A Additional Background

A.1 Total Variation Distance Properties

Recall the definition of total variation distance:

Definition 2.1. *The total variation distance between random variables (or distributions) P_1 and P_2 is*

$$\mathrm{TV}(P_1, P_2) = \sup_A |\Pr(P_1 \in A) - \Pr(P_2 \in A)|, \tag{2}$$

where A is any measurable set.

Lemma A.1 (Kelbert (2023))**.** *Properties of total variation distance:*

1. *For probability densities p_1 and p_2,*
$$\mathrm{TV}(p_1, p_2) = \frac{1}{2} \int |p_1(x) - p_2(x)| \, \mathrm{d}x. \tag{18}$$
2. *Total variation distance is a metric.*
3. *Pinsker's inequality: for distributions P_1 and P_2,*
$$\mathrm{TV}(P_1, P_2) \leq \sqrt{\frac{1}{2} \mathrm{KL}(P_1 \,||\, P_2)} \tag{19}$$
4. *Invariance to bijections: if f is a bijection and P_1 and P_2 are random variables,*
$$\mathrm{TV}(f(P_1), f(P_2)) = \mathrm{TV}(P_1, P_2) \tag{20}$$

We also occasionally write $\mathrm{TV}(p_1, p_2)$ for probability densities p_1 and p_2 as

$$\mathrm{TV}(p_1, p_2) = \sup_h \left| \int h(x) p_1(x) \, \mathrm{d}x - \int h(x) p_2(x) \, \mathrm{d}x \right| \tag{21}$$

where h is an indicator function of some measurable set A.

A.2 Bernstein–von Mises Theorem Regularity Conditions

The version of the Bernstein–von Mises theorem we use is from van der Vaart (1998). To state the regularity conditions, we need two definitions:

Definition A.2. *A parametric probability density p_Q is differentiable in quadratic mean at Q_0 if there exists a measurable vector-valued function $\dot{\ell}_{Q_0}$ such that, as $Q \to Q_0$,*

$$\int \left(\sqrt{p_Q(x)} - \sqrt{p_{Q_0}(x)} - \frac{1}{2}(Q - Q_0)^T \dot{\ell}_{Q_0}(x) \sqrt{p_{Q_0}(x)} \right)^2 \mathrm{d}x = o(||Q - Q_0||_2^2). \tag{22}$$

Definition A.3. *A randomised test is a function $\phi \colon \mathcal{X} \to [0, 1]$.*

The interepretation of $\phi(X)$ is the probability of rejecting some null hypothesis after observing data X.

Now we can state the regularity conditions of Theorem 2.2:

Condition A.4 (van der Vaart (1998))**.** *For true parameter value Q_0 and observed data X_n:*

1. *The datapoints of X_n are i.i.d.*
2. *The likelihood $p(x|Q)$ for a single datapoint x is differentiable in quadratic mean at Q_0.*
3. *The Fisher information matrix of $p(x|Q)$ is nonsingular at Q_0.*
4. *For every $\beta > 0$, there exists a sequence of randomised tests ϕ_n such that*
$$p(X_n|Q_0)\phi_n(X_n) \to 0, \qquad \sup_{||Q - Q_0||_2 \geq \beta} p(X_n|Q)(1 - \phi_n(X_n)) \to 0. \tag{23}$$
5. *The prior $p(Q)$ is absolutely continuous (as a measure) in a neighbourhood of Q_0 with a continuous positive density at Q_0.*

A.3 Differential Privacy and Noise-Aware Synthetic Data

Differential privacy (DP) (Dwork et al. 2006b) quantifies the privacy loss from releasing the results of analysing data. The quantification is done by looking at the output distributions of the analysis algorithm for two datasets that differ in a single data subject (Dwork and Roth 2014):

Definition A.5. *An algorithm $\mathcal{M}$ is (ϵ, δ)-DP if*

$$\Pr(\mathcal{M}(X) \in S) \leq e^{\epsilon} \Pr(\mathcal{M}(X') \in S) + \delta \tag{24}$$

for all measurable sets S and all datasets X, X' that differ in one data subject.

The choice of ϵ and δ is a matter of policy (Dwork 2008). One should set $\delta \ll \frac{1}{n}$ for n datapoints, as $\delta \approx \frac{1}{n}$ permits mechanisms that clearly violate privacy (Dwork and Roth 2014).

A common primitive for making an algorithm DP is the *Gaussian mechanism* (Dwork et al. 2006a), which simply adds Gaussian noise to the output of a function:

Definition A.6. *The Gaussian mechanism with noise variance σ_{DP}^2 and function f outputs $f(X) + \mathcal{N}(0, \sigma_{DP}^2 I)$ for input X.*

For a given (ϵ, δ)-bound and function f, the required value for σ_{DP}^2 can be computed tightly using the analytical Gaussian mechanism (Balle and Wang 2018).

Noise-Aware Private Synthetic Data To solve the uncertainty estimation problem for frequentist analyses from DP synthetic data, Räisä et al. (2023) developed a noise-aware algorithm for generating synthetic data called NAPSU-MQ. NAPSU-MQ takes discrete data, summarises it with marginal queries, releases the query values under DP with the Gaussian mechanism, and finally generates multiple synthetic datasets. The downstream analysis is done on each synthetic dataset, and the results are combined using Rubin's rules for synthetic data (Raghunathan et al. 2003; Rubin 1993), which use the multiple analysis results to account for the extra uncertainty coming from the synthetic data generation.

The synthetic data is generated by sampling the posterior predictive distribution

$$p(X^*|\tilde{s}) = \int p(X^*|\theta) p(\theta|\tilde{s}) \,\mathrm{d}\theta. \tag{25}$$

The conditioning on $\tilde{s}$ and including the Gaussian mechanism in the model is what makes NAPSU-MQ noise-aware, and allows Rubin's rules to accurately account for the synthetic data generation and DP noise in the downstream analysis.

A.4 Bayesian Inference with Gaussian Models

In this section, we collect well-known results on Bayesian inference of a Gaussian mean. See Gelman et al. (2014) for proofs.

Scaled inverse-chi-square distribution This parameterisation of the inverse gamma distribution is convenient in this setting.

$$\text{Inv-}\chi^2(\nu, s^2) = \text{Inv-Gamma}\left(\alpha = \frac{\nu}{2}, \beta = \frac{\nu}{2} s^2\right). \tag{26}$$

If $\theta \sim \text{Inv-}\chi^2(\nu, s^2)$, $\theta > 0$,

$$p(\theta) = \frac{(\frac{\nu}{2})^{\frac{\nu}{2}}}{\Gamma(\frac{\nu}{2})} s^{\nu} \theta^{-(\frac{\nu}{2}+1)} e^{-\frac{\nu s^2}{2\theta}} \tag{27}$$

$$\mathbb{E}(\theta) = \frac{\nu}{\nu - 2} s^2, \quad \nu > 2 \tag{28}$$

$$\mathrm{Var}(\theta) = \frac{2\nu^2}{(\nu - 2)^2 (\nu - 4)} s^4, \quad \nu > 4. \tag{29}$$

Gaussian Model with Known Variance When the variance of the data is known to be σ_k^2, and only the mean is unknown, the conjugate prior is another Gaussian, and we get the following inference problem:

$$\mu \sim \mathcal{N}(\mu_0, \sigma_0^2) \tag{30}$$
$$x_i|\mu \sim \mathcal{N}(\mu, \sigma_k^2). \tag{31}$$

The posterior with n datapoints with sample mean $\bar{X}$ is:

$$\mu|X \sim \mathcal{N}(\mu_n, \sigma_n^2) \tag{32}$$
$$\mu_n = \frac{\frac{1}{\sigma_0^2}\mu_0 + \frac{n}{\sigma_k^2}\bar{X}}{\frac{1}{\sigma_0^2} + \frac{n}{\sigma_k^2}} \tag{33}$$
$$\frac{1}{\sigma_n^2} = \frac{1}{\sigma_0^2} + \frac{n}{\sigma_k^2}. \tag{34}$$

Gaussian Model with Unknown Variance When the variance of the data is also unknown, the conjugate prior is a inverse-chi-squared for the variance, and Gaussian for the mean, which gives the following inference problem:

$$\sigma^2 \sim \text{Inv-}\chi^2(\nu_0, \sigma_0^2) \tag{35}$$
$$\mu|\sigma^2 \sim \mathcal{N}\left(\mu_0, \frac{\sigma^2}{\kappa_0}\right) \tag{36}$$
$$x_i|\mu, \sigma^2 \sim \mathcal{N}(\mu, \sigma^2). \tag{37}$$

The joint posterior of μ and σ^2 for n datapoints is:

$$\sigma^2|X \sim \text{Inv-}\chi^2(\nu_n, \sigma_n^2) \tag{38}$$
$$\mu|\sigma^2, X \sim \mathcal{N}\left(\mu_n, \frac{\sigma^2}{\kappa_n}\right) \tag{39}$$
$$\tag{40}$$

with

$$\bar{X} = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{41}$$
$$s^2 = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{X})^2 \tag{42}$$
$$\mu_n = \frac{\kappa_0}{\kappa_0 + n}\mu_0 + \frac{n}{\kappa_0 + n}\bar{X} \tag{43}$$
$$\kappa_n = \kappa_0 + n \tag{44}$$
$$\nu_n = \nu_0 + n \tag{45}$$
$$\nu_n\sigma_n^2 = \nu_0\sigma_0^2 + (n-1)s^2 + \frac{\kappa_0 n}{\kappa_0 + n}(\bar{X} - \mu_0)^2. \tag{46}$$

The marginal posterior of μ is

$$\mu|X \sim t_{\nu_n}\left(\mu_n, \frac{\sigma_n^2}{\kappa_n}\right). \tag{47}$$

B Missing Proofs

B.1 Consistency Proof

For ease of reference, we repeat Theorem 2.2 and Condition 3.2:

Theorem 2.2 (Bernstein–von Mises (van der Vaart 1998)). *Let n denote the size of the dataset X_n. Under regularity conditions stated in Condition A.4 in Supplemental Section A.2, for true parameter value Q_0, the posterior $\bar{Q}(X_n) \sim p(Q|X_n)$ satisfies*

$$\text{TV}\left(\sqrt{n}(\bar{Q}(X_n) - Q_0), \mathcal{N}(\mu(X_n), \Sigma)\right) \xrightarrow{P} 0 \tag{3}$$

as $n \to \infty$ for some $\mu(X_n)$ and Σ, that do not depend on the prior, where the convergence in probability is over sampling $X_n \sim p(X_n|Q_0)$.

