

Approximate Co-Sufficient Sampling for Goodness-of-fit Tests and Synthetic Data

Jordan Awan*

Department of Statistics, Purdue University

Zhanrui Cai†

Department of Statistics, The Pennsylvania State University

Abstract

Co-sufficient sampling refers to resampling the data conditional on a sufficient statistic, a useful technique for statistical problems such as goodness-of-fit tests, model selection, and confidence interval construction; it is also a powerful tool to generate synthetic data which limits the disclosure risk of sensitive data. However, sampling from such conditional distributions is both technically and computationally challenging, and is inapplicable in models without low-dimensional sufficient statistics.

We study an indirect inference approach to approximate co-sufficient sampling, which only requires an efficient statistic rather than a sufficient statistic. Given an efficient estimator, we prove that the expected KL divergence goes to zero between the true conditional distribution and the resulting approximate distribution. We also propose a one-step approximate solution to the optimization problem that preserves the original estimator with an error of $o_p(n^{-1/2})$, which suffices for asymptotic optimality. The one-step method is easily implemented, highly computationally efficient, and applicable to a wide variety of models, only requiring the ability to sample from the model and compute an efficient statistic. We implement our methods via simulations to tackle problems in synthetic data, hypothesis testing, and differential privacy.

Keywords: Asymptotic Sufficiency, Indirect Inference, Statistical Disclosure Control, Differential Privacy, simulation-based inference, bootstrap

*The first author gratefully acknowledges support from NSF SES-1853209 and NSF SES-153443 awarded to The Pennsylvania State University

†

1 Introduction

Suppose we have a sample $\underline{X} = \{X_1, \dots, X_n\} \stackrel{\text{i.i.d.}}{\sim} f_\theta$ from a parametric model, with unknown parameter $\theta \in \Theta$. If there exists a sufficient statistic $S(\underline{X})$ for this model, then *co-sufficient* sampling (CSS) refers to resampling the data, conditional on $S(\underline{X})$. We will write this distribution as $f_\theta^n(\underline{x} \mid S(\underline{X}))$, which does not depend on θ due to sufficiency. The term co-sufficient sampling was introduced in the context of goodness-of-fit tests [Lockhart et al., 2007]: by conditioning on a sufficient statistic for the null model, the resulting distribution of the test statistic can be used to derive optimal hypothesis tests with calibrated type I error rates. CSS has also been used to produce synthetic datasets, with the goal of limiting disclosure risk while maintaining data utility.

Goodness-of-fit Tests: We consider a *goodness-of-fit test* to be any hypothesis where the null hypothesis states that the distribution of the sample follows a parametric model, which we may write as follows: $H_0 : \underline{X} \sim f_\theta \mid \theta \in \Theta_0$, and where the alternative hypothesis can be of any form. While we use the term goodness-of-fit, there are in fact many types of hypothesis tests that fit within this framework: classical goodness-of-fit tests propose that the null model is a specified parametric model, model selection tests may propose that the null model is a subset of a full model, and conditional independence tests set the null model where the two random variables are independent given a third variable. Many classical results of uniformly most powerful unbiased tests show that when the null model is composite, an optimal test is based on the distribution of the test statistic conditional on a sufficient statistic for the null model [Schervish, 2012, Chapter 4].

However, working directly with the co-sufficient distribution may be difficult. CSS allows for an approximation to the true conditional sampling distribution and provides Monte Carlo p -values. Some recent applications of CSS to hypothesis tests are Agresti et al. [1992],

Engen and Lillegård [1997], Agresti [2001], O’Reilly and Gracia-Medrano [2006], Lockhart et al. [2007], Lindqvist and Rannestad [2011], Stephens [2012], Broniatowski and Caron [2019], Santos and Filho [2019]. Each of these works are limited to the case of a sufficient statistic, and also often restrict their attention to a particular models.

Barber and Janson [2020] expanded the applicability of CSS to models where low-dimensional sufficient statistics do not exist, using methods of approximate co-sufficient sampling (ACSS) based on the asymptotic sufficiency of efficient estimators [Le Cam et al., 1956]. However, implementing their methods requires computationally expensive Markov chain Monte Carlo (MCMC) techniques, besides the expensive run-time of these methods in the best case scenario, MCMC methods also require careful tuning and diagnostics to ensure convergence of the chains. Furthermore, MCMC methods require the ability to evaluate the likelihood, which may not be possible in models with intractable likelihood functions.

Another numerical approach to goodness-of-fit testing problems is the parametric bootstrap. With this approach, the model is fit under the null hypothesis, and samples are drawn from the fitted model to approximate the sampling distribution of a test statistic. While this approach is very popular due to its simplicity, Robins et al. [2000] show that in many cases tests based on the parametric bootstrap have overly conservative type I error rates, resulting in limited power. Barber and Janson [2020] on the other hand, showed that ACSS guarantees accurate type I error rates, resulting in more powerful tests.

Synthetic Data: With advances in modern technology, government and other research agencies collect massive amounts of data from individual respondents. These data are valuable for scientific progress and policy research, but they also come with increased privacy risk [Lane et al., 2014]. To publish useful information while preserving confidentiality of sensitive information, numerous methods of generating *synthetic data* have been proposed (see Hundepool et al. [2012, Chapter 3] for a survey). A major goal of synthetic data

is to communicate the data structure to non-technical audiences such as sociologists, demographers, and public policy experts who may not have a theoretic statistics background, but still desire to study, analyze, and visualize the data [Hundepool et al., 2012, p99]. For example, many area experts rely on synthetic data products published by the US Census.

Recently, there has been a push to provide formal privacy guarantees for synthetic data, and differential privacy (DP) [Dwork et al., 2006b] is the leading framework to develop formally private methods. For example, the 2020 decennial census will be released using differential privacy [Abowd, 2018, Dwork, 2019]. Methods which satisfy DP require the introduction of additional randomness, beyond sampling in order to obscure the effect of one individual on the output. Intuitively, DP ensures plausible deniability for those participating in the dataset. As the literature on DP has developed, there are now many privacy tools to perform many standard statistical tasks such as hypothesis testing [Awan and Slavković, 2018, Awan and Slavković, 2020, Gaboardi et al., 2016, Acharya et al., 2018, Canonne et al., 2019], empirical risk minimization [Chaudhuri et al., 2011, Reimherr and Awan, 2019, Bassily et al., 2014, Kifer et al., 2012], regression [Zhang et al., 2012, Fang et al., 2019, Chaudhuri and Monteleoni, 2008], and posterior inference [Williams and McSherry, 2010, Karwa et al., 2016] to name just a few areas. There have also been several works developing methods for DP synthetic data, which are included in the synthetic data literature below.

The two most common approaches to synthetic data are either the parametric bootstrap, which samples from a fitted model [Liew et al., 1985, Machanavajjhala et al., 2008, Zhang et al., 2017], or co-sufficient sampling, which produces samples conditional on a low-dimensional summary statistic [Muralidhar and Sarathy, 2003, Burridge, 2003, Mateo-Sanz et al., 2004, Ting et al., 2005, Chen et al., 2006, Karwa and Slavković, 2013, Slavković and Lee, 2010, Karwa and Slavković, 2012]. There are also Bayesian methods of producing synthetic data [McClure and Reiter, 2012, Liu, 2016], but the properties of these methods

are often similar to one of the above paradigms. Of these works, the following satisfy the formal privacy guarantee of DP [Machanavajjhala et al., 2008, Zhang et al., 2017, Karwa and Slavković, 2012, McClure and Reiter, 2012, Liu, 2016].

