \tag{2.5}$$

Finally, from Example 1 in Dwork et al. (2006), we know the Laplace mechanism satisfies DP, saying that given any function f: $\mathbb{N}^{|\mathcal{X}|} \to \mathbb{R}^k$, we have

$$\mathcal{M}_{\text{Laplace}}(x, f(\cdot), \varepsilon) = f(x) + (\eta_1, \ldots, \eta_n) \tag{2.6}$$

where η_i are i.i.d. random variables drawn from $\mathsf{Lap}(0, \Delta f/\varepsilon)$

Making the inference of binomial proportions under DP, we will utilize the arbitrary post-processing properties guaranteed by Dwork (2006), which is saying that if a mechanism $\mathcal{M}(x)$ satisfies ε-DP, then $\forall g : R \to R'$, any transformation applied to its output $g(\mathcal{M}(x)) := g \circ \mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \to R'$ also satisfies ε-DP. This property will allow the interval construction process to make the flexible transformation while also not losing the DP property.

2.3 Discrete Gaussian Mechanism

The motivation of the Rényi DP is from Rényi Divergence and Kullback-Leibler (KL) Divergence, where it takes advantage of the relative entropy distance to quantify the space between two density functions. Then, we utilize this property to examine Δf, which again is the ℓ_1-sensitivity to quantify the maximum deviation of the function change with only one observation change from x to x', where $x, x' \in \mathbb{N}^{|\mathcal{X}|}$.

2.3.1 Rényi Divergence

In the following, we follow Mironov (2017) for the definition of Rényi Divergence:

$$D_\alpha(F\|F') := \frac{1}{\alpha-1}\log \mathrm{E}_{x\sim F'}\left(\frac{F(x)}{F'(x)}\right)^\alpha, \tag{2.7}$$

where if we set $\alpha = 1$, then Rényi divergence is equivalent to the KL divergence

$$D_1(F\|F') = \mathrm{F}_{x\sim P}\log\frac{F(x)}{F'(x)}. \tag{2.8}$$

The derivation of the relationship is shown in many works of literature; for reference, this thesis mentions one of the derivations in Appendix.B. Finally, if we set $\alpha = \infty$,

$$D_\infty(F\|F') = \sup_{x\in\operatorname{supp} F'}\log\frac{F(x)}{F'(x)}. \tag{2.9}$$

2.3.2 Rényi Divergence under Gaussian Noise Mechanism

With a setup similar to that of ε-DP, but using the Rényi structure defined in Mironov (2017), we have

$$\mathcal{M}_{\text{Gaussian}}(x, f(\cdot), \varepsilon) = f(x) + (\eta_1, \ldots, \eta_n) \tag{2.10}$$

where η_i are i.i.d. random variables drawn from $\mathsf{N}(0, \sigma = \Delta f\sqrt{\frac{\alpha}{2\varepsilon(\alpha)}})$, and Δf is the ℓ_1-sensitivity to quantify the maximum deviation of the function change with only one observation change from x to x'. Then we say the Gaussian Mechanism preserves ε level of privacy of order α if

$$D_\alpha\left(\mathcal{M}(x, f(\cdot), \varepsilon) \| \mathcal{M}\left(x', f(\cdot), \varepsilon\right)\right) \leqslant \varepsilon. \tag{2.11}$$

Here, we will also utilize the post-processing property again for Rényi DP, and Corollary 3 in Mironov (2017) already proves that with the sensitivity 1, the Gaussian noise mechanism $\mathcal{M}_{\text{Gaussian}}$ satisfies $\left(\alpha, \alpha/\left(2\sigma^2\right)\right)$- Rényi DP. We utilize this result for incorporating the binomial count data to Rényi DP, as its maximum deviation for the adjacent data x and x' is 1, then $\Delta f = \max_{\substack{x,x' \in \mathbb{N}^{|\mathcal{X}|} \\ \|x - x'\|_1 = 1}} \|f(x) - f(x')\|_1$.

3. Failure of Plug-In Wald and Wilson Under DP

Under the binomial distribution, $\hat{p} = \frac{1}{n}\sum_{i=1}^{n} x_i$, which has sensitivity $\Delta(f) = \frac{1}{n}$. Hence, the Laplace mechanism adds η with scale parameter $\lambda = \frac{\Delta(f)}{\varepsilon} = \frac{1}{n\epsilon}$. Let $\eta \sim \mathsf{Lap}(0, \lambda = \frac{1}{n\varepsilon})$, we add Laplace noise to $\hat{p}$: $\hat{p}^* = \hat{p} + \eta$. Therefore, $\mathbb{E}[\hat{p}^*] = \mathbb{E}[\hat{p}] + \mathbb{E}[\eta]$. Since $\hat{p}$ is an unbiased estimator of p and the Laplace noise η has mean 0, $\mathbb{E}[\hat{p}^*] = p + 0 = p$, where $\hat{p}^*$ is an unbiased estimator of p. Also, $\text{Var}(\hat{p}) = \frac{p(1-p)}{n}$, and $\text{Var}(\eta) = 2\lambda^2 = 2\left(\frac{1}{n\epsilon}\right)^2 = \frac{2}{n^2\epsilon^2}$. So the total variance of $\hat{p}^* = \text{Var}(\hat{p}^*) = \frac{p(1-p)}{n} + \frac{2}{n^2\epsilon^2}$.

Thus, one possible version of a plug-in DP-Wald interval uses:

$$CI_{\text{Wald, DP}} = \hat{p}^* \pm z_{\alpha/2} \cdot \sqrt{\frac{\hat{p}^*(1-\hat{p}^*)}{n} + \frac{2}{n^2\varepsilon^2}} \tag{3.1}$$

The Wald interval without DP already shows the out-of-bound problem; for the plug-in method, we got the extra term $\frac{2}{n^2\varepsilon^2}$ within the square root. When the ε is small, given n is fixed, the adjusted DP-Wald interval bounds may fall outside $[0,1]$, particularly for proportions near these boundaries.

We can also find a version of an adjusted DP-Wilson interval by solving the equation,

$$\left(n + z_{\alpha/2}^2\right) p^2 - \left(2n\hat{p}^* + z_{\alpha/2}^2\right) p + \left(n\hat{p}^{*2} - \frac{2z_{\alpha/2}^2}{n\varepsilon^2}\right) \leqslant 0 \tag{3.2}$$

However, adding Laplace noise can lead to a DP-Wilson interval falling outside $[0,1]$ space, where the Wilson method is no longer the better alternative to Wald (proof in Appendix. A).

4. Alternative Intervals Under DP

In this section, we present alternative intervals under DP. In Section 4.1, we propose ε-DP Bayesian interval with a flat prior. In Section 4.2, we incorporate the Jeffreys prior into a similar Bayesian framework. In Section 4.3, we discuss an approximation to the fully Bayesian interval, which we call the sampling-based interval, that avoids MCMC estimation of the posterior distribution. In Section 4.4, we mention the early method, exact mechanism, and relatively straightforward to interpret the interval. We finally go through the Rényi framework in Section 4.5.