Recall that $\bar{Q}_n^+ \sim p(Q|Z, X_n^*)$, and $\bar{Q}_n \sim p(Q|X_n^*)$.

Condition 3.2. *For all Q there exist distributions D_n such that*

$$\text{TV}\left(\bar{Q}_n^+, D_n\right) \xrightarrow{P} 0 \quad \text{and} \quad \text{TV}\left(\bar{Q}_n, D_n\right) \xrightarrow{P} 0 \tag{10}$$

as $n \to \infty$, where the convergence in probability is over sampling $X_n^* \sim p(X_n^*|Z, Q)$.

Lemma 3.3. *If the assumptions of Theorem 2.2 (Condition A.4) and the following assumptions:*

(1) Z and X^* are conditionally independent given Q; and

(2) $p(Z|Q) > 0$ for all Q,

hold for the downstream analysis for all Q_0, then Condition 3.2 holds.

Proof. Under Assumption (1)

$$p(Q|Z, X_n^*) \propto p(X_n^*|Q)p(Z|Q)p(Q) \tag{48}$$

so we can view both $p(Q|Z, X_n^*)$ and $p(Q|X_n^*)$ as the posteriors for the same Bayesian inference problem with observed data X_n^*, and priors $p(Q|Z) \propto p(Z|Q)p(Q)$ and $p(Q)$, respectively. Due to Condition A.4 (5) and Assumption (2), $p(Q|Z)$ has an everywhere positive density. Recall that $\bar{Q}_n^+ \sim p(Q|Z, X_n^*)$ and $\bar{Q}_n \sim p(Q|X_n^*)$. Now, Theorem 2.2 gives

$$\text{TV}\left(\sqrt{n}(\bar{Q}_n^+ - Q_0), \mathcal{N}(\mu_n, \Sigma)\right) \xrightarrow{P} 0 \tag{49}$$

and

$$\text{TV}\left(\sqrt{n}(\bar{Q}_n - Q_0), \mathcal{N}(\mu_n, \Sigma)\right) \xrightarrow{P} 0 \tag{50}$$

as $n \to \infty$, where μ_n, Σ are equal in the two cases because they do not depend on the prior. The probability is over $X_n^* \sim p(X_n^*|Q_0)$. Because of Assumption (1), $p(X_n^*|Q_0) = p(X_n^*|Z, Q_0)$, so the convergence also holds with probability over $X_n^* \sim p(X_n^*|Z, Q_0)$. These hold for any Q_0. Because the function $f_n(q) = \sqrt{n}(q - Q_0)$ is a bijection and total variation distance is invariant to bijections, Condition 3.2 holds with D_n being the pushforward distribution $D_n = f_n^{-1} \circ \mathcal{N}(\mu_n, \Sigma)$, with the Q of Condition 3.2 being Q_0. Note that D_n is allowed to depend on Q in Condition 3.2 due to the order of quantifiers. □

Lemma B.1. *Under Condition 3.2,*

$$\text{TV}(\bar{Q}_n^+, \bar{Q}_n) \xrightarrow{P} 0 \tag{51}$$

as $n \to \infty$, with the probability over $X_n^* \sim p(X_n^*|Z)$.

Proof. Total variation distance is a metric, so

$$\text{TV}\left(\bar{Q}_n^+, \bar{Q}_n\right) \leq \text{TV}\left(\bar{Q}_n^+, D_n\right) + \text{TV}\left(\bar{Q}_n, D_n\right) \tag{52}$$

so by Condition 3.2

$$\text{TV}\left(\bar{Q}_n^+, \bar{Q}_n\right) \xrightarrow{P} 0 \tag{53}$$

as $n \to \infty$, with the probability over $X_n^* \sim p(X_n^*|Z, Q)$.

It remains to show (53) with the probability over $X_n^* \sim p(X_n^*|Z)$ instead of $X_n^* \sim p(X_n^*|Z, Q)$. With $X_n^* \sim p(X_n^*|Z)$, for any $\epsilon > 0$,

$$\Pr_{X_n^*|Z}(\text{TV}(\bar{Q}_n^+, \bar{Q}_n) > \epsilon) = \int \Pr_{X_n^*|Z,Q}(\text{TV}(\bar{Q}_n^+, \bar{Q}_n) > \epsilon) p(Q|Z) \, \mathrm{d}Q \tag{54}$$

(53) holds for any Q, so

$$\lim_{n\to\infty} \Pr_{X_n^*|Z,Q}(\mathrm{TV}(\bar{Q}_n^+, \bar{Q}_n) > \epsilon) = 0 \tag{55}$$

The dominated convergence theorem then implies that

$$\lim_{n\to\infty} \Pr_{X_n^*|Z}(\mathrm{TV}(\bar{Q}_n^+, \bar{Q}_n) > \epsilon) = 0 \tag{56}$$

so

$$\mathrm{TV}(\bar{Q}_n^+, \bar{Q}_n) \xrightarrow{P} 0 \tag{57}$$

as $n \to \infty$, with the probability over $X_n^* \sim p(X_n^*|Z)$. $\square$

Lemma B.2. *Let $y_n \sim U_n$ be an arbitrary sequence of continuous random variables and let $S(y_n)$, $T(y_n)$ be continuous random variables that depend on y_n. Let the density functions of $S(y_n)$, $T(y_n)$ and U_n be $f_{S(y_n)}, f_{T(y_n)}$ and f_{U_n}, respectively. If*

$$\mathrm{TV}(S(y_n), T(y_n)) \xrightarrow{P} 0 \tag{58}$$

as $n \to \infty$, where the probability is over $y_n \sim U_n$, then

$$\mathrm{TV}\left(\int f_{S(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n, \int f_{T(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n\right) \to 0 \tag{59}$$

as $n \to \infty$.

Proof. Let h be an indicator function of x over any measurable set and let $\epsilon > 0$. Then

$$\left|\int h(x) \int f_{S(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n\,\mathrm{d}x - \int h(x) \int f_{T(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n\,\mathrm{d}x\right| \tag{60}$$

$$= \left|\int h(x) \int f_{U_n}(y_n)\big(f_{S(y_n)}(x) - f_{T(y_n)}(x)\big)\,\mathrm{d}y_n\,\mathrm{d}x\right| \tag{61}$$

$$= \left|\int f_{U_n}(y_n) \int h(x)\big(f_{S(y_n)}(x) - f_{T(y_n)}(x)\big)\,\mathrm{d}x\,\mathrm{d}y_n\right| \tag{62}$$

$$\leq \int f_{U_n}(y_n) \left|\int h(x)\big(f_{S(y_n)}(x) - f_{T(y_n)}(x)\big)\,\mathrm{d}x\right|\,\mathrm{d}y_n \tag{63}$$

$$= \int f_{U_n}(y_n) \left|\int h(x) f_{S(y_n)}(x)\,\mathrm{d}x - \int h(x) f_{T(y_n)}(x)\,\mathrm{d}x\right|\,\mathrm{d}y_n \tag{64}$$

Because $\mathrm{TV}(S(y_n), T(y_n)) \xrightarrow{P} 0$, for large enough n, there is a set Y_n with $\mathrm{TV}(S(y_n), T(y_n)) < \frac{\epsilon}{2}$ for all $y_n \in Y_n$, and $\Pr(y_n \in Y_n^C) < \frac{\epsilon}{2}$. As

$$\mathrm{TV}(S(y_n), T(y_n)) = \sup_h \left|\int h(x) f_{S(y_n)}(x)\,\mathrm{d}x - \int h(x) f_{T(y_n)}(x)\,\mathrm{d}x\right| \leq 1 \tag{65}$$

now

$$\int f_{U_n}(y_n) \left|\int h(x) f_{S(y_n)}(x)\,\mathrm{d}x - \int h(x) f_{T(y_n)}(x)\,\mathrm{d}x\right|\,\mathrm{d}y_n \tag{66}$$

$$\begin{aligned} &= \int_{Y_n} f_{U_n}(y_n) \left|\int h(x) f_{S(y_n)}(x)\,\mathrm{d}x - \int h(x) f_{T(y_n)}(x)\,\mathrm{d}x\right|\,\mathrm{d}y_n \\ &+ \int_{Y_n^C} f_{U_n}(y_n) \left|\int h(x) f_{S(y_n)}(x)\,\mathrm{d}x - \int h(x) f_{T(y_n)}(x)\,\mathrm{d}x\right|\,\mathrm{d}y_n \end{aligned} \tag{67}$$

$$< \int_{Y_n} f_{U_n}(y_n) \frac{\epsilon}{2}\,\mathrm{d}y_n + \int_{Y_n^C} f_{U_n}(y_n)\,\mathrm{d}y_n \tag{68}$$

$$< \frac{\epsilon}{2} + \frac{\epsilon}{2} \tag{69}$$

$$= \epsilon \tag{70}$$

for large enough n. Now

$$\mathrm{TV}\left(\int f_{S(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y, \int f_{T(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n\right) \tag{71}$$

$$= \sup_h \left| \int h(x) \int f_{S(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n\,\mathrm{d}x - \int h(x) \int f_{T(y_n)}(x) f_{U_n}(y_n)\,\mathrm{d}y_n\,\mathrm{d}x \right| \tag{72}$$

$$< \epsilon \tag{73}$$

for any $\epsilon > 0$ with large enough n. □

Theorem 3.4. *Under congeniality and Condition 3.2,* $\mathrm{TV}\left(p(Q|Z), \bar{p}_n(Q)\right) \to 0$ *as* $n \to \infty$.

Proof. The claim follows from Lemma B.2 with $y_n = X_n^*$, $U_n = p(X_n^*|Z)$, $S(y_n) \sim p(Q|X_n^*)$ and $T(y_n) \sim p(Q|Z, X_n^*)$. These meet the condition for Lemma B.2 due to Lemma B.1. □

B.2 Convergence Rate

Definition 3.5. *A sequence of random variables X_n is uniformly integrable if*

$$\lim_{M\to\infty} \sup_n \mathbb{E}(|X_n| \mathbb{I}_{|X_n|>M}) = 0 \tag{12}$$

Lemma B.3. *If $|X_n| \le Y_n$ almost surely and Y_n is uniformly integrable, X_n is uniformly integrable.*

Proof.

$$0 \le \lim_{M\to\infty} \sup_n \mathbb{E}(|X_n| \mathbb{I}_{|X_n|>M}) \le \lim_{M\to\infty} \sup_n \mathbb{E}(Y_n \mathbb{I}_{Y_n>M}) = 0 \tag{74}$$

□

Lemma B.4 (Billingsley (1995), Section 16)**.** *If X_n and Y_n are uniformly integrable, $X_n + Y_n$ is uniformly integrable.*

Condition 3.6. *There exist distributions D_n such that for a sequence $R_1, R_2, \dots > 0$, $R_n \to 0$ as $n \to \infty$,*

$$\frac{1}{R_n} \mathrm{TV}\left(\bar{Q}_n^+, D_n\right) \quad \text{and} \quad \frac{1}{R_n} \mathrm{TV}\left(\bar{Q}_n, D_n\right) \tag{13}$$

are uniformly integrable when $X_n^* \sim p(X_n^*|Z)$.