The parametric bootstrap approach to synthetic data is limited by inferior asymptotics, as we show in Section 3.1, whereas CSS approaches synthetic data are limited by being either specific to particular distributions (e.g., Gaussian data [Burridge, 2003, Mateo-Sanz et al., 2004, Ting et al., 2005]) or are computationally burden-some (e.g., MCMC methods [Barber and Janson, 2020, Karwa and Slavković, 2013, Slavković and Lee, 2010]).

The synthetic data literature distinguishes between *partially synthetic data* and *fully synthetic data*: in partially synthetic data, certain attributes or statistics are preserved from the original dataset and the other attributes of the dataset are synthesized; for fully synthetic data, all attributes are synthesized and there are no attributes or statistics in common with the original dataset. Typically, the parametric bootstrap method of synthetic data results in fully synthetic data, whereas CSS results in partially synthetic data, where the sufficient statistic is held fixed. It is also possible to use CSS to produce fully synthetic data, by incorporating a differentially private estimate of the summary statistic. However, this is not possible with traditional CSS methods, as DP methods typically result in distributions without low-dimensional sufficient statistics. However, it is possible to produce asymptotically efficient estimators under differential privacy [Smith, 2011], for which ACSS techniques can be applied to produce DP fully synthetic data.

Indirect Inference: The approach to ACSS that we study in this paper is based on techniques from indirect inference. Indirect inference methods are tailored to models which are easy to sample from, but which may have intractable likelihood functions [Gourieroux et al., 1993]. Indirect inference is a simulation-based method of inference, which leverages the ability to fix the *seed* that generates a sample. Recall that all random variables are

measurable functions from a seed space to the space of possible observations. In real data, we only have access to the observations, and do not know the initial seed or the measurable function, which depends on the unknown parameter. However, in simulations, we can fix the seed and vary the parameter value, resulting in different observations for the same seed.

Indirect inference methods are based on the assumption that for well-behaved models, the resulting observations vary smoothly as the parameter value is varied. Determining the parameter value which makes the simulated observations most similar to the real-data observations gives a parameter estimate with certain properties such as unbiasedness or robustness [Gourieroux et al., 1993, Heggland and Frigessi, 2004, Genton and Ronchetti, 2003, Guerrier et al., 2019]. See Drovandi et al. [2011] for an exploration of how indirect inference is related to approximate Bayesian computing.

Co-sufficient Sampling via Indirect Inference: There have been some works that develop CSS methods using the techniques of indirect inference, which we refer to as indirect co-sufficient sampling (ICSS). However these works are not connected to the indirect inference literature and use different terminology such as conditional Monte Carlo [Lindqvist and Taraldsen, 2005, 2007, 2013], conditional parametric bootstrap [Lillegard and Engen, 1999], and stochastic simulations conditioned on sufficient statistics [Engen and Lillegård, 1997]. Besides the application of co-sufficient sampling, these methods have also been applied to produce confidence intervals [Lillegard and Engen, 1999], and fiducial inference Taraldsen and Lindqvist [2018].

The ICSS method used in these papers works as follows: given the original sample \underline{X} with sufficient statistic $S(\underline{X})$, sample a new seed and then determine the parameter value that produces a new sample \underline{Y} with sufficient statistic $S(\underline{Y}) = S(\underline{X})$. Lindqvist and Taraldsen [2005] show that under certain assumptions, this process does indeed result in co-sufficient samples. However, these previous works are greatly limited in their applicability.

They assume that the distribution of the seed variable is known and easy to sample from and that the transformation from seed space to the observation space is also known. However for non-trivial models, this is often not the case. These works also only study the case of exact CSS, which requires a low-dimensional sufficient statistic, which is also often not the case in more complex models. Finally, they assume that it is possible to find a parameter to solve $S(\underline{Y}) = S(\underline{X})$ exactly, however it may be non-trivial to solve this equation in θ , and a solution may not even exist. Due to these limitations, these CSS methods are generally limited to simple models such as Gaussian [Lindqvist and Taraldsen, 2005, Taraldsen and Lindqvist, 2018], exponential [Engen and Lillegård, 1997], truncated exponential [Lindqvist and Taraldsen, 2007], gamma [Engen and Lillegård, 1997, Taraldsen and Lindqvist, 2018], inverse Gaussian [Lindqvist and Taraldsen, 2007], and Bernoulli [Lindqvist and Taraldsen, 2005]. The most complex application of ICSS is to Poisson processes [Lindqvist and Taraldsen, 2013].

Our Contributions: In this paper, we consider the ICSS algorithm and expand its applicability as follows: 1) We show that rather than requiring the exact distribution of the seed random variable, we can instead use numerical seeds which are applicable to many more models. The use of numerical seeds is common in the indirect inference literature, but has not previously been applied to CSS. 2) Based on the principle of asymptotic sufficiency [Le Cam et al., 1956], we show that ICSS can be used for ACSS when there is an efficient estimator of the parameter, even if there is no low-dimensional sufficient statistic. We prove that when using an efficient estimator, the expected KL divergence between the distribution of the ICSS samples and the true co-sufficient sampling distribution goes to zero as the sample size increases. 3) We propose a simple one-step approximate solution to the ICSS objective, which we call One Approximate-Step Indirect co-Sufficient (OASIS) sampling. We show that OASIS sampling preserves an efficient statistic with error $o_p(n^{-1/2})$, which is

sufficient for optimal asymptotics. OASIS sampling is incredibly simple to implement and highly computationally efficient, with running time comparable to the parametric bootstrap. We apply our results to problems in goodness-of-fit testing, synthetic data generation, and differential privacy through several simulations, illustrating the above asymptotic results and demonstrating high finite sample utility as well.

From another perspective, our methodology can be viewed as an indirect inference approach to the ACSS methods of Barber and Janson [2020]. By working with an optimization-based version of the sampling problem, we greatly increase the tractability of the sampling problem, improving the ease of implementation as well as computational cost. In particular, the OASIS algorithm only requires the ability to sample from the model, and evaluate an efficient statistic for the model. As OASIS does not require any evaluations of the likelihood function, it is applicable to models with intractable likelihood functions, whereas the MCMC approach of Barber and Janson [2020] is inapplicable.

Because of the increased computational efficiency, and applicability to a wide-variety of models, our results allow for the powerful tools of ACSS to be applied to to produce hypothesis tests with reliable type I errors, as well as synthetic data with high utility. We also highlight applications of our methods to differential privacy: our results can be used to produce differentially private synthetic data by incorporating a differentially private estimator, and we develop valid hypothesis tests on differentially private releases, for which standard asymptotic tests fail in practical sample sizes.

Finally, we note that methods for CSS have been studied in various subfields, but that these fields tend to be disconnected. For instance, the literature on ICSS seems unfamiliar with the connection between their methods and indirect inference, and the works on CSS for synthetic data do not reference the other literature on CSS. It is also difficult for these fields to make these connections as they often use different terminology for the same concepts.