4.1 ε - DP Bayesian Interval with Flat Prior

We propose to construct the Bayesian interval for the true proportion p, where the true sample proportion $\hat{p}$ won't be released due to privacy concerns. Instead, $\hat{p}$ is perturbed by Laplace noise to create $\hat{p}^*$, the noisy, privatized version, which will be disseminated to the public. Here, our setup includes $\hat{p} \sim \mathsf{Binomial}(n, p), p \in (0, 1)$ and $\hat{p}^* = \hat{p} + \mathsf{Lap}(0, \lambda)$ where $\lambda = \frac{1}{n\varepsilon}$ is the scale parameter, and ε is the privacy-level parameter. Therefore, $\hat{p}^* \mid p \sim \mathsf{Lap}(\hat{p}, \lambda)$ and

$$f(\hat{p}^* \mid \hat{p}) = \frac{1}{2\lambda} \exp\left(-\frac{|\hat{p}^* - \hat{p}|}{\lambda}\right). \tag{4.1}$$

We choose the uniform distribution $p \sim \mathsf{Uniform}(0, 1)$ to construct a non-informative prior, reflecting the prior belief that any value of p between 0 and 1 is equally likely in the absence of prior knowledge about the true proportion or the observed data. Our objective is to estimate p based on the observed noisy proportion $\hat{p}^*$. The posterior distribution of p under Bayes' theorem is

$$f(p \mid \hat{p}^*) \propto f(\hat{p}^* \mid p) f(p), \tag{4.2}$$

and likelihood $f(\hat{p}^* \mid p)$ is computed by integrating over the unobserved $\hat{p}$

$$f(\hat{p}^* \mid p) = \int f(\hat{p}^* \mid \hat{p}) f(\hat{p} \mid p) \, d\hat{p}, \tag{4.3}$$

representing a convolution of the Laplace distribution (for $\hat{p}^* \mid \hat{p}$) and the binomial distribution (for $\hat{p} \mid p$). This integral does not have a closed-form solution. We, therefore, use the Gibbs sampler with a grid-based approach to sample from the posterior distribution in (4.2).

The likelihood for $\hat{p}$ comes from the binomial distribution: $f(\hat{p} \mid p) \propto p^k(1-p)^{n-k}$. The prior for p is uniform over the interval $(0,1)$: $f(p) = 1 \quad \text{for } p \in (0,1)$. To generate samples from the posterior distribution $f(p \mid \hat{p}^*)$, we use Gibbs sampling. The process alternates between sampling $\hat{p}$ and p from their full conditional distributions. Since we apply the grid-based method for sampling $\hat{p}$ due to the lack of a closed form, the integral in (4.3) becomes a finite sum. Then, the full posterior as a function of p is

$$f\left(p \mid \hat{p}^*\right) \propto \sum_{\hat{p}} \underbrace{\frac{1}{2\lambda} \exp\left(-\frac{|\hat{p}^* - \hat{p}|}{\lambda}\right)}_{\text{Laplace}} \underbrace{\binom{n}{k} p^k (1-p)^{n-k}}_{\text{Binomial likelihood for } \hat{p}=k/n} . \tag{4.4}$$

The full conditional for $\hat{p}$ combines both the Laplace mechanism and the binomial likelihood is

$$f(\hat{p} \mid p, \hat{p}^*) \propto \exp\left(-\frac{|\hat{p}^* - \hat{p}|}{\lambda}\right) \cdot p^{n\hat{p}}(1-p)^{n(1-\hat{p})}. \tag{4.5}$$

Instead of using a normal approximation, we first evaluate $\hat{p}$ over a grid of possible values between 0 and 1. This grid is spaced by $\frac{1}{n}$, so the possible values of $\hat{p}$ are $0, \frac{1}{n}, \frac{2}{n}, \ldots, 1$. Second, for each grid value $\hat{p}_i$, compute the convoluted likelihood $L(\hat{p}_i) = \exp\left(-\frac{|\hat{p}^* - \hat{p}_i|}{\lambda}\right) p^{n\hat{p}_i}(1-p)^{n(1-\hat{p}_i)}$. Third, we normalize likelihood values to form a probability distribution. Finally, we sample $\hat{p}$ from the grid based on these probabilities.

Given the updated value of $\hat{p}$, we know $f(p \mid \hat{p}) \propto p^k(1-p)^{n-k} \times 1$; hence, the full conditional distribution of p follows the Beta distribution

$$p \mid \hat{p} \sim \text{Beta}(\hat{p}n + 1, (1-\hat{p})n + 1) \tag{4.6}$$

We utilize the above two steps alternately in the Gibbs sampling process to generate posterior samples for p, which is summarised in Algorithm 1.

Algorithm 1 Two Steps Gibbs Sampling for Bayesian Interval under DP-ε

1. Sample $\hat{p} \mid p, \hat{p}^*$ Use the grid-based method to sample from the conditional distribution of $\hat{p}$, combining the Laplace mechanism and binomial likelihood.
2. Sample $p \mid \hat{p}$: Sample p from the Beta distribution: $p \sim \text{Beta}(\hat{p}n + 1, (1 - \hat{p})n + 1)$

After running Gibbs sampling for a sufficient number of iterations, we obtain a set of samples from the posterior distribution of p. We sort the posterior samples of p. The $(1 - \alpha) \times 100\%$ interval is given by the $\alpha/2$ percentile and the $1 - \alpha/2$ percentile of the posterior samples. For example, for a 95% credible interval: $\text{CI}_{95\%} = [p_{\alpha/2}, p_{1-\alpha/2}]$, where $p_{\alpha/2}$ is the 2.5 percentile and $p_{1-\alpha/2}$ is the 97.5 percentile of the posterior distribution, where this lower bound and upper bound finalize our DP Bayesian interval with flat prior.

4.2 ε - DP Bayesian Interval with Jeffreys Prior

A second uninformative prior chosen here is Jeffreys prior (Zanella-Béguelin et al., 2022). Under the structure of Jeffreys prior for the binomial likelihood:

$$p(p) \propto \sqrt{\frac{1}{p(1-p)}} = \frac{1}{\sqrt{p(1-p)}} \tag{4.7}$$

which corresponds to a Beta prior:

$$p \sim \text{Beta}(1/2, 1/2). \tag{4.8}$$

Our targeting a posteriori distribution of p is again $f(p \mid \hat{p}^*) \propto f(\hat{p}^* \mid p) f(p)$ and the likelihood function $f(\hat{p}^* \mid p)$ is computed by integrating over the unobserved $\hat{p}$: $f(\hat{p}^* \mid p) = \int f(\hat{p}^*, \hat{p} \mid p)\, d\hat{p} = \int f(\hat{p}^* \mid \hat{p})\, f(\hat{p} \mid p) d\hat{p}$. Here, the only difference with the flat prior is that the posterior distribution is based on Jeffreys prior.