Theorem 3.7. *Under congeniality and Condition 3.6,* $\mathrm{TV}\left(p(Q|Z), \bar{p}_n(Q)\right) = O(R_n)$.

Proof. Total variation distance is a metric, so

$$\frac{1}{R_n} \mathrm{TV}(\bar{Q}_n^+, \bar{Q}_n) \le \frac{1}{R_n} \mathrm{TV}\left(\bar{Q}_n^+, D_n\right) + \frac{1}{R_n} \mathrm{TV}\left(\bar{Q}_n, D_n\right). \tag{75}$$

Now Condition 3.6 and Lemmas B.3 and B.4 imply that

$$\frac{1}{R_n} \mathrm{TV}(\bar{Q}_n^+, \bar{Q}_n) \tag{76}$$

is uniformly integrable with $X_n^* \sim p(X_n^*|Z)$.

Recall that

$$\frac{1}{R_n} \mathrm{TV}(\bar{Q}_n^+, \bar{Q}_n) = \frac{1}{R_n} \sup_h \left| \int h(Q) p(Q|Z, X_n^*)\,\mathrm{d}Q - \int h(Q) p(Q|X_n^*)\,\mathrm{d}Q \right| \tag{77}$$

and

$$\frac{1}{R_n} \mathrm{TV}(p(Q|Z), \bar{p}_n(Q))$$

$$= \frac{1}{R_n} \sup_h \left| \int h(Q) \int p(Q|Z, X_n^*) p(X_n^*|Z)\,\mathrm{d}X_n^*\,\mathrm{d}Q - \int h(Q) \int p(Q|X_n^*) p(X_n^*|Z)\,\mathrm{d}X_n^*\,\mathrm{d}Q \right| \tag{78}$$

where h is an indicator function.

For any indicator function h, using the start of the proof of Lemma B.2 gives

$$\frac{1}{R_n}\left|\int h(Q)\int p(Q|Z,X_n^*)p(X_n^*|Z)\,\mathrm{d}X_n^*\,\mathrm{d}Q-\int h(Q)\int p(Q|X_n^*)p(X_n^*|Z)\,\mathrm{d}X_n^*\,\mathrm{d}Q\right| \tag{79}$$

$$\leq \int p(X_n^*|Z)\frac{1}{R_n}\left|\int h(Q)p(Q|Z,X_n^*)\,\mathrm{d}Q-\int h(Q)p(Q|X_n^*)\,\mathrm{d}Q\right|\mathrm{d}X_n^* \tag{80}$$

$$\leq \int p(X_n^*|Z)\frac{1}{R_n}\,\mathrm{TV}(\bar{Q}_n^+,\bar{Q}_n)\,\mathrm{d}X_n^* \tag{81}$$

Because $R_n^{-1}\,\mathrm{TV}(\bar{Q}_n^+,\bar{Q}_n)$ is uniformly integrable when $X_n^*\sim p(X_n^*|Z)$, there exists an M such that for all n,

$$\int_{Y_n} p(X_n^*|Z)\frac{1}{R_n}\left|\int h(Q)p(Q|Z,X_n^*)\,\mathrm{d}Q-\int h(Q)p(Q|X_n^*)\,\mathrm{d}Q\right|\mathrm{d}X_n^*\leq 1 \tag{82}$$

where $Y_n=\{X_n^*\mid \frac{1}{R_n}\,\mathrm{TV}(\bar{Q}_n^+,\bar{Q}_n)>M\}$.

Now, for all n

$$\frac{1}{R_n}\left|\int h(Q)\int p(Q|Z,X_n^*)p(X_n^*|Z)\,\mathrm{d}X_n^*\,\mathrm{d}Q-\int h(Q)\int p(Q|X_n^*)p(X_n^*|Z)\,\mathrm{d}X_n^*\,\mathrm{d}Q\right| \tag{83}$$

$$\leq \int p(X_n^*|Z)\frac{1}{R_n}\left|\int h(Q)p(Q|Z,X_n^*)\,\mathrm{d}Q-\int h(Q)p(Q|X_n^*)\,\mathrm{d}Q\right|\mathrm{d}X_n^* \tag{84}$$

$$\begin{aligned}&=\int_{Y_n} p(X_n^*|Z)\frac{1}{R_n}\left|\int h(Q)p(Q|Z,X_n^*)\,\mathrm{d}Q-\int h(Q)p(Q|X_n^*)\,\mathrm{d}Q\right|\mathrm{d}X_n^*\\&+\int_{Y_n^C} p(X_n^*|Z)\frac{1}{R_n}\left|\int h(Q)p(Q|Z,X_n^*)\,\mathrm{d}Q-\int h(Q)p(Q|X_n^*)\,\mathrm{d}Q\right|\mathrm{d}X_n^*\end{aligned} \tag{85}$$

$$\leq 1+\int_{Y_n^C} p(X_n^*|Z)M\,\mathrm{d}X_n^* \tag{86}$$

$$\leq 1+M \tag{87}$$

so $\mathrm{TV}\left(p(Q|Z),\bar{p}_n(Q)\right)=O(R_n)$. □

C Additional Examples

C.1 Gaussian with Known Variance Details

Checking where the mean and variance of $\mu^*\sim\bar{p}_n(\mu)$ converge when $n_{X^*}\to\infty$ in the Gaussian mean estimation example, when both parties use the known variance model:

$$\mathbb{E}(\mu^*)=\mathbb{E}(\mathbb{E}(\mu^*|X^*))=\mathbb{E}\left(\hat{\mu}_{n_{X^*}}\right) \tag{88}$$

$$=\mathbb{E}\left(\frac{\frac{1}{\hat{\sigma}_0^2}\hat{\mu}_0+\frac{n_{X^*}}{\hat{\sigma}_k^2}\bar{X}^*}{\frac{1}{\hat{\sigma}_0^2}+\frac{n_{X^*}}{\hat{\sigma}_k^2}}\right) \tag{89}$$

$$=\frac{\frac{1}{\hat{\sigma}_0^2}\hat{\mu}_0+\frac{n_{X^*}}{\hat{\sigma}_k^2}\mathbb{E}(\bar{X}^*)}{\frac{1}{\hat{\sigma}_0^2}+\frac{n_{X^*}}{\hat{\sigma}_k^2}} \tag{90}$$

$$\to\mathbb{E}(X^*)=\bar{\mu}_{n_X} \tag{91}$$

as $n_{X^*}\to\infty$.

$$\mathrm{Var}(\mu^*)=\mathbb{E}(\mathrm{Var}(\mu^*|X^*))+\mathrm{Var}(\mathbb{E}(\mu^*|X^*)) \tag{92}$$

$$=\mathbb{E}(\hat{\sigma}^2_{n_{X^*}})+\mathrm{Var}(\hat{\mu}_{n_{X^*}}) \tag{93}$$

$$\mathbb{E}(\hat{\sigma}^2_{n_{X^*}}) = \mathbb{E}\left(\frac{1}{\frac{n_{X^*}}{\hat{\sigma}_k^2} + \frac{1}{\hat{\sigma}_0^2}}\right) \to 0, n_{X^*} \to \infty \tag{94}$$

$$\mathrm{Var}\left(\hat{\mu}_{n_{X^*}}\right) = \mathrm{Var}\left(\frac{\frac{n_{X^*}}{\hat{\sigma}_k^2}\bar{X}^* + \frac{\hat{\mu}_0}{\hat{\sigma}_0^2}}{\frac{n_{X^*}}{\hat{\sigma}_k^2} + \frac{1}{\hat{\sigma}_0^2}}\right) = \left(\frac{\frac{n_{X^*}}{\hat{\sigma}_k^2}}{\frac{n_{X^*}}{\hat{\sigma}_k^2} + \frac{1}{\hat{\sigma}_0^2}}\right)^2 \mathrm{Var}(\bar{X}^*) \tag{95}$$

$$\mathrm{Var}(\bar{X}^*) = \mathbb{E}(\mathrm{Var}(\bar{X}^*|\bar{\mu})) + \mathrm{Var}(\mathbb{E}(\bar{X}^*|\bar{\mu})) = \frac{1}{n_{X^*}}\mathbb{E}(\mathrm{Var}(x_i^*)) + \mathrm{Var}(\bar{\mu}) \to \mathrm{Var}(\bar{\mu}) = \bar{\sigma}^2_{n_X} \tag{96}$$

as $n_{X^*} \to \infty$.

Putting these together,

$$\mathbb{E}(\mu^*) \to \bar{\mu}_{n_X} \tag{97}$$

$$\mathrm{Var}(\mu^*) \to \bar{\sigma}^2_{n_X} \tag{98}$$

as $n_{X^*} \to \infty$.

The plots of $p(\mu^*)$ in Figures 2, 3, S1, S2 and S3 are density functions of a mixture of Gaussians, where each mixture component is the Gaussian posterior distribution from one synthetic dataset.