We hope that this paper can help to bridge the gap between these sub-disciplines allowing for each subfield to benefit from the unique insights and applications of the other areas.

Organization: In Section 2, we review some background and set the notation for the paper. We recall some basic measure theory necessary to discuss indirect inference, and then introduce the ICSS algorithm in Section 3. In Section 3.1 we discuss the limitations of the parametric bootstrap, an alternative to CSS, for both synthetic data and hypothesis testing, and prove that the parametric bootstrap results in *inconsistent* synthetic data. We give our theoretical results in Section 4.1 measuring the closeness between the distribution of ICSS and the true conditional distribution. In particular, we prove that when using an efficient statistic, the expected KL divergence between the ICSS sampling distribution and the true conditional distribution goes to zero. In Section 4.2 we propose the OASIS sampling algorithm, which gives an approximate solution to the indirect inference optimization problem, preserving the efficient statistic with error $o_p(n^{-1/2})$, which is sufficient for asymptotic optimality. This is in contrast to the exact ICSS methods which do not specify how to solve the indirect inference optimization problem, and where an exact solution often does not even exist. In section 5, we demonstrate the performance of the OASIS algorithm on several simulations. In Sections 5.1 and 5.2, we produce partially synthetic data from the Burr distribution as well as a log-linear model. In Section 5.3, we produce differentially private fully synthetic data for the beta distribution. In Section 5.4 we use our methods to derive an asymptotically valid test for the difference of proportions under differential privacy. In Section 5.5, we compare the running time of OASIS compared to the MCMC approach of Barber and Janson [2020] when implemented on a linear regression problem. We end in Section 6 with some discussion. Proofs and technical lemmas are deferred to Section 7.2.

2 Background and notation

In this section, we review some background and notations that we use throughout the paper.

For a parametric random variable, we write $X \sim f_\theta$ to indicate that X has probability density function (pdf) f_θ . To indicate that a sequence of random variables from the model f_θ are independent and identically distributed (i.i.d.), we write $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f_\theta$. Depending on the context, we use the following notation interchangeably: $\underline{X} = (X_i)_{i=1}^n = (X_1, \dots, X_n)^\top$. We write $\ell(\theta | \underline{x}) = \sum_{i=1}^n \log f_\theta(x_i)$ to represent the log-likelihood function. Given a vector $\underline{\omega} = (\omega_1, \dots, \omega_n)^\top$ and a function F defined on each ω_i , we write $F(\underline{\omega}) := (F(\omega_1), \dots, F(\omega_n))^\top$.

Let A be a random vector, A_n be a sequence of random vectors, and r_n be a positive numerical sequence. We write $A_n \xrightarrow{d} A$ to denote that A_n converges in distribution to A . We write $A_n = o_p(r_n)$ to denote that $A_n/r_n \xrightarrow{d} 0$. We write $A_n = O_p(r_n)$ to denote that A_n/r_n is bounded in probability.

For multivariate derivatives, we will overload the $\frac{d}{d\theta}$ operator as follows. For a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$, we write $\frac{d}{d\theta} f(\theta)$ to denote the $p \times 1$ vector of partial derivatives $(\frac{\partial}{\partial \theta_j} f(\theta))_{j=1}^p$. For a function $g : \mathbb{R}^p \rightarrow \mathbb{R}^q$, we write $\frac{d}{d\theta} g(\theta)$ to denote the $p \times q$ matrix $(\frac{\partial}{\partial \theta_j} g_k(\theta))_{j,k=1}^{p,q}$.

For $X \sim f_\theta$, we denote the *score function* as $S(\theta, x) = \frac{d}{d\theta} \log f_\theta(x)$, and the *Fisher information* as $I(\theta) = \mathbb{E}_\theta [S(\theta, X)S^\top(\theta, X)]$. An estimator $\hat{\theta} : \mathcal{X}^n \rightarrow \Theta$ is *efficient* if for $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f_\theta$, we have $\sqrt{n}(\hat{\theta}(\underline{X}) - \theta) \xrightarrow{d} N(0, I^{-1}(\theta))$.

Let f and g be densities with respect to a common base measure μ . The *total variation distance* between f and g is $\text{TV}(f||g) = \frac{1}{2} \int_x |f(x) - g(x)| d\mu(x)$; the *Hellinger distance* between f and g is $H(f||g)^2 = \frac{1}{2} \int_x \left(\sqrt{f(x)} - \sqrt{g(x)} \right)^2 d\mu(x)$; the KL divergence from g to f is $\text{KL}(f||g) = \mathbb{E}_{x \sim f} \log(f(x)/g(x))$. For random variables $X \sim f$ and $Y \sim g$, we write $\text{TV}(X||Y)$, $H(X||Y)$ and $\text{KL}(X||Y)$ to mean $\text{TV}(f||g)$, $H(f||g)$ and $\text{KL}(f||g)$, respectively.

3 Methodology

In this section, we recall some basic measure theory about random variables in (R0) in order to formalize the approach of indirect inference and the ICSS algorithm. After properly introducing the ICSS algorithm, we discuss how numerical seeds can be used in computational platforms such as R to simulate the measure theory of (R0). We also give a pseudocode implementation of the ICSS algorithm using numerical seeds in Algorithm 1.

Consider sample $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f_\theta$ for a parametric model, where θ is unknown, and let $\hat{\theta}$ be an estimator for θ . Our goal is to produce a new sample Y_1, \dots, Y_n whose distribution is approximately $f_\theta^n(\underline{y} \mid \hat{\theta}(\underline{y}) = \hat{\theta}(\underline{X}))$. In Section 4, we focus on the case where $\hat{\theta}(\cdot)$ is an efficient estimator of θ . Recall that when $\hat{\theta}(\underline{X})$ is a sufficient statistic, that the above conditional distribution does not depend on θ ; in this case, it is at least conceivably possible to sample exactly from this distribution. If $\hat{\theta}(\underline{X})$ is an efficient estimator, then it is asymptotically sufficient [Le Cam et al., 1956], and we prove in Section 4.1 that ICSS can provide an asymptotic approximation to the desired distribution.

Barber and Janson [2020] produce approximate co-sufficient samples by substituting $\hat{\theta}(\underline{X})$ for θ . The density function of $f_{\hat{\theta}(\underline{X})}^n(\underline{y} \mid \hat{\theta}(\underline{y}) = \hat{\theta}(\underline{X}))$ can be expressed up to the normalizing constant, and then MCMC methods can be used to generate samples. there is still a computational challenge of sampling from this distribution, often involving MCMC techniques. Furthermore, for models with latent variables, the likelihood may be intractable rendering MCMC methods inapplicable.

In this paper, we study an indirect inference based method of sampling such distributions, which replaces the sampling problem with an optimization problem. This makes the problem much more computationally efficient and easier to implement, requiring only the ability to sample from the model and compute an efficient estimator for the model parameters.

Indirect inference is based on the idea that by first fixing the *seed*, we can study the deterministic process of the random variable, by modifying the parameter values. We give the formal definition of the seed in terms of measure theory in the following assumption.