$$f(p \mid \hat{p}^*) \propto f(\hat{p}^* \mid p)\, f(p) = \int f(\hat{p}^* \mid \hat{p})\, f(\hat{p} \mid p) f(p) d\hat{p} \tag{4.9}$$

where we have

$$f(p \mid \hat{p}) \propto f(\hat{p} \mid p) f(p). \tag{4.10}$$

The binomial likelihood is:

$$f(\hat{p} \mid p) \propto p^k (1-p)^{n-k}. \tag{4.11}$$

The Jeffreys prior is:

$$f(p) \propto p^{-1/2}(1-p)^{-1/2}. \tag{4.12}$$

Then, the multiplication of both terms:

$$\begin{aligned} f(p \mid \hat{p}) &\propto p^k(1-p)^{n-k} \cdot p^{-1/2}(1-p)^{-1/2} \\ &= p^{k-1/2}(1-p)^{n-k-1/2}. \end{aligned} \tag{4.13}$$

From the kernel, we know it follows a Beta distribution and $k = n\hat{p}$:

$$\begin{aligned} p \mid \hat{p} &\sim \mathsf{Beta}(k+1/2, n-k+1/2) \\ p \mid \hat{p} &\sim \mathsf{Beta}(n\hat{p}+1/2, n(1-\hat{p})+1/2) \end{aligned} \tag{4.14}$$

The full posterior as a function of p under a grid-based approach is:

$$f(p \mid \hat{p}^*) \propto \sum_{\hat{p}} \frac{1}{2\lambda} \exp\left(-\frac{|\hat{p}^* - \hat{p}|}{\lambda}\right) \binom{n}{k} p^{k-1/2}(1-p)^{n-k-1/2}. \tag{4.15}$$

The Gibbs sampler steps remain the same as in Algorithm 1, except that in sampling $p \mid \hat{p}$, we now use the Jeffreys prior-based posterior: $p \mid \hat{p} \sim \mathsf{Beta}(n\hat{p}+1/2, n(1-\hat{p})+1/2)$.

Finally, we obtain a set of samples from the posterior distribution of p. The Bayesian credible interval with Jeffreys prior is constructed using these samples where $p_{\alpha/2}$ is the 2.5 percentile and $p_{1-\alpha/2}$ is the 97.5 percentile of the posterior distribution.

4.3 ε - DP Sampling-Based Interval

This method could be considered an ad-hoc strategy but is straightforward. We know the public could not directly observe $\hat{p}$ under the DP mechanism. However, for any observed $\hat{p}^*$, if we can sample $\hat{p}^{(i)}|\hat{p}^*$ to reflect the randomness of $\hat{p}$, given its guarantee to be [0,1], then we can construct the inference of true p based on each $\hat{p}^{(i)}$.

The logic is as follows: $\hat{p}^* = \hat{p} + \eta$, and $\eta \sim \mathsf{Lap}\ (0, \lambda)$ where $\lambda = \frac{1}{n\varepsilon}$. We know $\forall X \in \{0, 1, \ldots, n\}$, $\hat{p} = \frac{X}{n}$ must be in the set $\left\{0, \frac{1}{n}, \ldots, \frac{n}{n}\right\}$. For each integer k from 0 to n, let $\hat{p}_k = \frac{k}{n}$, and the mechanism is $\hat{p}^*|\hat{p}_k \sim \mathsf{Lap}(\hat{p}_k, \lambda = \frac{1}{n\varepsilon})$; hence, if we just use the noise distribution to learn about $\hat{p}_k$, the density function is

$$p\left(\hat{p}^* \mid \hat{p}_k\right) = \frac{1}{2\lambda} \exp\left(-\frac{|\hat{p}^* - \hat{p}_k|}{\lambda}\right). \tag{4.16}$$

Thus, an unnormalized weight for each k is:

$$w(k) = p\left(\hat{p}^* \mid \hat{p}_k\right). \tag{4.17}$$

Summing up weights over $k = 0, \ldots, n$ to be the denominator of the normalizing process:

$$W = \sum_{k=0}^{n} w(k). \tag{4.18}$$

Then the discrete probability of $\hat{p} = \frac{k}{n}$ is

$$P\left(\hat{p} = \frac{k}{n} \middle| \hat{p}^*\right) = \frac{w(k)}{W}. \tag{4.19}$$

Hence we have a mass function over the $(n+1)$ discrete values $\frac{0}{n}, \frac{1}{n}, \ldots, 1$. Finally, to get N Monte Carlo samples $\hat{p}^{(1)}, \ldots, \hat{p}^{(N)}$ from $\hat{p} \mid \hat{p}^*$:

- Sample an integer $k_i \in \{0, \ldots, n\}$ with probabilities $P\left(k_i\right) = \frac{w(k_i)}{W}$.
- Set $\hat{p}^{(i)} = k_i/n$.

The intuition here for this technical step is that we want to know how likely we will observe $\hat{p}^*$, given the original sample proportion $\hat{p}$. Hence, once we observed $\hat{p}^*$, we assign the probability for all the possible $\hat{p}^{(i)}$ based on (4.16). If a possible $\hat{p}$ is close to $\hat{p}^*$, it is more likely to have generated $\hat{p}^*$, so it gets a higher probability.

After we observe enough sample size of $\hat{p}^{(i)}$, for each $\hat{p}^{(i)}$, we translate the sampling as if we had observed it directly from the binomial sampling, then mapping back to the binomial count data:

$$X_i = n \times \hat{p}^{(i)} \quad \text{(rounded to an integer)}, \quad X_i \sim \text{Binomial}(n, p).$$

Hence each draw $\hat{p}^{(i)}$ is treated like a "real sample proportion" from an underlying binomial. Now we bring in a standard binomial Wilson confidence interval. Then we

Algorithm 2 Constructing the Distribution of $\hat{p}*$, given p

1. Simulate Sample Proportion $\hat{p} = \frac{X}{n}$, $X \sim$ Binomial(n, p), where $\mathbb{E}[\hat{p}] = p$, andVar$(\hat{p}) = \frac{p(1-p)}{n}$
2. Adding Laplace noise $\eta \sim$ Laplace $(0, \lambda)$ with mean 0 and scale parameter $\lambda = \frac{1}{n\epsilon}$ to follow the DP. Hence, obtaining the noisy sample proportion $\hat{p}^* = \hat{p} + \eta$
3. Repeat the above two processes many times (e.g., 1,000 or more) to obtain a large sample of $\hat{p}^*$ values, which reflects both the sampling variability in $\hat{p}$ and the added Laplace noise.

will get n plausible Wilson intervals, $\left[L^{(i)}, U^{(i)}\right]$ as an interval for p. We finally merge all $\left[L^{(1)}, U^{(1)},\right]$, $\left[L^{(2)}, U^{(2)},\right]$, $\left[L^{(3)}, U^{(3)},\right]$, $\ldots$ $\left[L^{(n)}, U^{(n)},\right]$ by taking: the $\alpha/2$ quantile of all $\left\{L^{(i)}\right\}$ as the overall lower bound of p; the $(1-\alpha/2)$ quantile of all $\left\{U^{(i)}\right\}$ as the overall upper bound of p.

4.4 ε - DP Exact Interval

This method is motivated by the Clopper-Pearson interval. Without the DP noise, the observed sample proportion $\hat{p}$ is directly connected to the true proportion p through a simple binomial distribution. However, under the DP structure, we have to take into account the extra randomness.

We begin this method by considering a range of candidate values for p from adding noise to the sample proportion, covering the interval from 0 to 1 in small increments. For each candidate value p, we assess if it could realistically explain the observed noisy proportion $\hat{p}^*_{\text{obs}}$.