C.2 Gaussian with Unknown Variance Upstream, Known Variance Downstream

When the synthetic data is generated from the unknown variance model, $p(X^*|X)$ is

$$\bar{\sigma}^2|X \sim \text{Inv-}\chi^2(\bar{\nu}_{n_X}, \bar{\sigma}^2_{n_X}) \tag{99}$$

$$\bar{\mu}|\bar{\sigma}^2, X \sim \mathcal{N}\left(\bar{\mu}_{n_X}, \frac{\bar{\sigma}^2}{\bar{\kappa}_{n_X}}\right) \tag{100}$$

$$x_i^*|\bar{\mu}, \bar{\sigma}^2 \sim \mathcal{N}(\bar{\mu}, \bar{\sigma}^2). \tag{101}$$

$$\tag{102}$$

When downstream analysis is the model with known variance $\hat{\sigma}_k^2$, $p(\mu^*|X^*)$ is

$$\mu^*|X^* \sim \mathcal{N}(\hat{\mu}_{n_{X^*}}, \hat{\sigma}^2_{n_{X^*}}) \tag{103}$$

$$\hat{\mu}_{n_{X^*}} = \frac{\frac{1}{\hat{\sigma_0}^2}\hat{\mu}_0 + \frac{n_{X^*}}{\hat{\sigma}_k^2}\bar{X}^*}{\frac{1}{\hat{\sigma}_0^2} + \frac{n_{X^*}}{\hat{\sigma}_k^2}} \tag{104}$$

$$\frac{1}{\hat{\sigma}^2_{n_{X^*}}} = \frac{1}{\hat{\sigma}_0^2} + \frac{n_{X^*}}{\hat{\sigma}_k^2}. \tag{105}$$

Checking where the mean and variance of $\bar{p}_n(\mu)$ converge when $n_{X^*} \to \infty$:

$$\mathbb{E}(\mu^*) = \mathbb{E}(\mathbb{E}(\mu^*|X^*)) = \mathbb{E}\left(\hat{\mu}_{n_{X^*}}\right) \tag{106}$$

$$= \mathbb{E}\left(\frac{\frac{1}{\hat{\sigma}_0^2}\hat{\mu}_0 + \frac{n_{X^*}}{\hat{\sigma}_k^2}\bar{X}^*}{\frac{1}{\hat{\sigma}_0^2} + \frac{n_{X^*}}{\hat{\sigma}_k^2}}\right) \tag{107}$$

$$= \frac{\frac{1}{\hat{\sigma}_0^2}\hat{\mu}_0 + \frac{n_{X^*}}{\hat{\sigma}_k^2}\mathbb{E}(\bar{X}^*)}{\frac{1}{\hat{\sigma}_0^2} + \frac{n_{X^*}}{\hat{\sigma}_k^2}} \tag{108}$$

$$\to \mathbb{E}(X^*) = \bar{\mu}_{n_X} \tag{109}$$

as $n_{X^*} \to \infty$.

$$\mathrm{Var}(\mu^*) = \mathbb{E}(\mathrm{Var}(\mu^*|X^*)) + \mathrm{Var}(\mathbb{E}(\mu^*|X^*)) \tag{110}$$

$$= \mathbb{E}(\hat{\sigma}^2_{n_{X^*}}) + \mathrm{Var}(\hat{\mu}_{n_{X^*}}) \tag{111}$$

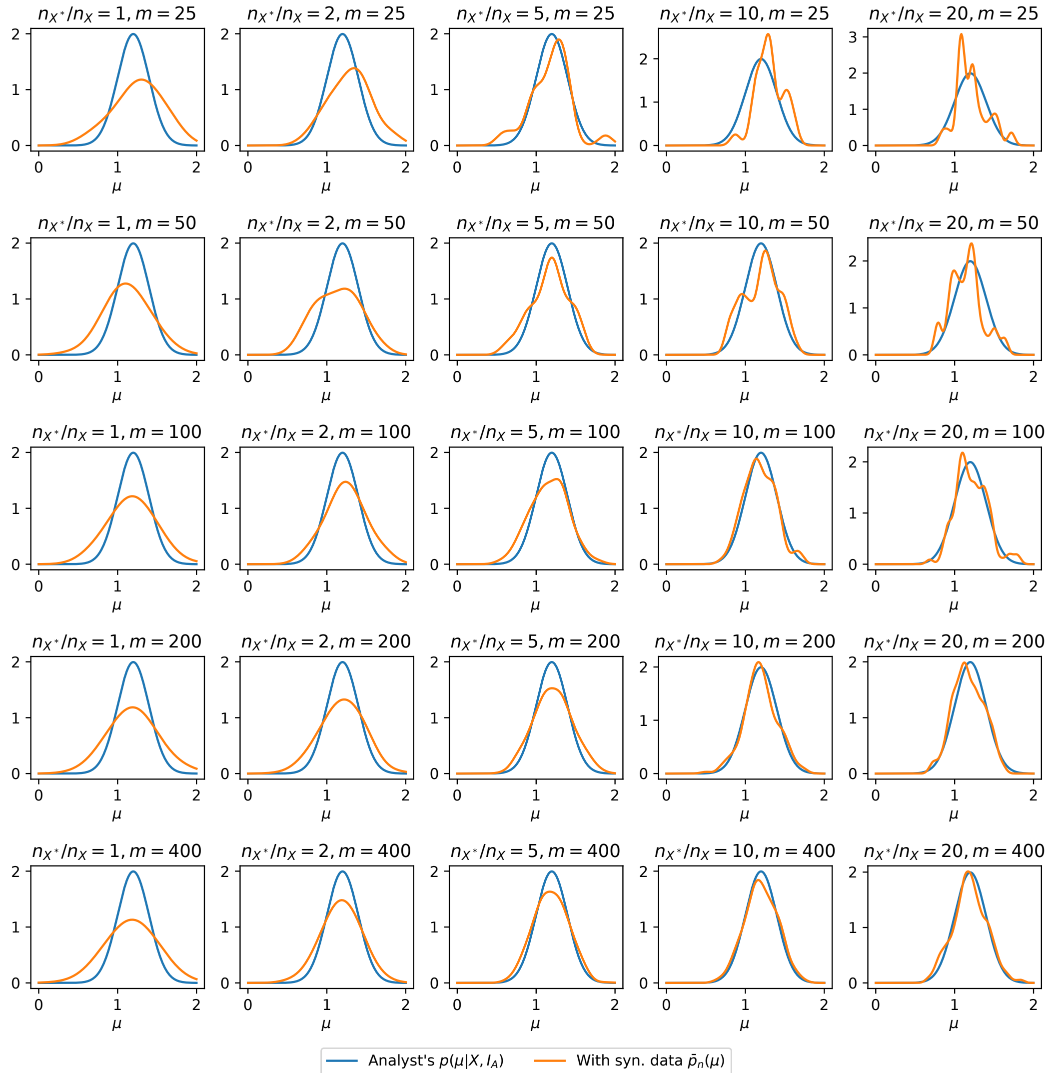

Figure S1: Convergence of the mixture of synthetic data posteriors (in orange) with different values of m and n_{X^*} in Gaussian mean estimation with known variance.

$$\mathbb{E}(\hat{\sigma}^2_{n_{X^*}}) = \mathbb{E}\left(\frac{1}{\frac{n_{X^*}}{\hat{\sigma}_k^2} + \frac{1}{\hat{\sigma}_0^2}}\right) \to 0, n_{X^*} \to \infty \tag{112}$$

$$\mathrm{Var}\left(\hat{\mu}_{n_{X^*}}\right) = \mathrm{Var}\left(\frac{\frac{n_{X^*}}{\hat{\sigma}_k^2}\bar{X}^* + \frac{\hat{\mu}_0}{\hat{\sigma}_0^2}}{\frac{n_{X^*}}{\hat{\sigma}_k^2} + \frac{1}{\hat{\sigma}_0^2}}\right) = \left(\frac{\frac{n_{X^*}}{\hat{\sigma}_k^2}}{\frac{n_{X^*}}{\hat{\sigma}_k^2} + \frac{1}{\hat{\sigma}_0^2}}\right)^2 \mathrm{Var}(\bar{X}^*) \tag{113}$$

$$\mathrm{Var}(\bar{X}^*) = \mathbb{E}(\mathrm{Var}(\bar{X}^*|\bar{\mu}, \bar{\sigma}^2)) + \mathrm{Var}(\mathbb{E}(\bar{X}^*|\bar{\mu}, \bar{\sigma}^2)) = \frac{1}{n_{X^*}}\mathbb{E}(\bar{\sigma}^2) + \mathrm{Var}(\bar{\mu}) \to \mathrm{Var}(\bar{\mu}) = \frac{\bar{\sigma}_0^2}{\bar{\kappa}_{n_X}} \tag{114}$$

as $n_{X^*} \to \infty$.

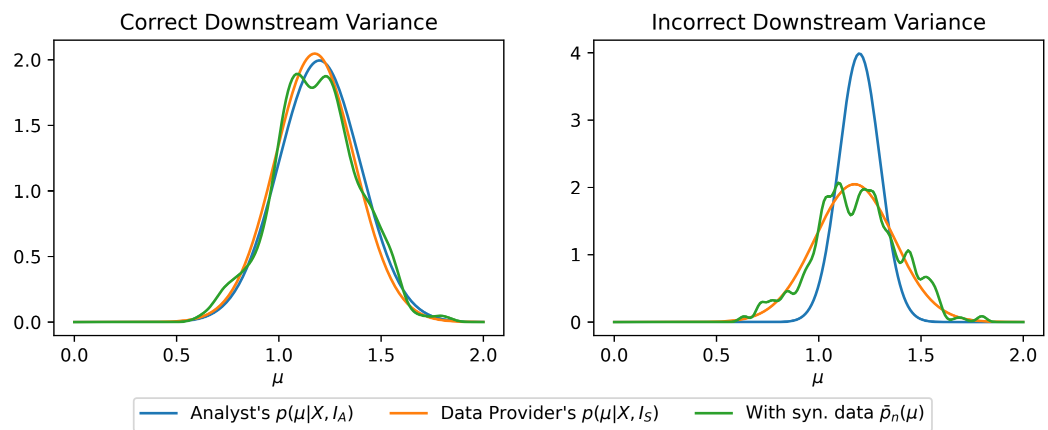

Figure S2: Results when the synthetic data is generated from the unknown variance Gaussian mean estimation model, and the analyst uses the model with known variance. On the left, the analyst's known variance is correct, on the right it is incorrect. In both cases, the mixture of synthetic data posteriors converges to the data provider's posterior. In both panels, $m = 400$ and $\frac{n_{X^*}}{n_X} = 20$.

Putting these together,

$$\mathbb{E}(\mu^*) \to \bar{\mu}_{n_X} \tag{115}$$

$$\mathrm{Var}(\mu^*) \to \frac{\bar{\sigma}_0^2}{\bar{\kappa}_{n_X}} \tag{116}$$

as $n_{X^*} \to \infty$, so μ^* asymptotically has the same mean and variance as the marginal posterior of μ in the synthetic data model, which is not the same as the downstream posterior distribution $p(\mu|X, I_A)$ on the real data.