Assumption 3.1. (R0) Let (Ω, \mathcal{F}, P) be a probability space of the *seed* ω . Let $X_\theta : \Omega \rightarrow \mathcal{X}$ be a measurable function, where $(\mathcal{X}, \mathcal{G})$ is a measurable space and $\theta \in \Theta \subset \mathbb{R}^p$. We assume that there exists a measure μ on $(\mathcal{X}, \mathcal{G})$ which dominates PX_θ^{-1} for all $\theta \in \Theta$. Then there exist densities $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{\geq 0}$ such that $\int_A dPX_\theta^{-1} = \int_A f_\theta d\mu$ for all $A \in \mathcal{G}$.

Assumption (R0) tells us that to produce a sample $X \sim f_\theta$, we can first sample the seed $\omega \sim P$ and then transform the seed into $Y := X_\theta(\omega)$. This procedure results in a sample equal in distribution: $X \stackrel{d}{=} Y$. Next, assuming some regularity conditions on f_θ (formalized in Assumption 4.1), we have that the mapping $X_\theta(\cdot)$ is smooth as θ is varies.

The ICSS algorithm works as follows: given $\hat{\theta}(\underline{X})$ computed from the original dataset, we first sample the seeds $\omega_1, \dots, \omega_n \stackrel{\text{i.i.d.}}{\sim} P$, and then while holding $\omega_1, \dots, \omega_n$ fixed, solve for the value θ^* which satisfies:

$$\hat{\theta}(X_{\theta^*}(\underline{\omega})) = \hat{\theta}(\underline{X}), \tag{1}$$

thereby ensuring that the new sample has the same value of $\hat{\theta}$ as the original sample \underline{X} . Finally, we can produce our synthetic data $Y_i = X_{\theta^*}(\omega_i)$, which can be used for limited data disclosure or for statistical inference.

When implementing the ICSS algorithm in computing software, the use of numerical seeds is incredibly useful. For example in R, the command `set.seed` fixes the initial numerical seed, which is used to generate pseudo-random numbers. After setting the seed, sampling from a probability model with parameter θ is a deterministic function of the seed. Importantly, under the conditions of Assumption 4.1, for a fixed seed the sample

varies smoothly as θ is varied. See Algorithm 1 for R pseudo-code demonstrating how to use numerical seeds to solve Equation (1). In Algorithm 2, we give R pseudo-code for the one-step approximate solution to (1), which we develop in Section 4.2.

Algorithm 1: Indirect approximate co-sufficient sampling Pseudo-code in R

Input: Seed ω , efficient estimator $\hat{\theta}(\underline{X})$, function `theta_hat(y)` to compute $\hat{\theta}(y)$, function `rsample(theta)` to sample n i.i.d. samples from f_θ .

```

1 Function objective(theta):
2   |   set.seed(omega)
3   |   Y = rsample(theta)
4   |   theta_hat(Y) = theta_hat(Y)
5   |   return ||theta_hat(Y) - theta_hat(X)||^2
6 theta* = argMinimum(objective)
7 set.seed(omega)
8 Y = rsample(theta*)
Output: Y1, ..., Yn

```

Remark 3.2 (Related Methods). While using different terminology and notation than we do here, similar methods to Equation (1) have been proposed previously in the literature [Lindqvist and Taraldsen, 2005, 2007, 2013, Lillegard and Engen, 1999, Engen and Lillegård, 1997]. One key difference between these past works and the present paper is that they are limited to the case of sufficient statistics, whereas we show in Theorem 4.5 that the conditional distribution can be approximated even when $\hat{\theta}$ is an efficient statistic. The previous works also assume that the distribution P is known and easily sampled, whereas in Algorithm 1 we show that numerical seeds can be more easily used. Lastly, the previous methods do not discuss algorithmic methods of solving Equation (1). We show in Example 4.6 that an exact solution may not exist, and propose in Algorithm 2 an efficient one-step algorithm that produces an approximate solution to Equation (1).

We end this section with two examples illustrating the ICSS algorithm in simple settings such as location-scale families and continuous real-valued random variables.

Example 3.3 (Location-Scale Example). In the case where f_θ is a location-scale family, it is easy to illustrate the seed, the function X_θ , and how to compute the parameter value θ^* . Let $\theta = (m, s)$ be the parameter vector, where m is the location parameter and s is the scale parameter. Let $\theta_0 = (0, 1)$ be a default parameter. To produce a sample from $f_{\theta=(m,s)}$, we can draw a seed $\omega \sim f_{\theta_0}$ and apply the function $X_{\theta=(m,s)}(\omega) = \omega \cdot s + m$. Given a sample $\underline{x} = (x_1, \dots, x_n)^\top$, let $\hat{m}(\underline{x})$ and $\hat{s}(\underline{x})$ be location and scale estimators, respectively, which satisfy $\hat{m}(a\underline{x}+b) = a\hat{m}(\underline{x})+b$ and $\hat{s}(a\underline{x}+b) = a\hat{s}(\underline{x})$, where $(a\underline{x}+b) := (ax_1+b, \dots, ax_n+b)^\top$.

Suppose we are given $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f_\theta$ and $\omega_1, \dots, \omega_n \stackrel{\text{i.i.d.}}{\sim} f_{\theta_0}$. It is easy to verify that $\theta^* = (\hat{m}(\underline{X}) - \hat{m}(\underline{\omega})\hat{s}(\underline{X})/\hat{s}(\underline{\omega}), \hat{s}(\underline{X})/\hat{s}(\underline{\omega}))$ is the solution to Equation (1). So, we produce the sample $Y_i = X_{\theta^*}(\omega_i)$, which satisfies $\hat{m}(\underline{Y}) = \hat{m}(\underline{X})$ and $\hat{s}(\underline{Y}) = \hat{s}(\underline{X})$.

Example 3.4 (Continuous Real-Valued Example). For continuous real-valued random variables, we can be more explicit about the “seeds.” Recall that for $U \sim U(0, 1)$, $F_\theta^{-1}(U) \sim f_\theta$ where $F_\theta^{-1}(\cdot)$ is the quantile function. So in this case, the distribution P can be taken as $U(0, 1)$, and $X_\theta(\cdot) = F_\theta^{-1}(\cdot)$.

3.1 Limitations of Parametric Bootstrap

The parametric bootstrap is a popular alternative to co-sufficient sampling, which has been popularly used both for synthetic data as well as for hypothesis testing. However, the parametric bootstrap has considerable weaknesses compared to co-sufficient sampling, in that it results in significantly worse approximations to the true sampling distribution. In this section, we prove that the parametric bootstrap results in “inconsistent” synthetic data, where we show that the total variation distance between the true distribution and the

parametric bootstrap approximation does not go to zero as $n \rightarrow \infty$. We also recall other limitations of the parametric bootstrap as related to hypothesis testing.

From the perspective of synthetic data, the goal is to produce a new dataset \underline{Y} such that $\underline{Y} \stackrel{d}{\approx} \underline{X}$, where the approximation could be measured by a metric or divergence. At minimum, we may expect that the distribution of \underline{Y} approaches the distribution of \underline{X} as the sample size n grows. We begin with an example that show that the parametric bootstrap results in suboptimal asymptotics, calling into question the appropriateness of the parametric bootstrap for the generation of synthetic.