For each candidate p, we simulate noisy estimates $\hat{p}*$ using Algorithm 2. Also, we compare $\hat{p}^*_{\text{obs}}$ to the simulated $\hat{p}^*$ distribution. For the upper tail, we calculate the probability of observing a value at least as large as $\hat{p}^*_{\text{obs}}$ under the simulated distribution of $\hat{p}^*$ for each candidate p :

$$P\left(\hat{p}^*_{\text{sim}} \geqslant \hat{p}^*_{\text{obs}} \mid p\right) \tag{4.20}$$

where $\hat{p}^*_{\text{sim}}$ are the simulated values of $\hat{p}^*$ for the candidate p. For the lower tail, similarly, we calculate the probability of observing a value at least as small as $\hat{p}^*_{\text{obs}}$ under the simulated distribution of $\hat{p}^*$:

$$P\left(\hat{p}_{\text{sim}}^* \leqslant \hat{p}_{\text{obs}}^* \mid p\right). \tag{4.21}$$

Finally, we use the numerical method in R to find the bound p_L, and p_U that match the following conditions and form our final confidence interval $[p_L, p_U]$. The lower bound p_L is the smallest candidate p for which the upper tail probability is less than or equal to $\alpha/2$:

$$p_L = \min\left\{p : P\left(\hat{p}_{\text{sim}}^* \geqslant \hat{p}_{\text{obs}}^* \mid p\right) \leqslant \alpha/2\right\}. \tag{4.22}$$

The upper bound p_U is the largest candidate p for which the lower tail probability is less than or equal to $\alpha/2$:

$$p_U = \max\left\{p : P\left(\hat{p}_{\text{sim}}^* \leqslant \hat{p}_{\text{obs}}^* \mid p\right) \leqslant \alpha/2\right\}. \tag{4.23}$$

4.5 Rényi DP Bayesian Interval

Under a Rényi DP setup with a discrete Gaussian mechanism, we have $X \sim \mathsf{Binomial}(n, p)$, where $p \in (0,1)$. The government can observe $x_1 \ldots x_n$ and the proportion $\hat{p} = \frac{x}{n}$; then, they add noise to the counts x:

$$\hat{x}^* = x + G, \quad G \sim \mathsf{DiscreteGaussian}\left(0, \sigma^2\right), \text{where } \hat{p}^* = \frac{\hat{x}^*}{n}. \tag{4.24}$$

Given a sensitivity Δf ($\Delta f = 1$ for binomial counts) and an Rényi DP parameter α, the Gaussian mechanism achieves $(\alpha, \varepsilon(\alpha))$-Rényi DP if $\varepsilon(\alpha) = \frac{\alpha \Delta f^2}{2\sigma^2}$ as proved in Corollary 3 of (Mironov, 2017). Hence, $\sigma = \Delta\sqrt{\frac{\alpha}{2\varepsilon(\alpha)}}$. In many simplified examples, we can choose α so that $\varepsilon(\alpha)$ aligns with a desired privacy level. Here, to gain some idea of the connection between ε-DP and Rényi DP, we set a more straightforward approximate relationship $\sigma = \frac{1}{n\varepsilon}$, where it is analogous to the Laplace scale parameter $\lambda = 1/(n\varepsilon)$. However, the more rigorous relationship ties σ, Rényi DP parameters α, and $\varepsilon(\alpha)$ worth further studies.

G follows the discrete Gaussian distribution with mean zero and variance σ^2 has pmf:

$$f_G(g) = \frac{\exp\left(-\frac{g^2}{2\sigma^2}\right)}{\sum_{m=-\infty}^{\infty} \exp\left(-\frac{m^2}{2\sigma^2}\right)}. \tag{4.25}$$

The likelihood of $\hat{x}^*$:

$$f(\hat{x}^* \mid p) = \sum_{x=0}^{n} f(\hat{x}^* \mid x) f(x \mid p), \text{ where } f(x \mid p) = \binom{n}{x} p^x (1-p)^{n-x}. \tag{4.26}$$

Since G is discrete Gaussian with mean zero, we just replace G by $(\hat{x}^* - x)$ in the pmf of G because

$$f(\hat{x}^* \mid x) = P(\hat{x}^* = \hat{x}^* \mid x) = P(x + G = \hat{x}^*) = P(G = \hat{x}^* - x). \tag{4.27}$$

Hence,

$$f(\hat{x}^* \mid x) = f_G(\hat{x}^* - x) = \frac{\exp\left(-\frac{(x^* - x)^2}{2\sigma^2}\right)}{\sum_{m=-\infty}^{\infty} \exp\left(-\frac{m^2}{2\sigma^2}\right)}. \tag{4.28}$$

We know $\hat{p}^* = \hat{x}^*/n$; hence,

$$f(\hat{p}^* \mid p) = f(\hat{x}^* \mid p). \tag{4.29}$$

Since $\hat{x}^* = x + G$ and

$$\hat{p}^* = \hat{x}^*/n = \hat{p} + \underbrace{(G/n)}_{\text{scaled discrete Gaussian}}, \tag{4.30}$$

the distribution of $\hat{p}^* \mid \hat{p} \sim \mathsf{DiscreteGaussian}\left(\hat{p}, \frac{\sigma^2}{n^2}\right)$ and

$$f(\hat{p}^* \mid \hat{p}) \propto \exp\left(-\frac{(\hat{p}^* - \hat{p})^2}{2(\sigma^2/n^2)}\right) = \exp\left(-\frac{n^2(\hat{p}^* - \hat{p})^2}{2\sigma^2}\right). \tag{4.31}$$

As in Section 4.1, we assume the noninformative flat prior $p \sim \mathsf{Uniform}(0,1)$. Hence, $f(p) = 1$ for $p \in (0,1)$. Again, $\hat{p}$ (or x) is not directly observed by the public due to privacy-preserving mechanisms. We cannot directly simplify to a Beta distribution. With Bayes' theorem, the posterior is $f(p \mid \hat{p}^*) \propto f(\hat{p}^* \mid p) f(p)$. We plug in the likelihood and prior:

$$f(p \mid \hat{p}^*) \propto \sum_{x=0}^{n} \exp\left(-\frac{(\hat{x}^* - x)^2}{2\sigma^2}\right) \binom{n}{x} p^x (1-p)^{n-x}. \tag{4.32}$$

Algorithm 3 Two Steps Gibbs Sampling for Bayesian Interval under Rényi-DP

1. Sampling $\hat{p}$ given p and $\hat{p}^*$: $f(\hat{p} \mid p, \hat{p}^*) \propto f(\hat{p}^* \mid \hat{p}) f(\hat{p} \mid p)$
   - The discrete Gaussian likelihood:

$$f(\hat{p}^* \mid \hat{p}) \propto \exp\left(-\frac{n^2 (\hat{p}^* - \hat{p})^2}{2\sigma^2}\right). \tag{4.33}$$

   - The binomial prior component in terms of $\hat{p}$:

$$f(\hat{p} \mid p) \propto p^{n\hat{p}} (1-p)^{n(1-\hat{p})}. \tag{4.34}$$

   We evaluate this over a grid $\{0, 1/n, 2/n, \ldots, 1\}$, normalize, and sample from it.
2. Sampling p given $\hat{p}$: Once conditional on $\hat{p} = k/n$ is known:

$$p \mid \hat{p} \sim \text{Beta}(n\hat{p} + 1, n(1-\hat{p}) + 1). \tag{4.35}$$

Then, we introduce $\hat{p} = x/n$ as a latent variable and perform two steps of Gibbs sampling as we did in Sections 4.1 and 4.2 to generate posterior samples for $p \mid \hat{p}^*$. The procedure is summarized in Algorithm 3.