We verify this with the simulation in Figure S2, where the synthetic data is generated from the model with unknown variance, while the analyst uses the known variance model. The setting is otherwise identical to the case where both used the known variance model in Figure 2. The mixture of synthetic data posteriors converges to the data provider's posterior, even when the analyst uses an incorrect value for the known variance $\hat{\sigma}_k^2$.

C.3 Size of the Synthetic Dataset

In the preceding analysis, most of the approximations hold when n_{X^*} is large, even when $n_{X^*} \approx n_X$. However, based on the experiment with different values of n_{X^*} and m in Figure S1, $n_{X^*} \gg n_X$ is needed for all of the approximations to hold.

This is explained by looking at $\mathrm{Var}(\bar{X}^*)$. In the case where both parties use the known variance model,

$$\begin{aligned}
\mathrm{Var}(\bar{X}^*) &= \frac{1}{n_{X^*}}\mathbb{E}(\mathrm{Var}(x_i^*)) + \mathrm{Var}(\bar{\mu}) && (117)\\
&= \frac{1}{n_{X^*}}(\bar{\sigma}_k^2 + \bar{\sigma}_{n_X}^2) + \bar{\sigma}_{n_X}^2 && (118)\\
&= \frac{1}{n_{X^*}}\bar{\sigma}_k^2 + \left(1 + \frac{1}{n_{X^*}}\right)\frac{1}{\frac{1}{\bar{\sigma}_0^2} + \frac{n_X}{\bar{\sigma}_k^2}} && (119)
\end{aligned}$$

If $n_X \approx n_{X^*}$ and both are large, $1 + \frac{1}{n_{X^*}} \approx 1$ and $\frac{1}{\bar{\sigma}_0^2} + \frac{n_X}{\bar{\sigma}_k^2} \approx \frac{n_X}{\bar{\sigma}_k^2}$, so

$$\mathrm{Var}(\bar{X}^*) \approx \frac{\bar{\sigma}_k^2}{n_{X^*}} + \frac{\bar{\sigma}_k^2}{n_X} \approx \frac{2\bar{\sigma}_k^2}{n_X} \tag{120}$$

With these approximations,

$$\mathrm{Var}(\bar{\mu}) \approx \frac{\bar{\sigma}_k^2}{n_X} \tag{121}$$

so

$$\mathrm{Var}(\bar{X^*}) \approx 2\mathrm{Var}(\bar{\mu}) \tag{122}$$

while the $n_{X^*} \to \infty$ limit is $\mathrm{Var}(\bar{X^*}) \to \mathrm{Var}(\bar{\mu})$. This means that $n_{X^*} \gg n_X$ is required.

The same happens when the synthetic data is generated from the unknown variance model:

$$\mathrm{Var}(\bar{X^*}) = \frac{1}{n_{X^*}}\mathbb{E}(\bar{\sigma}^2) + \mathrm{Var}(\bar{\mu}) \tag{123}$$

$$= \frac{1}{n_{X^*}}\frac{\bar{\nu}_0 + n_X}{\bar{\nu}_0 + n_X - 2}\bar{\sigma}^2_{n_X} + \frac{\bar{\sigma}^2_{n_X}}{\bar{\kappa}_0 + n_X} \tag{124}$$

If $n_X \approx n_{X^*}$ and both are large, $\frac{\bar{\nu}_0+n_X}{\bar{\nu}_0+n_X-2} \approx 1$ and $\bar{\kappa}_0 + n_X \approx n_X$, so

$$\mathrm{Var}(\bar{X^*}) \approx \frac{\bar{\sigma}^2_{n_X}}{n_{X^*}} + \frac{\bar{\sigma}^2_{n_X}}{n_X} \approx \frac{2\bar{\sigma}^2_{n_X}}{n_X} \tag{125}$$

With these approximations,

$$\mathrm{Var}(\bar{\mu}) \approx \frac{\bar{\sigma}^2_{n_X}}{n_X} \tag{126}$$

so

$$\mathrm{Var}(\bar{X^*}) \approx 2\mathrm{Var}(\bar{\mu}). \tag{127}$$

C.4 Further Approximation

When n_{X^*} is large,

$$\mathrm{Var}(\mu^*) \approx \mathbb{E}(\hat{\sigma}^2_{X^*}) + \mathrm{Var}(\bar{X^*}) \tag{128}$$

If $n_{X^*} = cn_X$ for some $c > 0$, from the analyses in Section C.3, we get

$$\mathrm{Var}(\bar{X}) \approx \left(1 + \frac{1}{c}\right)\mathrm{Var}(\bar{\mu}) \tag{129}$$

so

$$\mathrm{Var}(\mu^*) \approx \mathbb{E}(\hat{\sigma}^2_{n_{X^*}}) + \left(1 + \frac{1}{c}\right)\mathrm{Var}(\bar{\mu}) \tag{130}$$

Solving for $\mathrm{Var}(\mu)$ gives

$$\mathrm{Var}(\bar{\mu}) \approx \left(1 + \frac{1}{c}\right)^{-1}\left(\mathrm{Var}(\mu^*) - \mathbb{E}(\hat{\sigma}^2_{n_{X^*}})\right) \tag{131}$$

which gives a Rubin's rules-like approximation of $\mathrm{Var}(\mu)$ that can be computed from smaller synthetic datasets with $n_{X^*} \approx n_X$.

We validate this with the experiment in Figure S3, which shows that approximating $p(\mu|X, I_A)$ with a Gaussian with variance from (131) is closer to the real data posterior than the mixed posterior approximation from Section 3.

C.5 Logistic Regression

Setting Details In the toy data setting of Räisä et al. (2023), the real dataset consists of $n_X = 2000$ i.i.d. samples of three binary variables. The first two variables are sampled with independent coinflips, and the third is sampled from logistic regression on the other two, with coefficients $(1, 0)$.

We generate synthetic data with the NAPSU-MQ algorithm (Räisä et al. 2023), instructing the algorithm to generate m synthetic datasets of size n_{X^*}. DP-GLM doesn't use synthetic data, so we run it directly on the real data. For the privacy bounds, we vary ϵ, and set $\delta = n_X^{-2} = 2.5 \cdot 10^{-7}$, which is how DP mechanisms are typically evaluated.

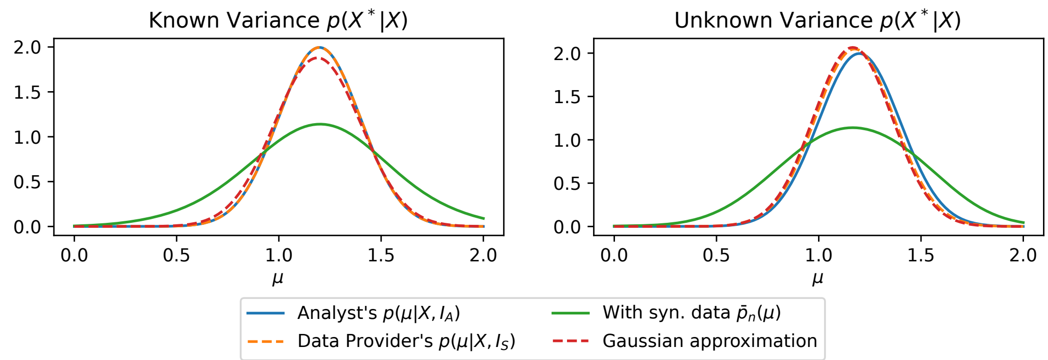

Figure S3: Results with the Gaussian approximation with $n_{X^*} = n_X$, showing that the Gaussian approximation is closer to the real data posterior than the mixture of synthetic data posteriors. On the left, the synthetic data is generated from the known variance model, and on the right, the synthetic data is generated from the unknown variance model. In both cases, the known variances for both parties are correct, and $m = 400$.

Hyperparameters For NAPSU-MQ, we use the hyperparameters of Räisä et al. (2023), except we used NUTS (Hoffman and Gelman 2014) with 200 warmup samples and 500 kept samples for $\epsilon \in \{0.5, 1\}$, and 1500 kept samples for $\epsilon = 0.1$, as the posterior sampling algorithm. The prior is $\mathcal{N}(0, 10^2 I)$, and the marginal queries are the full set of 3-way marginals of all three variables.

The hyperparameters of DP-GLM are the L_2-norm upper bound R for the covariates of the logistic regression, a coefficient norm upper bound s, and the parameters of the posterior sampling algorithm DP-GLM uses. We set $R = \sqrt{2}$ so that the covariates do not get clipped, and set $s = 5$ after some preliminary runs. The posterior sampling algorithm is NUTS (Hoffman and Gelman 2014) with 1000 warmup iterations and 1000 kept samples from 4 parallel chains.

Plotting Details The plotted density of DP-GLM in Figure 4 is a kernel density estimate from the posterior samples DP-GLM returns. The non-DP density is a Laplace approximation. Both synthetic data methods use Laplace approximations in the downstream analysis, so their posteriors are mixtures of these Laplace approximations for each synthetic dataset. This was also used in Figure S5.

Sampling the exact posterior In order to sample the exact posterior $p(Q|\tilde{s})$, we use another decomposition:

$$p(Q|\tilde{s}) = \int p(Q|\tilde{s}, X)p(X|\tilde{s})\,\mathrm{d}X = \int p(Q|X)p(X|\tilde{s})\,\mathrm{d}X, \tag{132}$$

where $p(Q|\tilde{s}, X) = p(Q|X)$ due to the independencies of the graphical model in Figure 1. It remains to sample $p(X|\tilde{s})$. This is not tractable in general, but is possible in the toy data setting due to using the full set of 3-way marginals that covers all possible values of a datapoint, and the simplicity of the toy data.

We can decompose

$$p(X|\tilde{s}) = \int p(s|\tilde{s})p(X|s)\,\mathrm{d}\theta\,\mathrm{d}X = \int p(X|s) \int p(s, \theta|\tilde{s})\,\mathrm{d}\theta\,\mathrm{d}X, \tag{133}$$

so we can sample $(s, \theta) \sim p(s, \theta|s)$ and then sample $X \sim p(X|s)$ to obtain a sample from $p(X|\tilde{s})$. Due to the simplicity of the toy data setting, sampling both $p(s, \theta|s)$ and $p(X|s)$ is possible.