Example 3.5. Suppose that $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} N(\mu, 1)$. We estimate $\hat{\mu}(\underline{X}) = n^{-1} \sum_{i=1}^n X_i$ and draw $Z_1, \dots, Z_n | \hat{\mu}(\underline{X}) \stackrel{i.i.d.}{\sim} N(\hat{\mu}(\underline{X}), 1)$. We can compute $\text{Var}(\hat{\mu}(\underline{X})) = n^{-1}$, whereas $\text{Var}(\hat{\mu}(\underline{Z})) = 2n^{-1}$. By using the synthetic data \underline{Z} , we have lost half of the effective sample size. We can also derive $(Z_1, \dots, Z_n)^\top \sim N(\mu \underline{1}, I_n + n^{-1} \underline{1} \underline{1}^\top)$, where $\underline{1} = (1, \dots, 1)^\top$ is a vector of length n . It is easily calculated that the Hellinger distance between the samples \underline{X} and \underline{Z} is $H(\underline{X} || \underline{Z}) = 1 - 2^{1/4} / (3/2)^{1/2} \approx .029$, by using the formula in the proof of Proposition 3.6. This demonstrates that the samples \underline{X} and \underline{Z} are fundamentally different, and that their distributions do not approach each other as n grows.

While a simple example, the implications of Example 3.5 are quite general. In fact, when $\hat{\theta}(\cdot)$ is an efficient estimator, we show that a similar result holds in Proposition 3.6: $\hat{\theta}(\underline{Z})$ is an inefficient estimator for θ , and that the distribution of \underline{Z} is “inconsistent” in that the distance between the distributions of \underline{Z} and \underline{X} does not go to zero as $n \rightarrow \infty$.

Proposition 3.6. *Let $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} f_\theta$, let $\hat{\theta}(\underline{X})$ be an efficient estimator for θ . Sample $Z_1, \dots, Z_n | \hat{\theta}(\underline{X}) \stackrel{i.i.d.}{\sim} f_{\hat{\theta}(\underline{X})}$. Then*

1. $\sqrt{n}(\hat{\theta}(\underline{Z}) - \theta) \xrightarrow{d} N(0, 2I^{-1}(\theta))$, whereas $\sqrt{n}(\hat{\theta}(\underline{X}) - \theta) \xrightarrow{d} N(0, I^{-1}(\theta))$, and

2. $\text{TV}(X_1, \dots, X_n || Z_1, \dots, Z_n)$ does not converge to zero as $n \rightarrow \infty$.

Proposition 3.6 demonstrates that the parametric bootstrap is not ideal for the generation of synthetic data. We show in Section 4 that our approach to approximate so-sufficient sampling offers improved asymptotics over the parametric bootstrap.

The parametric bootstrap also has limitations when applied to hypothesis testing problems. Indeed, Robins et al. [2000] show that if the test statistic $T(X)$ is asymptotically normal with mean $a(\theta)$, then the parametric bootstrap p -values are asymptotically conservative if $a(\theta)$ is not a constant in θ ; this results in tests with low power. On the other hand, if $a(\theta)$ is constant in θ , then the p -values are asymptotically uniform under the null. See Boos et al. [2003] for an intuitive discussion of this phenomenon.

For complex models, it may be difficult to construct a test statistic with the necessary properties to ensure that the p -values are appropriately calibrated. On the other hand, Barber and Janson [2020] showed that using approximate co-sufficient sampling methods similar to those in this paper guarantee asymptotically uniform p -values under the null hypothesis, no matter which test statistic is used. We show in Section 4.1 that our method also results in approximate co-sufficient samples, and demonstrate in Section 5.4 that for a DP testing problem, samples from our method have calibrated p -values whereas samples from the parametric bootstrap result in overly conservative p -values and limited power.

4 Main results

In this section, we prove the main results related to the indirect inference approach to approximate co-sufficient sampling. The results fit into two categories: in section 4.1, we assume an exact solution to Equation (1) and prove distributional results of the resulting ICSS samples, when using an efficient estimator. In Section 4.2, we address the practical

problem of finding a numerical solution by proposing a one-step approximate solution to (1), which only requires fitting the model, sampling from the model, and computing an efficient statistic. We show that whether there exists an exact solution to (1) or not, the one-step algorithm is guaranteed to preserve the original efficient statistic $\hat{\theta}(\underline{X})$ up to $o_p(n^{-1/2})$, thereby maintaining efficiency.

4.1 Distributional results

In this section, we prove results comparing the distribution of the ICSS synthetic data when using an efficient estimator and the original sample. We begin by introducing additional assumptions on the data generating distribution, which ensure that an efficient estimator exists, as well as an assumption ensuring that the efficient estimator is well-behaved. Next, we prove a basic result expressing the phenomenon of asymptotic sufficiency [Le Cam et al., 1956] in terms of KL divergence in Theorem 4.2. This result essentially shows that the KL divergence between ACSS samples and the true data distribution goes to zero, so long as the plug-in value satisfies $\theta_n = \theta + O(n^{-1/2})$. After this preliminary work, we prove two lemmas about the ICSS algorithm in the case of an efficient estimator, expressing the conditional distribution of the ICSS samples as well as the asymptotic performance of θ^* , the solution to Equation (1). Finally, we combine our lemmas with Theorem 4.2 to show that the expected KL divergence between the approximate samples via ICSS and the true CSS distribution goes to zero as the sample size goes to infinity.

Assumptions (R1)-(R3) below consists of standard regularity conditions to ensure that an efficient estimator exists, see for example Serfling [1980]. These assumptions will be required to establish properties of the conditional distribution $f_{\theta}^n(x_1, \dots, x_n \mid \hat{\theta}(\underline{x}))$, and the efficiency will imply asymptotic sufficiency. While (R4) is not a standard regularity condition, it ensures that the efficient estimator is well-behaved.

Assumption 4.1. (R1) Let $\theta_0 \in \Theta \subset \mathbb{R}^p$ be the true parameter, where Θ is a compact set.

Assume there exists an open ball $B(\theta_0) \subset \Theta$ about θ_0 , the model f_θ is identifiable, and that the set $\{x \in \mathcal{X} \mid f_\theta(x) > 0\}$ does not depend on θ .

(R2) The pdf $f_\theta(x)$ has three derivatives in θ for all x and there exist functions $g_i(x)$, $g_{ij}(x)$, $g_{ijk}(x)$ for $i, j, k = 1, \dots, p$ such that for all x and all $\theta \in B(\theta_0)$,

$$\left| \frac{\partial f_\theta(x)}{\partial \theta_i} \right| \leq g_i(x), \quad \left| \frac{\partial^2 f_\theta(x)}{\partial \theta_i \partial \theta_j} \right| \leq g_{ij}(x), \quad \left| \frac{\partial^3 f_\theta(x)}{\partial \theta_i \partial \theta_j \partial \theta_k} \right| \leq g_{ijk}(x).$$

We further assume that each g satisfies $\int g(x) dx < \infty$ and $\mathbb{E}_\theta g_{ijk}(X) < \infty$ for $\theta \in B(\theta_0)$.