Finally, the rest are all very similar where the Bayesian credible interval is constructed using these posterior samples p; $p_{\alpha/2}$ is the 2.5 percentile and $p_{1-\alpha/2}$ is the 97.5 percentile of the posterior distribution.

5. Simulations

We conduct the simulation studies in Section 5, aiming to examine the repeated sampling performances of the DP plug-in methods introduced in Section 3 and the alternative method proposed in Section 4, focusing on coverage and interval properties.

5.1 Descriptions

To ensure all simulation results are comparable, we fix the following constants. Here, $n = 100$ represents the sample size for each iteration. To test for most possible outcomes, we use the true proportions $p \in \{0.1, 0.2, 0.5, 0.8\}$, paired with $\varepsilon \in \{0.1, 0.3, 0.5, 5\}$. The reason for choosing these specific values is that when the true p is closer to the boundaries of $[0, 1]$, interval inference tends to be less robust, leading to longer intervals, lower coverage rate, and a higher frequency of out-of-bound problems. Similar issues arise with privacy-preserving levels—when ε is close to 0.1 or lower, the data is blurred more aggressively to enhance privacy and confidentiality. We run all simulations 5000 times for all the intervals, 1000 posterior draws from MCMC iterations for the Bayesian method, and 1000 draws for the sampling interval and the exact interval to ensure they accurately represent real scenarios. Finally, we set the significance level at $\alpha = 0.05$, which is commonly used in statistical inference.

5.2 Results

All the results discussed here fall under a DP framework. In Table 5.1, the plug-in Wald interval severely overshoots, especially for $\varepsilon < 1$, which matches expectations. We also observe that the plug-in Wilson interval, while less prone to overshooting than the Wald interval, still suffers from out-of-bound issues and fails to remain within $[0, 1]$. In particular, when $\varepsilon = 0.1$, the out-of-bound rate exceeds 50% for $p = 0.1$ and $p = 0.8$. Consequently, it is not a reasonable alternative for inference under these settings.

From Table 5.2, the Bayesian methods with a uniform prior and a Jeffreys prior show similar coverage and average length. This is anticipated because they share nearly identical structures and use uninformative priors whose differences do not significantly influence the posterior. Meanwhile, the sampling-based interval and the exact interval have similar structures for evaluating the likelihood of observing $\hat{p}^*$ given $\hat{p}$. However, when $\varepsilon = 0.1, 0.3, 0.5$, the sampling-based method shows a higher-than-nominal coverage rate ($\alpha = 0.05$), as illustrated in Figure 5.1, where we can observe the purple line having a higher coverage than the others. This comes at the cost of a wider interval, which is arguably too conservative, as seen in Figure 5.2. The most pronounced differences appear when $\varepsilon = 0.1$ and $p = 0.1$, where we observe the coverage ordering Sample > Exact > Bayes-J ≈ Bayes-U > 95.0%, while the average interval length follows the order Exact > Sample > Bayes-J ≈ Bayes-U. Hence, when ε is very low—indicating a scenario with a very strict privacy level—it is worthwhile to adopt the Bayesian method, as it offers a good balance between coverage and interval length. Overall, under ε-DP, the Bayesian methods, the sampling-based interval, and the exact interval all yield intervals whose average length decreases as ε increases; this could be observed in Figure 5.2, where the upper-left ranges between 0.4 to 0.7; the upper-right range between 0.2 to 0.4; the lower-left range between 0.1 to 0.3; and the lower-right range between 0.1 to 0.2.

Finally, the Rényi DP results in Table 5.3 confirm that it remains stable under challenging conditions, where small ε or when p is near 0 or 1. Though Rényi DP is a more relaxed privacy framework and its stability is expected, the choice between Rényi DP and ε-DP depends on the specific context and degree of privacy required.

Table 5.1: Coverage (%), average length, and out-of-bound rate of plug-in Wald and plug-in Wilson confidence intervals for ε-DP with Laplace noise, evaluated across different values of ε and p.

Settings		**Coverage (%)**		**Average Length**		**Out-of-Bound Rate**	
ε	p	Wald	Wilson	Wald	Wilson	Wald	Wilson
0.1	0.1	95.6	95.1	.567	.557	.892	.879
0.1	0.2	94.6	94.4	.574	.564	.748	.725
0.1	0.5	93.7	94.0	.585	.574	.141	.122
0.1	0.8	97.0	96.9	.574	.564	.785	.752
0.3	0.1	93.8	97.1	.218	.216	.592	.441
0.3	0.2	94.1	94.6	.241	.237	.066	.039
0.3	0.5	94.3	94.7	.268	.263	.000	.000
0.3	0.8	94.1	94.6	.241	.237	.067	.038
0.5	0.1	93.6	94.7	.160	.160	.272	.128
0.5	0.2	94.1	94.9	.191	.189	.006	.002
0.5	0.5	94.5	94.9	.224	.220	.000	.000
0.5	0.8	94.1	94.9	.191	.189	.006	.002
5.0	0.1	93.1	94.7	.116	.118	.010	.000
5.0	0.2	93.7	94.7	.156	.155	.000	.000
5.0	0.5	94.4	94.7	.195	.192	.000	.000
5.0	0.8	93.8	94.8	.156	.155	.000	.000

Table 5.2: Coverage (%) and average length of Bayesian credible intervals with uniform (Bayes-U) and Jeffreys (Bayes-J) priors, a sampling-based method (Sample), and the exact interval (Exact) for ε-DP with Laplace noise, evaluated across different values of ε and p.

Settings		**Coverage (%)**				**Average Length**			
ε	p	Bayes-U	Bayes-J	Sample	Exact	Bayes-U	Bayes-J	Sample	Exact
0.1	0.1	96.0	96.0	98.4	97.3	.434	.438	.528	.555
0.1	0.2	97.2	97.2	98.5	97.9	.482	.487	.586	.548
0.1	0.5	95.4	96.4	98.3	95.5	.564	.571	.684	.605
0.1	0.8	97.5	96.6	98.7	97.4	.481	.488	.585	.552
0.3	0.1	97.0	97.7	99.3	97.7	.201	.203	.278	.244
0.3	0.2	95.3	95.8	99.2	95.6	.238	.240	.333	.243
0.3	0.5	95.3	94.7	99.2	95.0	.265	.268	.378	.268
0.3	0.8	95.0	95.5	99.2	94.9	.238	.239	.333	.244
0.5	0.1	95.8	95.7	99.3	96.0	.158	.156	.222	.167
0.5	0.2	95.2	95.9	99.0	94.9	.189	.189	.265	.190
0.5	0.5	95.0	95.3	99.3	94.4	.220	.222	.304	.222
0.5	0.8	94.4	93.7	99.2	95.2	.189	.189	.265	.191
5.0	0.1	94.2	95.7	94.9	95.0	.118	.116	.122	.117
5.0	0.2	94.2	94.1	95.1	95.0	.155	.155	.159	.155
5.0	0.5	94.6	94.8	94.7	95.0	.192	.193	.196	.194
5.0	0.8	94.9	95.6	94.4	94.4	.155	.154	.159	.155

Table 5.3: Coverage (%) and average length of Bayesian credible intervals for Rényi DP with discrete Gaussian noise, evaluated across different values of ε and p.