NAPSU-MQ uses the following Bayesian inference problem:

$$\theta \sim \text{Prior} \tag{134}$$
$$X \sim \text{MED}_\theta^n \tag{135}$$
$$s = a(X) \tag{136}$$
$$\tilde{s} \sim \mathcal{N}(s, \sigma_{DP}^2), \tag{137}$$

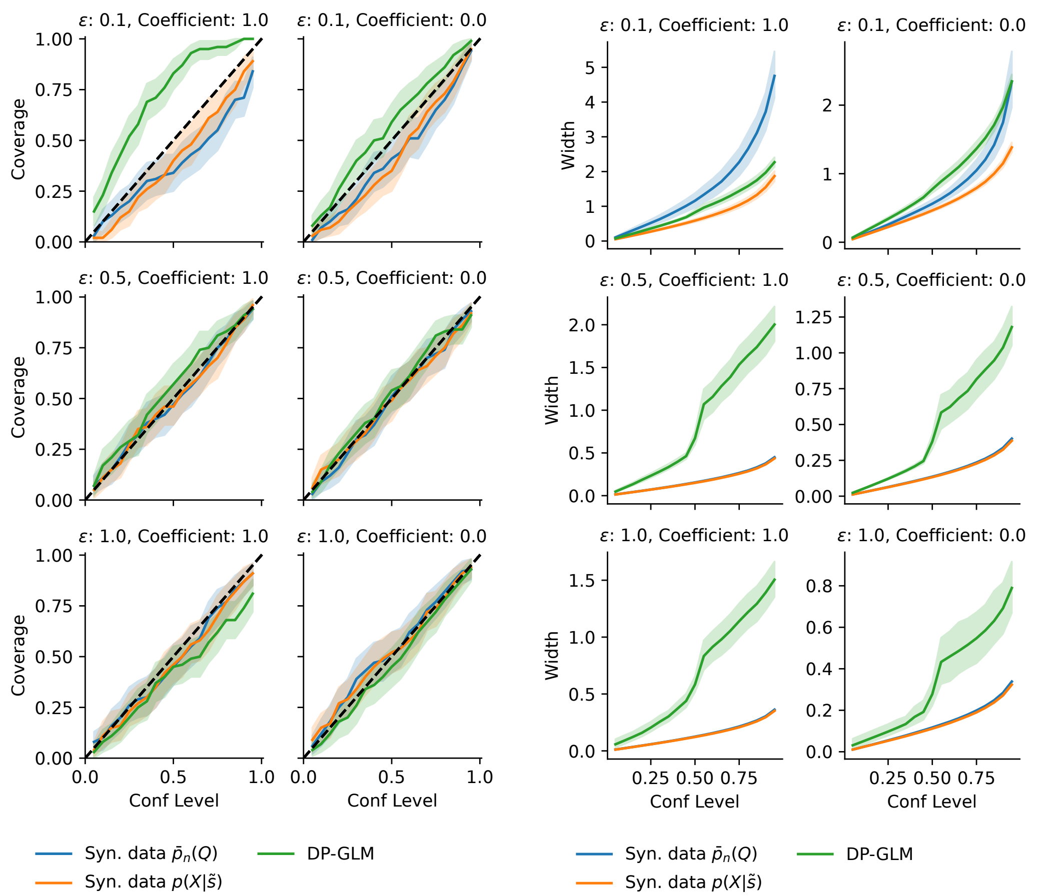

(a) (b)

Figure S4: (a) Coverages of credible intervals in the toy data experiment. The mixture of synthetic data posteriors is accurate, except with $\epsilon = 0.1$, where it may not have converged yet. (b) Widths of credible intervals in the toy data experiment. DP-GLM produces much wider intervals than other methods, except with $\epsilon = 0.1$.

where a are the marginal queries, σ_{DP}^2 is the noise variance of the Gaussian mechanism, and MED_θ^n is the maximum entropy distribution (Räisä et al. 2023) with point probability

$$p(x) = \exp(\theta^T a(x) - \theta_0(\theta)), \tag{138}$$

where θ_0 is the log-normalising constant.

In the toy data setting, a is the full set of 3-way marginals for all of the 3 variables. In other words, $a(x)$ is the one-hot coding of x, so $s = a(X)$ is a vector of counts of how many times each of the 8 possible values is repeated in X. This means that sampling $p(X|s)$ is simple:

1. For each possible value of a datapoint, find the corresponding count from s, and repeat that datapoint according the that count.
2. Shuffle the datapoints to a random order.

As the downstream analysis $p(Q|X)$ doesn't depend on the order of the datapoints, the second step is not actually needed.

To sample $p(s, \theta|\tilde{s})$, we use a Metropolis-within-Gibbs sampler (Gilks et al. 1995) that sequentially updates s and θ while keeping the other fixed. The proposal for θ is obtained from Hamiltonian Monte Carlo (HMC) (Duane et al. 1987; Neal 2011). The proposal for s is obtained by repeatedly choosing a random index in s to increment and another to decrement. It is possible to obtain negative values in

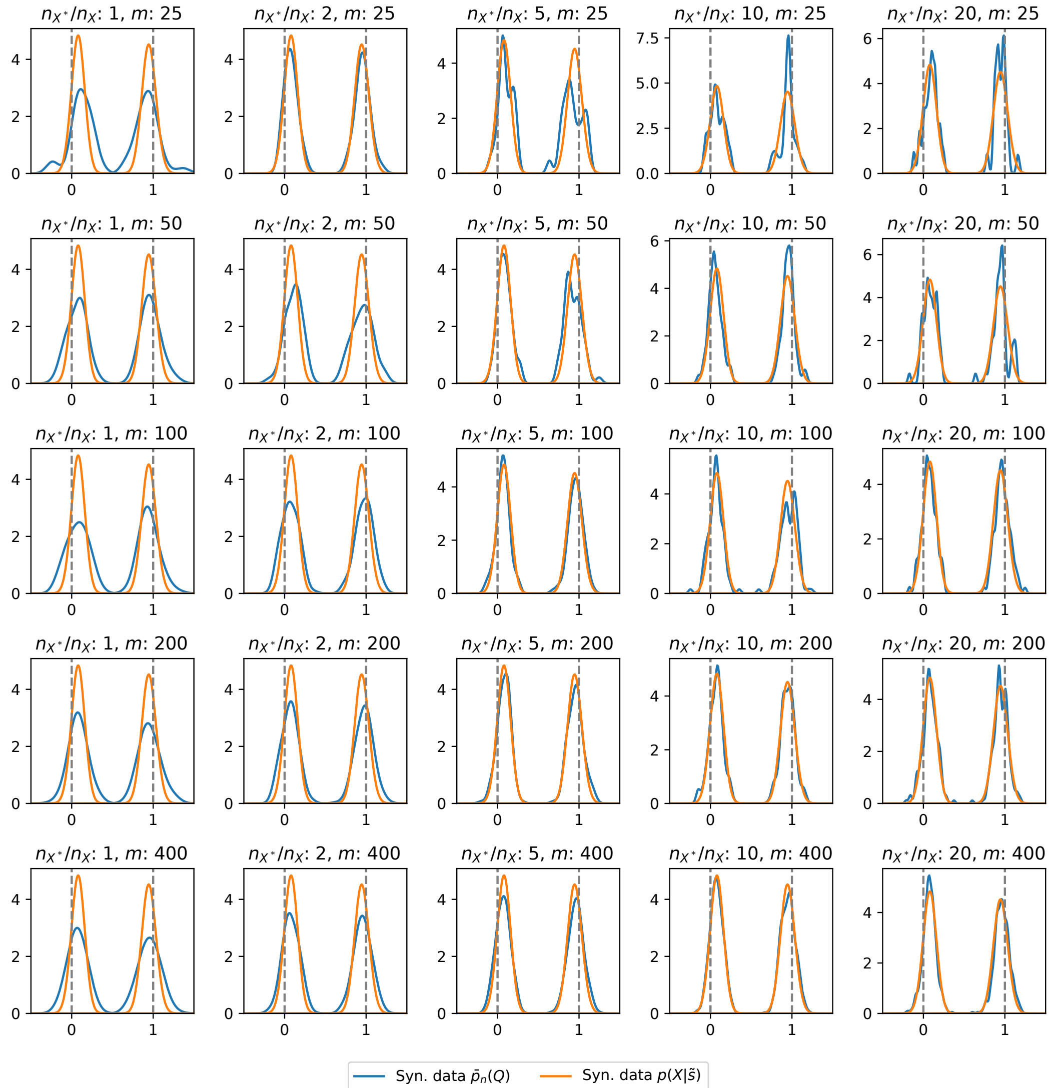

Figure S5: Convergence of the mixture of synthetic data posteriors (in blue) with different values of m and n_{X^*} in the toy data logistic regression experiment.

s from this proposal, but those will always be rejected by the acceptance test, as the likelihood for them is 0.

To initialise the sampler, we pick an initial value for θ from a Gaussian distribution, and pick the initial s by rounding $\tilde{s}$ to integer values, changing the rounded values such that they sum to n while ensuring that all values are non-negative.

The step size for the HMC we used is 0.05, and the number of steps is 20. In the s proposal, we repeat the combination of an increment and a decrement 30 times. We take 20000 samples in total from 4 parallel chains, and drop the first 20% as warmup samples.

The method described in this section is similar to the noise-aware Bayesian inference method of Ju et al. (2022). The difference between the two is that Ju et al. (2022) use X instead of s as the auxiliary variable, and they sample the X proposals from the model, changing one datapoint at a time. This makes their algorithm more generalisable.

D Gaussian Known Variance Convergence Rate

Theorem D.1. *When the up- and downstream models are Gaussian mean estimations with known variance, when $D_n = \mathcal{N}(\bar{X}, n^{-1}\sigma_k^2)$,*

$$\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n, D_n\right) \tag{139}$$

and

$$\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n^+, D_n\right) \tag{140}$$

are uniformly integrable when $X_n^* \sim p(X_n^*|X)$.