(R3) The Fisher information matrix $I(\theta) = \mathbb{E}_\theta[(\frac{d}{d\theta} \log f_\theta(X))(\frac{d}{d\theta} \log f_\theta(X))^\top]$ consists of finite entries, and is positive definite for all $\theta \in B(\theta_0)$.

(R4) Let $\hat{\theta}(\underline{X})$ be an efficient estimator of θ . Let $g_{\theta,n}(\hat{\theta})$ be the marginal distribution of $\hat{\theta}(\underline{X})$ based on the sample $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f_\theta$. Assume there exists functions $G_{ijk}(\hat{\theta})$ such that $\left| \frac{\partial^3}{\partial \theta_i \partial \theta_j \partial \theta_k} \log g_{\theta,n}(\hat{\theta}) \right| \leq n G_{ijk}(\hat{\theta})$ for all $\hat{\theta}$, all $n \geq 1$, and all $\theta \in B(\theta_0)$, where $\mathbb{E}_{\hat{\theta} \sim \theta} G_{ijk}(\hat{\theta}) < \infty$.

Theorem 4.2 shows that when using ACSS with plug-in value of θ_n for θ , the KL distance between the approximate samples and the true sampling distribution goes to zero so long as $\theta_n = \theta + O(n^{-1/2})$. This result heavily relies on the efficiency of the estimator. Recall that Le Cam et al. [1956] show that an efficient estimator is asymptotically sufficient. Theorem 4.2 illustrates this phenomenon by showing that the dependence on θ_n is small, as n grows.

Theorem 4.2. *Under assumptions (R0)-(R4), let $\theta \in \Theta$ and let θ_n be a sequence of values in Θ . Let $\hat{\theta}(\cdot)$ be a randomized estimator based on a sample $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} f_\theta$, with conditional*

distribution $g_n(\hat{\theta} \mid \underline{X})$ and marginal distribution $g_{\theta,n}(\hat{\theta})$. Then the KL divergence between the sample X_1, \dots, X_n and $Y_1, \dots, Y_n \sim f_{\theta_n}^n(y_1, \dots, y_n \mid \hat{\theta}(\underline{y}) = \hat{\theta}(\underline{X}))$ is

$$\text{KL}(X_1, \dots, X_n \parallel Y_1, \dots, Y_n) = o(n)\|\theta_n - \theta\|^2 + O(n)\|\theta_n - \theta\|^3.$$

In particular, if $\theta_n - \theta = O(n^{-1/2})$, then the above KL divergence goes to zero as $n \rightarrow \infty$.

Next we return to the ICSS algorithm. In Lemma 4.3, we give a formula for the conditional distribution of the ICSS algorithm, and show that it can be written in terms of the data generating distribution conditioned on the efficient statistic.

Lemma 4.3. *We assume (R0) and use the notation therein. Let $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} f_{\theta_0}$ and let $\hat{\theta} \in \Theta$. We assume that there exists a unique solution θ to the equation $\hat{\theta}(X_\theta(\omega)) = \hat{\theta}$ for all values of $\omega_1, \dots, \omega_n \in \Omega$. Let $\omega_1, \dots, \omega_n \stackrel{i.i.d.}{\sim} P$, let $\theta^* \in \Theta$, and call $Y_i^{\theta^*} = X_{\theta^*}(\omega_i)$. Then $Y_1^{\theta^*}, \dots, Y_n^{\theta^*} \mid \left(\hat{\theta}(\underline{Y}^{\theta^*}) = \hat{\theta} \right) \sim f_{\theta^*}^n(y_1, \dots, y_n \mid \hat{\theta}(\underline{y}) = \hat{\theta})$.*

We will combine Lemma 4.3 and Theorem 4.2 to address the performance of the ICSS algorithm, but first we need to characterize the rate of convergence of θ^* to θ . It will be convenient to call $h_{\theta,n}(\theta^*)$ the marginal distribution of θ^* , where the random variable θ^* is generated as follows: Assuming (R0)-(R4), that there exists a unique solution to Equation (1), and using the notation of Theorem 4.2, sample $\hat{\theta} \sim g_{\theta,n}(\hat{\theta})$, where $g_{\theta,n}(\cdot)$, and sample $\omega_1, \dots, \omega_n \stackrel{i.i.d.}{\sim} P$. Then θ^* is defined to be the solution to the equation $\hat{\theta}(X_{\theta^*}(\underline{\omega})) = \hat{\theta}$.

Lemma 4.4. *Assume (R0-R4) and that there exists a unique solution to Equation (1). Let $\theta^* \sim h_{\theta,n}(\theta^*)$. Then $\theta^* = \theta + O_p(n^{-1/2})$.*

Finally, we combine our results to prove Theorem 4.5. Connecting Theorem 4.5 with Lemma 4.3, we interpret Theorem 4.5 roughly as follows: the conditional distribution generated by ICSS using an efficient statistic approaches the true conditional distribution in

terms of the expected KL divergence, as the sample size goes to infinity. Based on Theorem 4.5, we expect that any test which tries to distinguish between \underline{Y} and \underline{X} will have power approaching the type I error as $n \rightarrow \infty$. We confirm this through simulations in Section 5.1.

Theorem 4.5. *Assume (R0)-(R4) and that there is a unique solution to Equation (1). Let $\theta^* \sim h_{\theta,n}(\theta^*)$. Then*

$$\mathbb{E}_{\hat{\theta} \sim g_{\theta,n}} \text{KL} \left(X_1, \dots, X_n \mid \hat{\theta}(\underline{X}) = \hat{\theta} \mid X_{\theta^*}(\omega_1), \dots, X_{\theta^*}(\omega_n) \mid \hat{\theta}(X_{\theta^*}(\underline{\omega})) = \hat{\theta} \right) = o_p(1),$$

where $X_i \stackrel{i.i.d.}{\sim} f_{\theta}$, $\omega_i \stackrel{i.i.d.}{\sim} P$, and $g_{\theta,n}(\cdot)$ is defined in Theorem 4.2.

Theorem 4.5 follows from the proof of Theorem 4.2 in combination with Lemma 4.4. Note that in Theorem 4.5, θ^* is not included in the expectations on the left hand side, so the expression is a function of the random variable θ^* . Similarly, the $o_p(1)$ term depends on the randomness in the random variable θ^* .

4.2 A computationally efficient one-step estimator

So far, we have developed distributional properties of the ICSS algorithm when applied to efficient statistics. However, solving the optimization problem in Equation (1) is can be challenging. In fact, as we demonstrate in Example 4.6 it may be the case that an exact solution does not exist for many models. In this section, we propose a One Approximate-Step Indirect co-Sufficient Sampling method (OASIS Sampling), which gives a “one-step” approximate solution to equation (1). We prove in Theorem 4.7 that OASIS results in a sample with $\hat{\theta}(\underline{Y}) = \hat{\theta}(\underline{X}) + o_p(n^{-1/2})$.