Settings		**Coverage (%)**	**Average Length**
ε	p	Rényi DP	Rényi DP
0.1	0.1	96.1	.165
0.1	0.2	95.2	.195
0.1	0.5	94.8	.226
0.1	0.8	95.2	.195
0.3	0.1	95.3	.136
0.3	0.2	95.0	.169
0.3	0.5	95.3	.204
0.3	0.8	94.9	.169
0.5	0.1	95.6	.129
0.5	0.2	95.2	.164
0.5	0.5	95.5	.199
0.5	0.8	94.7	.164
5.0	0.1	94.1	.118
5.0	0.2	95.0	.155
5.0	0.5	94.6	.192
5.0	0.8	94.1	.155

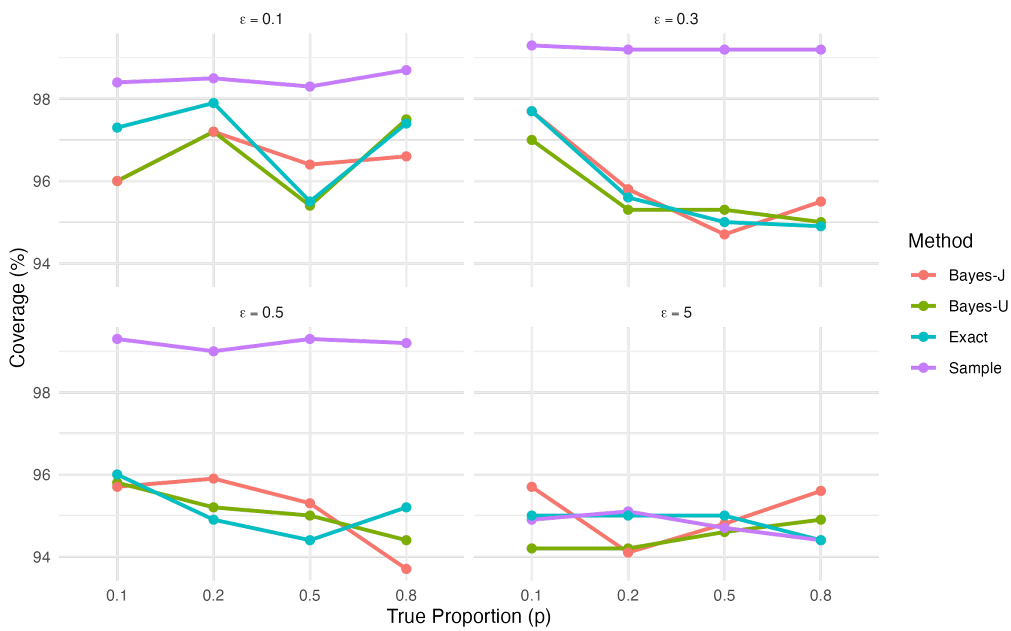

FIGURE 5.1: Coverage (%) of Bayesian credible intervals with uniform (Bayes-U) and Jeffreys (Bayes-J) priors, a sampling-based method (Sample), and the exact interval (Exact) for ε-DP with Laplace noise across different values of ε and p.

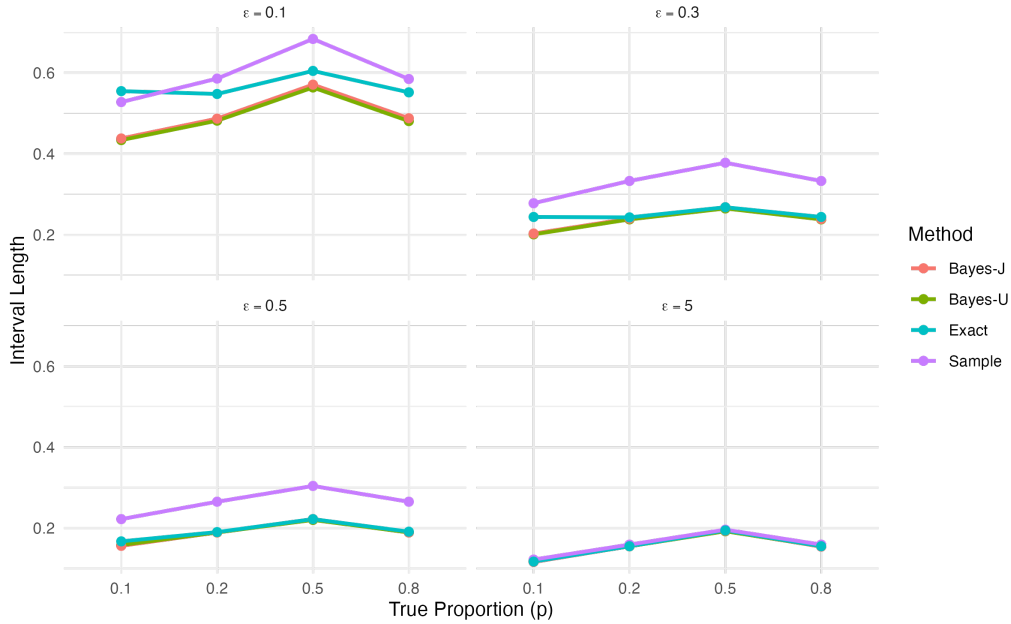

FIGURE 5.2: Average length of Bayesian credible intervals with uniform (Bayes-U) and Jeffreys (Bayes-J) priors, a sampling-based method (Sample), and the exact interval (Exact) for ε-DP with Laplace noise across different values of ε and p.

6. Conclusions

This thesis discusses the multiple interval construction methods for binomial proportions under the ε-DP and Rényi DP framework. From the simulation, we can conclude that under $\varepsilon = 0.1$, which is more interesting under the privacy context, the Bayesian interval dominates other methods with a greater balance between the narrower length and sufficient coverage. However, the sampling-based interval is obviously too conservative and should be ruled out if we get other methods to use. So to speak, we couldn't find the universal best method under ε-DP; there is a tradeoff between coverage and interval length over all the proposed methods. However, one point worth mentioning is that the Bayesian credible interval for Rényi DP with discrete Gaussian noise is the most stable, with good coverage, and the length is remarkably consistent even under the very small ε, though this is the more generalized differential privacy criteria.

For future studies, we have several feasible improvements. First, we need a better method for inferring the interval for the true p under a stringent privacy level (small ε), maintaining sufficient coverage—ideally reaching the nominal rate—without resulting in an overly conservative interval length. In other words, we can incorporate more advanced interval construction mechanisms, as many sophisticated methods have been developed, even if they are not yet widely adopted. Second, to enhance practicality, we also aim to extend the current privacy-preserving framework to other types of noise, such as geometric noise, and truncated or censored noise distributions. As the DP field continues to grow rapidly, we believe that inference methods should evolve alongside its development.