Proof. When the downstream model is Gaussian mean estimation with known variance,

$$\bar{Q}_n = \mathcal{N}(\mu_n, \sigma_n^2) \tag{141}$$

$$\mu_n = \frac{\frac{1}{\sigma_0^2}\mu_0 + \frac{n}{\sigma_k^2}\bar{X}}{\frac{1}{\sigma_0^2} + \frac{n}{\sigma_k^2}} \tag{142}$$

$$\frac{1}{\sigma_n^2} = \frac{1}{\sigma_0^2} + \frac{n}{\sigma_k^2} \tag{143}$$

We start with the proof for $\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n, D_n\right)$. By Pinsker's equality and the formula for KL-divergence between Gaussians,

$$\begin{aligned}
\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n, D_n\right) &\leq \sqrt{\frac{1}{2} n \mathrm{KL}(\bar{Q}_n \,||\, D_n)} && (144)\\
&= \sqrt{\frac{1}{4} n \left(\frac{\sigma_k^2}{n\sigma_n^2} + \frac{(\mu_n - \bar{X})^2}{\sigma_n^2} - 1 + \ln\frac{n\sigma_n^2}{\sigma_k^2}\right)} && (145)\\
&\leq \sqrt{\left|\frac{1}{4} n \left(\frac{\sigma_k^2}{n\sigma_n^2} - 1\right)\right| + \frac{1}{4} n \frac{(\mu_n - \bar{X})^2}{\sigma_n^2} + \left|\frac{1}{4} n \ln\frac{n\sigma_n^2}{\sigma_k^2}\right|} && (146)\\
&\leq \sqrt{\left|\frac{1}{4} n \left(\frac{\sigma_k^2}{n\sigma_n^2} - 1\right)\right|} + \sqrt{\frac{1}{4} n \frac{(\mu_n - \bar{X})^2}{\sigma_n^2}} + \sqrt{\left|\frac{1}{4} n \ln\frac{n\sigma_n^2}{\sigma_k^2}\right|} && (147)
\end{aligned}$$

The last inequality can be deduced from the fact that the L_2-norm is upper bounded by the L_1 norm.

Denote

$$C_1(n) = n\left(\frac{\sigma_k^2}{n\sigma_n^2} - 1\right) = \left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma_k^2}\right)\sigma_k^2 - n = \frac{\sigma_k^2}{\sigma_0^2} + n - n = \frac{\sigma_k^2}{\sigma_0^2} \tag{148}$$

and

$$\begin{aligned}
C_2(n) &= n \ln\frac{n\sigma_n^2}{\sigma_k^2} && (149)\\
&= -n \ln\frac{\sigma_k^2}{n\sigma_n^2} && (150)\\
&= -n \ln\left(\left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma_k^2}\right)\frac{\sigma_k^2}{n}\right) && (151)\\
&= -n \ln\left(\frac{\sigma_k^2}{n\sigma_0^2} + 1\right) && (152)\\
&= -\frac{u\sigma_k^2}{\sigma_0^2} \ln\left(\frac{1}{u} + 1\right) && (153)\\
&= -\frac{\sigma_k^2}{\sigma_0^2} \ln\left(\frac{1}{u} + 1\right)^u && (154)
\end{aligned}$$

with $n = \frac{u\sigma_k^2}{\sigma_0^2}$. Because

$$\lim_{u\to\infty}\left(1+\frac{1}{u}\right)^u = e \tag{155}$$

we have

$$\lim_{u\to\infty} -\frac{\sigma_k^2}{\sigma_0^2}\ln\left(\frac{1}{u}+1\right)^u = -\frac{\sigma_k^2}{\sigma_0^2} \tag{156}$$

which implies that $C_2(n)$ is bounded.

Futhermore,

$$\sqrt{\frac{1}{4}n\frac{(\mu_n-\bar{X})^2}{\sigma_n^2}} = \frac{1}{2}\sqrt{n\left(\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}\right)}|\mu_n-\bar{X}| \tag{157}$$

Denote

$$s_n = \frac{1}{2}\sqrt{n\left(\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}\right)} \tag{158}$$

Note that $s_n = O(n)$.

Then

$$\sqrt{\frac{1}{4}n\frac{n(\mu_n-\bar{X})^2}{n\sigma_n^2}} = s_n|\mu_n-\bar{X}| \tag{159}$$

$$\begin{aligned}
s_n|\mu_n-\bar{X}| &= s_n\left|\frac{\frac{1}{\sigma_0^2}\mu_0}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}+\frac{\frac{n}{\sigma_k^2}}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}\bar{X}-\bar{X}\right| &(160)\\
&\le s_n\frac{\frac{1}{\sigma_0^2}\mu_0}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}+s_n\left|\frac{\frac{n}{\sigma_k^2}}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}-1\right||\bar{X}| &(161)\\
&= s_n\frac{\frac{1}{\sigma_0^2}\mu_0}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}+s_n\left|\frac{\frac{n}{\sigma_k^2}}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}-\frac{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}\right||\bar{X}| &(162)\\
&= s_n\frac{\frac{1}{\sigma_0^2}\mu_0}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}+s_n\frac{\frac{1}{\sigma_0^2}}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}|\bar{X}|. &(163)
\end{aligned}$$

Denote

$$C_3(n) = s_n\frac{\frac{1}{\sigma_0^2}\mu_0}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}} \tag{164}$$

and

$$C_4(n) = s_n\frac{\frac{1}{\sigma_0^2}}{\frac{1}{\sigma_0^2}+\frac{n}{\sigma_k^2}}. \tag{165}$$

Because $s_n = O(n)$, we have $C_3(n) = O(1)$ and $C_4(n) = O(1)$, so $C_3(n)$ and $C_4(n)$ are bounded.

We now have

$$\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n, D_n\right) \le \sqrt{\frac{1}{4}|C_1(n)|}+\sqrt{\frac{1}{4}|C_2(n)|}+C_3(n)+C_4(n)\bar{X}. \tag{166}$$

By Lemmas B.3 and B.4, it suffices to show that each of the terms on the right is uniformly integrable. The terms containing C_1, C_2 and C_3 are non-random and bounded in n, so they are uniformly integrable. It remains to show that $C_4(n)\bar{X}$ is uniformly integrable. $C_4(n)$ is bounded, so we only need to show that $\bar{X}$ is uniformly integrable.

To bound the expectation in the definition of uniform integrability for $|\bar{X}|$, we need some background facts. For geometric series, with $a \in \mathbb{R}$ and $|r| < 1$,

$$\sum_{i=0}^{\infty} ar^i = \frac{a}{1-r}, \tag{167}$$

and differentiating both sides with regards to r gives

$$\sum_{i=0}^{\infty} a(i+1)r^i = \frac{a}{(1-r)^2}. \tag{168}$$

For a Gaussian random variable Y with mean μ and variance σ, $\Pr(Y > \mu + t) \leq e^{-\frac{t^2}{2\sigma^2}}$. $\bar{X} \sim \mathcal{N}\left(\mu, \frac{1}{n}\sigma_k^2\right)$, so this tail bound gives

$$\Pr(\bar{X} > t + \mu) \leq 2e^{-\frac{nt^2}{2\sigma_k^2}}. \tag{169}$$

By the symmetry of the Gaussian distribution,

$$\Pr(\bar{X} < \mu - t) \leq 2e^{-\frac{nt^2}{2\sigma_k^2}}. \tag{170}$$

Now

$$\lim_{M\to\infty} \sup_n \mathbb{E}\big(|\bar{X}|\mathbb{I}_{|\bar{X}|>M}\big) \tag{171}$$

$$= \lim_{M\to\infty} \sup_n \mathbb{E}_\mu\big(\mathbb{E}(|\bar{X}|\mathbb{I}_{|\bar{X}|>M}|\mu)\big) \tag{172}$$

$$= \lim_{M\to\infty} \sup_n \mathbb{E}_\mu\left(\sum_{i=0}^{\infty} \mathbb{E}\big(|\bar{X}|\mathbb{I}_{M+i<|\bar{X}|\leq M+i+1}|\mu\big)\right) \tag{173}$$

$$\leq \lim_{M\to\infty} \sup_n \mathbb{E}_\mu\left(\sum_{i=0}^{\infty} \mathbb{E}\big((M+i+1)\mathbb{I}_{|\bar{X}|>M+i}|\mu\big)\right) \tag{174}$$

$$= \lim_{M\to\infty} \sup_n \mathbb{E}_\mu\left(\sum_{i=0}^{\infty} (M+i+1)\Pr\left(|\bar{X}| > M+i|\mu\right)\right) \tag{175}$$

$$= \lim_{M\to\infty} \sup_n \mathbb{E}_\mu\left(\sum_{i=0}^{\infty} (M+i+1)\Big(\Pr\left(\bar{X} > M+i|\mu\right) + \Pr\left(\bar{X} < -M-i|\mu\right)\Big)\right) \tag{176}$$

$$\leq \lim_{M\to\infty} \sup_n \mathbb{E}_\mu\left(\sum_{i=0}^{\infty} (M+i+1)\left(e^{-\frac{n(M+i-\mu)^2}{2\sigma_k^2}} + e^{-\frac{n(\mu+M+i)^2}{2\sigma_k^2}}\right)\right) \tag{177}$$

$$\leq \lim_{M\to\infty} \mathbb{E}_\mu\left(\sum_{i=0}^{\infty} (M+i+1)\left(e^{-\frac{(M+i-\mu)^2}{2\sigma_k^2}} + e^{-\frac{(\mu+M+i)^2}{2\sigma_k^2}}\right)\right). \tag{178}$$

When $|M+i-\mu| \geq 1$, $(M+i-\mu)^2 \geq M+i-\mu$. It is possible that $|M+i-\mu| < 1$ for exactly two values of i that depend on μ. Let $i_{\mu 1}$ and $i_{\mu 2}$ be those values. We know that $i_{\mu j} < 1 + \mu - M$ for $j \in \{1, 2\}$. Now

$$\begin{aligned}\mathbb{E}_\mu\left(\sum_{i=0}^{\infty}(M+i+1)e^{-\frac{(M+i-\mu)^2}{2\sigma_k^2}}\right) &= \mathbb{E}_\mu\left(\sum_{i=0, i\neq i_{\mu 1}, i\neq i_{\mu 2}}^{\infty}(M+i+1)e^{-\frac{(M+i-\mu)^2}{2\sigma_k^2}}\right)\\ &\quad + \mathbb{E}_\mu\left((M+i_{\mu 1}+1)e^{-\frac{(M+i_{\mu 1}-\mu)^2}{2\sigma_k^2}}\right)\\ &\quad + \mathbb{E}_\mu\left((M+i_{\mu 2}+1)e^{-\frac{(M+i_{\mu 2}-\mu)^2}{2\sigma_k^2}}\right),\end{aligned} \tag{179}$$

Now we can upper bound the series using $(M+i-\mu)^2 \geq M+i-\mu$ and the formulas for geometric series.