Example 4.6. [No solution to Equation (1)] Suppose that $X_i = W_i + N_i$, where $W_i \stackrel{i.i.d.}{\sim} \text{Bern}(\theta)$ and $N_i \stackrel{i.i.d.}{\sim} U(0, 1)$ and $W_i \perp\!\!\!\perp N_i$. Note that $\hat{\theta}(\underline{X}) = n^{-1} \sum_{i=1}^n X_i - \frac{1}{2}$ is an efficient

estimator. Using the notation of (R0), $\Omega = U(0, 1)^2$ and $X_\theta : \Omega \rightarrow \mathbb{R}$, defined by $X_\theta(u_1, u_2) = F_\theta^{-1}(u_1) + u_2$, where $F_\theta^{-1}(\cdot)$ is the quantile function for $\text{Bern}(\theta)$. Then, for any value $\theta^* \in [0, 1]$, $\hat{\theta}(\underline{Y}) = \hat{\theta}(X_{\theta^*}(u_1^{(1)}, u_2^{(1)}), \dots, X_{\theta^*}(u_1^{(n)}, u_2^{(n)}))$ only takes $n + 1$ possible values: $\left\{ n^{-1} \sum_{i=1}^n u_2^{(i)} - \frac{1}{2} + \frac{i}{n} \mid i \in \{0, 1, 2, \dots, n\} \right\}$. As such, given $\hat{\theta}(\underline{X})$ from an independent sample, the probability that $\hat{\theta}(\underline{X})$ is in this set is zero. From another perspective, the decimal value of $n\hat{\theta}(\underline{Y})$ is completely determined by $n^{-1} \sum_{i=1}^n u_2^{(i)}$, and is not influenced by θ^* .

Based on Example 4.6, we cannot expect to always find an exact solution to Equation (1). In Theorem 4.7 and Algorithm 2, we propose a one-step approximate solution which is based on the following intuition: Call $Z_i = X_{\hat{\theta}(\underline{X})}(\omega_i)$ and $Y_i = X_{\theta^*}(\omega_i)$, where θ^* is the exact solution to Equation (1). Since \underline{Z} and \underline{Y} use the same seeds, we expect that the differences $\hat{\theta}(\underline{Y}) - \theta^* = \hat{\theta}(\underline{X}) - \theta^*$ and $\hat{\theta}(\underline{Z}) - \hat{\theta}(\underline{X})$ should be similar, especially for larger sample sizes. Solving $\hat{\theta}(\underline{X}) - \theta^* = \hat{\theta}(\underline{Z}) - \hat{\theta}(\underline{X})$ for θ^* gives the solution $\theta^* = 2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z})$, which is exactly the one-step estimator of Theorem 4.7.

The following Theorem shows that regardless of whether a solution to Equation (1) exists, the OASIS sampling algorithm preserves the efficient statistic up to $o_p(n^{-1/2})$. Note that the assumption (R4) is not necessary for Theorem 4.7.

Theorem 4.7. *Assume that (R0)-(R3) hold. Let $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} f_{\theta_0}$ and let $\omega_1, \dots, \omega_n \stackrel{i.i.d.}{\sim} P$. Set $\theta_1^* = \text{Proj}_\Theta(2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z}))$, where $\text{Proj}_\Theta(\cdot)$ is the projection operator onto the set Θ , $\hat{\theta}$ is an efficient estimator, and $(Z_i)_{i=1}^n = (X_{\hat{\theta}(\underline{X})}(\omega_i))_{i=1}^n$. Then for $(Y_i)_{i=1}^n = (X_{\theta_1^*}(\omega_i))_{i=1}^n$, we have $\hat{\theta}(\underline{Y}) = \hat{\theta}(\underline{X}) + o_p(n^{-1/2})$.*

In combination with the results of Section 4.1, Theorem 4.7 is very powerful. With only a constant multiple increase in computational time compared to the parametric bootstrap, OASIS sampling is able to produce samples which preserve an efficient statistic with

asymptotically negligible error. In Algorithm 2, we give an R pseudo-code implementation of the OASIS method, demonstrating how the method is very easily implemented.

In Theorem 4.7 and Algorithm 2, the role of the projection operator is to address parameter spaces that may not be closed under linear combinations. For example, in multinomial models, the parameter space is a probability simplex. The projection operator ensures that θ_1^* always lies within Θ . Note that the exact projection used does not affect the asymptotic result of Theorem 4.7, since for large n , $2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z})$ will lie in Θ with probability one.

Algorithm 2: OASIS Sampling Pseudo-code in R

Input: Seed ω , efficient estimator $\hat{\theta}(\underline{X})$, function `theta_hat(y)` to compute $\hat{\theta}(y)$,
function `rsample(θ)` to sample n i.i.d. samples from f_θ .

```

1 set.seed( $\omega$ )
2  $\underline{Z} = \text{rsample}(\hat{\theta}(\underline{X}))$ 
3  $\hat{\theta}(\underline{Z}) = \text{theta\_hat}(\underline{Z})$ 
4  $\theta_1^* = 2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z})$ 
5 if  $\theta_1^* \notin \Theta$  then
6   |  $\theta_1^* = \text{Proj}_\Theta(\theta_1^*)$ 
7 set.seed( $\omega$ )
8  $\underline{Y} = \text{rsample}(\theta_1^*)$ 

```

Output: Y_1, \dots, Y_n

Remark 4.8 (Intractable Likelihood Functions). Another strength of ICSS, and the OASIS method in particular, is that it can be applied to problems which are easy to sample, but which have intractable or expensive likelihood functions. This may occur when there are latent variables that must be marginalized to evaluate the likelihood. A notable example of this is in differential privacy: the private sample X_1, \dots, X_n is drawn from f_θ , and then the privatized result S is drawn from a privacy mechanism $Q(S|X_1, \dots, X_n)$. At this point, the analyst only has access to the privatized result S . Deriving the marginal distribution of S can be computationally intractable, as it requires integrating over all

possible datasets X_1, \dots, X_n . This phenomenon of differential privacy was first noted in Williams and McSherry [2010]. In Karwa et al. [2016], a variational approximation was used to approximate the intractable likelihood for a privatized log-linear model. Note that the MCMC methods of Barber and Janson [2020] are inapplicable in such settings as MCMC methods require evaluations of the likelihood function for each step of the Markov chain.

5 Examples and simulations

5.1 Burr type XII distribution

The Burr Type XII distribution, denoted $\text{Burr}(c, k)$, also known as the Singh–Maddala distribution, is a useful model for income [McDonald, 2008]. The distribution has pdf $f(x) = ckx^{c-1}(1+x^c)^{-(k+1)}$, with support $x > 0$. Both c and k are positive. The Burr distribution was chosen for our first simulation because 1) it is one-dimensional, allowing for the Kolmogorov-Smirnov (K-S) test to be applied, and 2) as it is not exponential family or location-scale, the conditional distribution is non-trivial. The goal of the following simulation is to verify that the samples generated using OASIS sampling to preserve the efficient statistic with negligible error, and that the samples are indistinguishable from the original unknown distribution, as tested via the K-S test.

For the simulation, we set $c = 2$ and $k = 4$, and denote $\theta = (c, k)$. Let $\hat{\theta}_{MLE}$ be the MLE. We draw $X_i \stackrel{\text{i.i.d.}}{\sim} \text{Burr}(2, 4)$, $Z_i \stackrel{\text{i.i.d.}}{\sim} \text{Burr}(\hat{\theta}_{MLE}(\underline{X}))$, and $(Y_i)_{i=1}^n$ from Algorithm 2. The simulation is conducted for $n \in \{100, 1000, 10000\}$ with results averaged over 10000 replicates for each n .