Appendix A. Wilson Confidence Interval Under Differential Privacy is Likely Out of Bounds

Under DP with Laplace noise, a plug-in version of a Wilson confidence interval for a proportion p can lead to interval bounds outside the $[0,1]$ range. This proof directly analyzes the quadratic function derived from the Wilson interval under DP to demonstrate why this out-of-bound issue occurs.

The inequality for the Wilson confidence interval under DP is:

$$(\hat{p}^* - p)^2 \leqslant z_{\alpha/2}^2 \left(\frac{p(1-p)}{n} + \frac{2}{n^2 \varepsilon^2} \right), \tag{A.1}$$

where $\hat{p}^*$ represents the observed noisy proportion with Laplace noise added in, while p denotes the true proportion. The sample size is given by n, and $z_{\alpha/2}$ is the critical value from the standard normal distribution. Last, ε is the privacy parameter that control the privacy level. Then we expand both sides:

- Left-Hand Side(LHS):

$$(\hat{p}^* - p)^2 = \hat{p}^{*2} - 2\hat{p}^* p + p^2. \tag{A.2}$$

- Right-Hand Side(RHS):

$$z_{\alpha/2}^2 \left(\frac{p(1-p)}{n} + \frac{2}{n^2 \varepsilon^2} \right) = \frac{z_{\alpha/2}^2}{n}(p - p^2) + \frac{2z_{\alpha/2}^2}{n^2 \varepsilon^2}. \tag{A.3}$$

Subtract RHS from LHS to bring all terms to one side:

$$\begin{gathered} p^2 + \frac{z_{\alpha/2}^2}{n} p^2 - \frac{z_{\alpha/2}^2}{n} p - 2\hat{p}^* p + \hat{p}^{*2} - \frac{2z_{\alpha/2}^2}{n^2 \varepsilon^2} \leqslant 0 \\ \left(1 + \frac{z_{\alpha/2}^2}{n}\right) p^2 - \left(2\hat{p}^* + \frac{z_{\alpha/2}^2}{n}\right) p + \left(\hat{p}^{*2} - \frac{2z_{\alpha/2}^2}{n^2 \varepsilon^2}\right) \leqslant 0. \end{gathered} \tag{A.4}$$

Noticeably, times all terms with n, will match (A.4) and (3.2)

$$\left(n + z_{\alpha/2}^2\right) p^2 - \left(2n\hat{p}^* + z_{\alpha/2}^2\right) p + \left(n\hat{p}^{*2} - \frac{2z_{\alpha/2}^2}{n\epsilon^2}\right) \leqslant 0. \tag{A.5}$$

To solve this, we separate the inequality into three parts: (1.) quadratic terms in p^2: $\left(n+z_{\alpha/2}^2\right)p^2$; (2.) linear terms in p: $-\left(2n\hat{p}^*+z_{\alpha/2}^2\right)p$; (3.) constant terms: $\left(n\hat{p}^{*2}-\frac{2z_{\alpha/2}^2}{n\epsilon^2}\right)$.

The inequality can now be written in the standard quadratic form:

$$Ap^2+Bp+C\leqslant 0, \tag{A.6}$$

where: A $=n+z_{\alpha/2}^2$; B $=-\left(2n\hat{p}^*+z_{\alpha/2}^2\right)$; and C $=n\hat{p}^{*2}-\frac{2z_{\alpha/2}^2}{n\varepsilon^2}$. We use the quadratic formula to analyze its properties:

$$p=\frac{-B\pm\sqrt{B^2-4AC}}{2A}, \tag{A.7}$$

where the discriminant is:

$$D=B^2-4AC. \tag{A.8}$$

Then, we assess the coefficients, especially for its sign:

- Coefficient A:

$$A=n+z_{\alpha/2}^2>0, \tag{A.9}$$

- Coefficient C:

$$C=n\hat{p}^{*2}-\frac{2z_{\alpha/2}^2}{n\varepsilon^2}. \tag{A.10}$$

When ε is small (strong privacy), the term $\frac{2z_{\alpha/2}^2}{n\varepsilon^2}$ becomes large, potentially making C negative. Compute D to analyze the roots:

$$\begin{aligned}
D&=B^2-4AC && \text{(A.11)}\\
&=\left(-\left(2n\hat{p}^*+z_{\alpha/2}^2\right)\right)^2-4(n+z_{\alpha/2}^2)\left(n\hat{p}^{*2}-\frac{2z_{\alpha/2}^2}{n\varepsilon^2}\right) && \text{(A.12)}\\
&=\left(2n\hat{p}^*+z_{\alpha/2}^2\right)^2-4(n+z_{\alpha/2}^2)n\hat{p}^{*2}+\frac{8z_{\alpha/2}^2(n+z_{\alpha/2}^2)}{n\varepsilon^2} && \text{(A.13)}\\
&=4n\hat{p}^*z_{\alpha/2}^2+z_{\alpha/2}^4-4nz_{\alpha/2}^2\hat{p}^{*2}+\frac{8z_{\alpha/2}^2\left(n+z_{\alpha/2}^2\right)}{n\varepsilon^2} && \text{(A.14)}\\
&=4nz_{\alpha/2}^2\hat{p}^*\left(1-\hat{p}^*\right)+z_{\alpha/2}^4+\frac{8z_{\alpha/2}^2\left(n+z_{\alpha/2}^2\right)}{n\varepsilon^2}. && \text{(A.15)}
\end{aligned}$$

The discriminant $D > 0$ ensures real roots exist.

Plug in everything, the roots of the quadratic equation are:

$$p = \frac{2n\hat{p}^* + z_{\alpha/2}^2 \pm \sqrt{D}}{2(n + z_{\alpha/2}^2)}. \tag{A.16}$$

Now, we demonstrate the out-of-bound issues for two extremes, $\hat{p}^* \to 0$ and $\hat{p}^* \to 1$:

- Case 1: $\hat{p}^* \to 0$ When $\hat{p}^*$ is close to 0:

$$C \approx -\frac{2z_{\alpha/2}^2}{n\varepsilon^2} < 0, \tag{A.17}$$

$$B \approx -z_{\alpha/2}^2 < 0, \tag{A.18}$$

$$D \approx z_{\alpha/2}^4 + \frac{8z_{\alpha/2}^2(n + z_{\alpha/2}^2)}{n\varepsilon^2} > 0. \tag{A.19}$$

The lower root:

$$p_{\text{lower}} = \frac{z_{\alpha/2}^2 - \sqrt{D}}{2(n + z_{\alpha/2}^2)} < 0 \tag{A.20}$$

since $\sqrt{D} > z_{\alpha/2}^2$.