$$\mathbb{E}_\mu\left(\sum_{i=0, i\neq i_{\mu 1}, i\neq i_{\mu 2}}^{\infty}(M+i+1)e^{-\frac{(M+i-\mu)^2}{2\sigma_k^2}}\right) \tag{180}$$

$$\leq \mathbb{E}_\mu\left(\sum_{i=0, i\neq i_{\mu 1}, i\neq i_{\mu 2}}^{\infty}(M+i+1)e^{-\frac{(M+i-\mu)}{2\sigma_k^2}}\right) \tag{181}$$

$$\leq \mathbb{E}_\mu\left(\sum_{i=0}^{\infty}(M+i+1)e^{-\frac{(M+i-\mu)}{2\sigma_k^2}}\right) \tag{182}$$

$$= \mathbb{E}_\mu\left(\sum_{i=0}^{\infty}(M+i+1)e^{-\frac{(M-\mu)}{2\sigma_k^2}}\left(e^{-\frac{1}{2\sigma_k^2}}\right)^i\right) \tag{183}$$

$$= \mathbb{E}_\mu\left(\sum_{i=0}^{\infty}Me^{-\frac{(M-\mu)}{2\sigma_k^2}}\left(e^{-\frac{1}{2\sigma_k^2}}\right)^i\right) + \mathbb{E}_\mu\left(\sum_{i=0}^{\infty}(i+1)e^{-\frac{(M-\mu)}{2\sigma_k^2}}\left(e^{-\frac{1}{2\sigma_k^2}}\right)^i\right) \tag{184}$$

$$= \mathbb{E}_\mu\left(\frac{Me^{-\frac{(M-\mu)}{2\sigma_k^2}}}{1-e^{-\frac{1}{2\sigma_k^2}}}\right) + \mathbb{E}_\mu\left(\frac{Me^{-\frac{(M-\mu)}{2\sigma_k^2}}}{\left(1-e^{-\frac{1}{2\sigma_k^2}}\right)^2}\right) \tag{185}$$

$$= \left(\frac{M}{1-e^{-\frac{1}{2\sigma_k^2}}} + \frac{M}{\left(1-e^{-\frac{1}{2\sigma_k^2}}\right)^2}\right)\mathbb{E}_\mu\left(e^{-\frac{(M-\mu)}{2\sigma_k^2}}\right), \tag{186}$$

For the expectation, we have

$$\mathbb{E}_\mu\left(e^{-\frac{M-\mu}{2\sigma_k^2}}\right) = e^{-\frac{M}{2\sigma_k^2}}\mathbb{E}_\mu\left(e^{\frac{1}{2\sigma_k^2}\mu}\right). \tag{187}$$

$\mathbb{E}_\mu\left(e^{\frac{1}{2\sigma_k^2}\mu}\right)$ is finite, as it is an evaluation of the moment generating function of μ, which means that

$$\lim_{M\to\infty}\left(\frac{M}{1-e^{-\frac{1}{2\sigma_k^2}}} + \frac{M}{\left(1-e^{-\frac{1}{2\sigma_k^2}}\right)^2}\right)\mathbb{E}_\mu\left(e^{-\frac{(M-\mu)}{2\sigma_k^2}}\right) = 0. \tag{188}$$

For the two other terms on the RHS of (179)

$$\mathbb{E}_\mu\left((M+i_{\mu j}+1)e^{-\frac{(M+i_{\mu j}-\mu)^2}{2\sigma_k^2}}\right) \leq \mathbb{E}_\mu\left((M+1+\mu-M+1)e^{-\frac{(M+i_{\mu j}-\mu)^2}{2\sigma_k^2}}\right) \tag{189}$$

$$= \mathbb{E}_\mu\left((\mu+2)e^{-\frac{(M+i_{\mu j}-\mu)^2}{2\sigma_k^2}}\right) \tag{190}$$

$$\to 0 \tag{191}$$

as $M \to \infty$ by the dominated convergence theorem, as

$$(\mu+2)e^{-\frac{(M+i_{\mu j}-\mu)^2}{2\sigma_k^2}} \leq (\mu+2)e^{-\frac{0}{2\sigma_k^2}}, \tag{192}$$

and the right-hand-side has a finite expectation.

We have now shown that one part of the limit in (178) is 0. For the other part, setting $\mu' = -\mu$, and using the reasoning above with μ replaced by μ', we have

$$\lim_{M\to\infty} \mathbb{E}_{\mu}\left(\sum_{i=0}^{\infty}(M+i+1)\left(e^{-\frac{(\mu+M+i)^2}{2\sigma_k^2}}\right)\right) \tag{193}$$

$$= \lim_{M\to\infty} \mathbb{E}_{\mu'}\left(\sum_{i=0}^{\infty}(M+i+1)\left(e^{-\frac{(M+i-\mu')^2}{2\sigma_k^2}}\right)\right) \tag{194}$$

$$= 0. \tag{195}$$

so the limit in (178) is 0.

We have now shown that

$$\lim_{M\to\infty} \sup_n \mathbb{E}\big(|\bar{X}|\mathbb{I}_{|\bar{X}|>M}\big) = 0, \tag{196}$$

or, in other words, that $|\bar{X}|$ is uniformly integrable. As shown earlier, this concludes the proof that

$$\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n, D_n\right) \tag{197}$$

is uniformly integrable when $X_n^* \sim p(X_n^*|X)$.

To show that $\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n^+, D_n\right)$ is uniformly integrable, as in the proof of Lemma 3.3

$$p(Q|X, X^*) \propto p(X^*|Q)p(X|Q)p(Q), \tag{198}$$

so we can view both $p(Q|X, X_n^*)$ and $p(Q|X_n^*)$ as the posteriors for the same Bayesian inference problem with observed data X^*, and priors $p(Q|X) \propto p(X|Q)p(Q)$ and $p(Q)$, respectively. $p(Q|X)$ is Gaussian, so the uniform integrability of

$$\sqrt{n}\,\mathrm{TV}\left(\bar{Q}_n^+, D_n\right) \tag{199}$$

follows from the previous case with different values for μ_0 and σ_0^2. □

E Finite Number of Synthetic Datasets

In practice, we only have a finite number of synthetic datasets, so we must further approximate

$$p(Q|Z) \approx \int p(Q|X^*)p(X^*|Z)\,\mathrm{d}X^* \approx \frac{1}{m}\sum_{i=1}^{m} p(Q|X^* = X_i^*), \tag{200}$$

with $X_i^* \sim p(X^*|Z)$.

From the strong law of large numbers, for any n and Q,

$$\frac{1}{m}\sum_{i=1}^{m} p(Q|X^* = X_{i,n}^*) \to \int p(Q|X_n^*)p(X_n^*|Z)\,\mathrm{d}X_n^* = \bar{p}_n(Q) \tag{201}$$

almost surely as $m \to \infty$ when $X_{i,n}^* \sim p(X_n^*|Z)$.

Total variation distance is a metric, so

$$\begin{aligned} &\mathrm{TV}\left(\frac{1}{m}\sum_{i=1}^{m} p(Q|X_{i,n}^*), p(Q|Z)\right) \\ &\leq \mathrm{TV}\left(\frac{1}{m}\sum_{i=1}^{m} p(Q|X_{i,n}^*), \bar{p}_n(Q)\right) + \mathrm{TV}\left(\bar{p}_n(Q), p(Q|Z)\right). \end{aligned} \tag{202}$$

Theorem 3.4 gives

$$\lim_{n\to\infty} \mathrm{TV}\left(\bar{p}_n(Q), p(Q|Z)\right) = 0. \tag{203}$$

(201) implies (van der Vaart 1998, Corollary 2.30)

$$\lim_{m\to\infty} \mathrm{TV}\left(\frac{1}{m}\sum_{i=1}^{m} p(Q|X_{i,n}^*), \bar{p}_n(Q)\right) = 0 \tag{204}$$

for almost all $X^*_{i,n}$, so

$$\lim_{n\to\infty}\lim_{m\to\infty} \mathrm{TV}\left(\frac{1}{m}\sum_{j=1}^{m} p(Q|X^*_{j,n}), p(Q|Z)\right) = 0 \tag{205}$$

almost surely when $X^*_{j,n} \sim p(X^*_n|Z)$.

Based on the experiment in Figure S1, it looks like

$$\lim_{m\to\infty}\lim_{n\to\infty} \mathrm{TV}\left(\frac{1}{m}\sum_{j=1}^{m} p(Q|X^*_{j,n}), p(Q|Z)\right) \neq 0 \tag{206}$$

because the distributions $p(Q|X^*_{i,n})$ become narrower as n increases, so a fixed number of them is not enough to cover $p(Q|Z)$.

F Relation to Missing Data Imputation

In the missing data setting, only a part X_{obs} of the complete dataset X is observed, while a part X_{mis} is missing (Rubin 1987). To facilitate downstream analysis, the missing data is imputed by sampling $X_{mis} \sim p(X_{mis}|X_{obs}, I_I)$. Analogously with synthetic data, I_I represents the imputer's background knowledge.

Like with synthetic data, we have the decomposition (Gelman et al. 2014)

$$p(Q|X_{obs}, I_A) = \int p(Q|X_{obs}, X_{mis}, I_A)p(X_{mis}|X_{obs}, I_A)\,\mathrm{d}X_{mis}. \tag{207}$$

If the analyst's and imputer's models are congenial in the sense that

$$p(Q|X_{obs}, I_A) = p(Q|X_{obs}, I_I) \tag{208}$$

and

$$p(Q|X, I_A) = p(Q|X, I_I) \tag{209}$$

for any complete dataset X, then

$$\begin{aligned} p(Q|X_{obs}, I_A) = p(Q|X_{obs}, I_I) &= \int p(Q|X_{obs}, X_{mis}, I_I)p(X_{mis}|X_{obs}, I_I)\,\mathrm{d}X_{mis} \\ &= \int p(Q|X_{obs}, X_{mis}, I_A)p(X_{mis}|X_{obs}, I_I)\,\mathrm{d}X_{mis}, \end{aligned} \tag{210}$$

so sampling $p(Q|X_{obs}, I_A)$ can be done by sampling $X_{mis} \sim p(X_{mis}|X_{obs}, I_I)$ multiple times, sampling $p(Q|X_{obs}, X_{mis}, I_A)$ for each X_{mis}, and combining the samples. Unlike with synthetic data, where sampling $p(Q|X, X^*, I_A)$ would require the original data and defeat the purpose of using synthetic data, in the missing data setting, sampling $p(Q|X_{obs}, X_{mis}, I_A)$ is simply the analysis for a complete dataset, so generating large imputed datasets is not required.