- Case 2: $\hat{p}^* \to 1$ When $\hat{p}^*$ is close to 1:

$$C \approx n - \frac{2z_{\alpha/2}^2}{n\varepsilon^2}, \tag{A.21}$$

$$B \approx -(2n + z_{\alpha/2}^2) < 0, \tag{A.22}$$

$$D \approx z_{\alpha/2}^4 + \frac{8z_{\alpha/2}^2(n + z_{\alpha/2}^2)}{n\varepsilon^2} > 0. \tag{A.23}$$

The upper root:

$$p_{\text{upper}} = \frac{2n + z_{\alpha/2}^2 + \sqrt{D}}{2(n + z_{\alpha/2}^2)} > 1 \tag{A.24}$$

since the numerator exceeds the denominator.

To sum up, the analysis shows that when $C < 0$, which is likely under strong ε-DP privacy restriction (small ε), the roots of the quadratic equation can be outside the $[0, 1]$

interval. Specifically, when $\hat{p}^*$ is close to 0, the lower bound of the confidence interval is negative; when $\hat{p}^*$ is close to 1, the upper bound exceeds 1. Thus, the plug-in Wilson interval under ε-DP with Laplace noise mechanism is possibly and likely to be out of bounds.

Appendix B. Rényi Divergence and the Relationship to KL Divergence

First, the Rényi divergence is defined as in Mironov (2017):

$$D_\alpha(F|F') := \frac{1}{\alpha-1}\log \mathbb{E}_{x\sim F'}\left(\frac{F(x)}{F'(x)}\right)^\alpha \tag{B.1}$$

where α controls the level of emphasis on different regions of the density function, and F and F' are the two distributions under comparison.

We are interested in

$$\lim_{\alpha\to 1} D_\alpha(F\|F') = \lim_{\alpha\to 1}\frac{\log \mathbb{E}_{x\sim F'}\left(\frac{F(x)}{F'(x)}\right)^\alpha}{\alpha-1} \tag{B.2}$$

We set $r(\alpha) = \frac{F(x)}{F'(x)}$. Hence,

$$\lim_{\alpha\to 1} D_\alpha(F\|F') = \lim_{\alpha\to 1}\frac{\log \mathbb{E}_{x\sim F'}\left(r(x)\right)^\alpha}{\alpha-1}. \tag{B.3}$$

Now, we use L'Hôpital's rule, where the derivative of the denominator is 1; for the numerator, we set $g(\alpha) = \mathbb{E}_{x\sim F'}\left(r(x)\right)^\alpha$, then

$$\frac{\partial}{\partial\alpha}\log\left[g(\alpha)\right] = \frac{g'(\alpha)}{g(\alpha)} = \frac{\frac{\partial}{\partial\alpha}\mathbb{E}_{x\sim F'}\left(r(x)\right)^\alpha}{\mathbb{E}_{x\sim F'}\left(r(x)\right)^\alpha} = \frac{\mathbb{E}_{x\sim F'}\left[\frac{\partial}{\partial\alpha}\left(r(x)\right)^\alpha\right]}{\mathbb{E}_{x\sim F'}\left(r(x)\right)^\alpha}. \tag{B.4}$$

We know

$$\mathbb{E}_{x\sim F'}\left[\frac{\partial}{\partial\alpha}r(x)^\alpha\right] = \mathbb{E}_{x\sim F'}\left[\frac{\partial}{\partial\alpha}e^{\log r(x)^\alpha}\right] = \mathbb{E}_{x\sim F'}\left[e^{\log r(x)^\alpha}\log r(x)\right] = \mathbb{E}_{x\sim F'}\left[r(x)^\alpha \log r(x)\right]. \tag{B.5}$$

Hence, (B.3) equals to

$$\frac{\mathbb{E}_{x\sim F'}\left[r(x)^1\log r(x)\right]}{\mathbb{E}_{x\sim F'}\left[r(x)^1\right]}. \tag{B.6}$$

Finally, we rearrange and simplify the formula,

$$\mathbb{E}_{x\sim F'}\left[r(x)^1\log r(x)\right] = \int F'(x)\frac{F(x)}{F'(x)}\log\frac{F(x)}{F'(x)}dx = \int F(x)\log\frac{F(x)}{F'(x)}dx, \tag{B.7}$$

and

$$\mathbb{E}_{x\sim F'}\left[r(x)\right] = \int F'(x)\frac{F(x)}{F'(x)}dx = \int F(x)dx = 1. \tag{B.8}$$

Thus, (B.6) equals

$$\frac{\int F(x)\log\frac{F(x)}{F'(x)}dx}{1} = \mathbb{E}_{x\sim F}\log\frac{F(x)}{F'(x)}, \tag{B.9}$$

and we recover the KL divergence:

$$D_1(F\|F') = \mathbb{E}_{x\sim F}\log\frac{F(x)}{F'(x)}. \tag{B.10}$$

Bibliography

Abowd, J., & Schmutte, I. (2015). Economic analysis and statistical disclosure limitation. *Brookings Papers on Economic Activity*, *2015*, 221–293.

Awan, J., & Slavković, A. (2020). Differentially private inference for binomial data. *Journal of Privacy and Confidentiality*, *10*(1), 725.

Barrientos, A. F., Bolton, A., Balmat, T., Reiter, J. P., de Figueiredo, J. M., Machanavajjhala, A., Chen, Y., Kneifel, C., & DeLong, M. (2018). Providing access to confidential research data through synthesis and verification: An application to data on employees of the U.S. federal government. *The Annals of Applied Statistics*, *12*(2), 1124–1156.

Brown, L. D., Cai, T. T., & DasGupta, A. (2001). Interval estimation for a binomial proportion. *Statistical Science*, *16*(2), 101–133.

Differential Privacy Team, Apple. (2017). Learning with privacy at scale.

Drechsler, J. (2023). Differential privacy for government agencies—are we there yet? *Journal of the American Statistical Association*, *118*(541), 761–773.

Dwork, C. (2006). Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone, & I. Wegener (Eds.), *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, july 10–14, 2006. Proceedings, Part ii* (pp. 1–12). Springer.

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings*, 265–284.

Lott, A., & Reiter, J. P. (2020). Wilson confidence intervals for binomial proportions with multiple imputation for missing data. *The American Statistician*, *74*(2), 109–115.

Mironov, I. (2017). Rényi differential privacy. *CoRR*, *abs/1702.07476*.

Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. *2008 IEEE Symposium on Security and Privacy (sp 2008)*, 111–125.

Near, J. (2018). Differential privacy at scale: Uber and berkeley collaboration. *Enigma 2018 (Enigma 2018)*.

Newcombe, R. G. (1998). Two-sided confidence intervals for the single proportion: Comparison of seven methods. *Statistics in Medicine*, *17*(8), 857–872.

Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, *57*, 1701.

Raghunathan, T. E. (2021). Synthetic data. *Annual Review of Statistics and Its Application*, *8*, 129–150.

Reiter, J. P. (2019). Differential privacy and federal data releases. *Annual Review of Statistics and Its Application*, *6*, 85–101.

Wallis, S. (2013). Binomial confidence intervals and contingency tests: Mathematical fundamentals and the evaluation of alternative methods. *Journal of Quantitative Linguistics*, *20*(3), 178–208.

Wilson, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, *22*(158), 209–212.

Zanella-Béguelin, S., Wutschitz, L., Tople, S., Salem, A., Rühle, V., Paverd, A., Naseri, M., Köpf, B., & Jones, D. (2022). Bayesian estimation of differential privacy.