Over the replicates, we compute the MLE and report the average squared ℓ_2 -distance to the true parameters, which estimates the variance. The results are in Table 1. When

sampling from the the fitted model, $\hat{\theta}(\underline{Z})$ has about twice the variance as $\hat{\theta}(\underline{X})$, whereas $\hat{\theta}(\underline{Y})$ has very similar variance as $\hat{\theta}(\underline{X})$.

We also calculate the empirical power of the K-S test, comparing each sample with the true distribution Burr(2, 4), at type I error .05. The results are presented in Table 2. We see that the (X_i) have empirical power approximately .05, confirming that the type I error is appropriately calibrated. We also see that the K-S test using (Y_i) has power approximately .05, indicating that the empirical distribution of the OASIS samples (Y_i) is very close to the true distribution. On the other hand, we see that the K-S test with (Z_i) has power .15, significantly higher than the type I error, indicating that the parametric bootstrap samples (Z_i) are from a fundamentally different distribution than the (X_i) . This result is in agreement with Proposition 3.6 and the results of Section 4.

Table 1: Average squared ℓ_2 -distance between the MLE and the vector (2, 4). (X_i) are drawn i.i.d. from Burr(2, 4), (Z_i) are i.i.d. from Burr($\hat{\theta}(\underline{X})$), and (Y_i) are from Algorithm 2. Results are averaged over 10000 replicates, for each n . The first and third lines are accurate up to approximately ± 2 in the third digit of each value with 95% confidence. The second line has error ± 4 in the third digit.

n :	100	1000	10000
$\hat{\theta}(\underline{X})$	2.6252×10^{-1}	2.2254×10^{-2}	2.1992×10^{-3}
$\hat{\theta}(\underline{Z})$	5.8542×10^{-1}	4.4763×10^{-2}	4.4149×10^{-3}
$\hat{\theta}(\underline{Y})$	2.6211×10^{-1}	2.2178×10^{-2}	2.1994×10^{-3}

5.2 Log-linear model

This example is based on a dataset of of 68,694 passengers in automobiles and light trucks involved in accidents in the state of Maine in 1991. Table 3 reports the number of passengers according to gender (G), location (L), seatbelt status (S), and injury status (I).

Table 2: Empirical power of the Kolmogorov-Smirnov test for the distribution Burr(2, 4) at type I error .05. (X_i) are drawn i.i.d from Burr(2, 4), (Z_i) are drawn i.i.d from Burr($\hat{\theta}(\underline{X})$), and (Y_i) are from Algorithm 2. Results are averaged over 10000 replicates, for each n . Standard errors are approximately 0.0022 for lines 1 and 3, and 0.0036 for line 2.

n :	100	1000	10000
(X_i)	0.0471	0.0464	0.0503
(Z_i)	0.1524	0.1541	0.1493
(Y_i)	0.0544	0.0489	0.0485

As in Agresti [2003], we fit a hierarchical log-linear model based on all one-way effects and two-way interactions. The model is summarized in Equation (2), where μ_{ijkl} represents the expected count in bin i, j, k, ℓ . The parameter λ_i^G represents the effect of Gender, and parameter λ_{ij}^{GL} represents the interaction between Gender and Location. The other main effects and interactions are analogous.

$$\log \mu_{ijkl} = \lambda + \lambda_i^G + \lambda_j^L + \lambda_k^S + \lambda_\ell^I + \lambda_{ij}^{GL} + \lambda_{ik}^{GS} + \lambda_{i\ell}^{GI} + \lambda_{jk}^{LS} + \lambda_{j\ell}^{LI} + \lambda_{k\ell}^{SI} \quad (2)$$

For our simulations, we treat the fitted parameters as the true parameters, to ensure that model assumptions are met. We simulate from the fitted model at sample sizes $n \in \{10^2, 10^3, 10^4, 10^5\}$ and compare the performance in terms of the fitted probabilities for each bin of the contingency table. The results are plotted in Figure 1a, with both axes on log-scale. The “mean error” is the average squared ℓ_2 distance between the estimated parameter vector and the true parameter vector, averaged over 200 replicates. To interpret the plot, note that if the error is of the form $\text{error} = cn^{-1}$, where c is a constant, then $\log(\text{error}) = c + (-1)\log(n)$. So, the slope represents the convergence rate, and the vertical offset represents the asymptotic variance. In Figure 1a, we see that the curve for $\hat{\theta}(\underline{Y})$, based on OASIS samples, approaches the curve for $\hat{\theta}(\underline{X})$, indicating that they have the same asymptotic rate and variance. On the other hand, the curve for $\hat{\theta}(\underline{Z})$, based on parametric

Table 3: Injury, Seat-Belt Use, Gender, and Location. Source: Agresti [2003, Table 8.8]. Originally credited to Cristanna Cook, Medical Care Development, Augusta, Maine.


Gender	Location	Seatbelt	Injury	
			No	Yes
Female	Urban	No	7,287	996
		Yes	11,587	759
	Rural	No	3,246	973
		Yes	6,134	757
Male	Urban	No	10,381	812
		Yes	10,969	380
	Rural	No	6,123	1,084
		Yes	6,693	513

bootstrap samples, has the same slope, but does not approach the $\hat{\theta}(\underline{X})$ curve, indicating that $\hat{\theta}(\underline{Z})$ has the same rate but inflated variance.

Recall that our procedure approximately preserves the sufficient statistics, similar to sampling from a conditional distribution. Previous work has proposed procedures to sample directly from conditional distributions for contingency table data. However, these approaches require sophisticated tools from algebraic statistics, and are computationally expensive (e.g., MCMC) [Karwa and Slavković, 2013]. In contrast, our approach is incredibly simple to implement and highly computationally efficient. Our approach is also applicable for a wide variety of models, whereas the techniques to sample directly from the conditional distribution often require a tailored approach for each setting.

5.3 DP beta distributed synthetic data

In this example, we assume that $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(\alpha, \beta)$, where $\alpha, \beta \geq 1$, and our goal is to produce differentially private synthetic data. Often, to ensure finite sensitivity, the



(a) Simulations corresponding to the log-linear model with two-way interactions from Section 5.2

(b) Simulations for the beta distribution from Section 5.3. $\hat{\theta}(\underline{X})$ is the MLE. $\hat{\theta}(\underline{Z})$ and $\hat{\theta}(\underline{Y})$ both satisfy 1-DP.

Figure 1: Both figures plot the average squared ℓ_2 -distance between the estimated parameters and the true parameters on the log-scale. Averages are over 200 replicates for both plots. $\hat{\theta}(\underline{X})$ is from the true model, $\hat{\theta}(\underline{Z})$ from the fitted model, and $\hat{\theta}(\underline{Y})$ from Algorithm 2.

data are clamped to artificial bounds $[a, b]$, introducing bias in the DP estimate. Typically, these bounds are fixed in n , resulting in asymptotically negligible noise, but $O_p(1)$ bias. However, we will show that it is possible to increase the bounds in n to produce both noise and bias of order $o_p(n^{-1/2})$, resulting in an efficient DP estimator. We show through simulations that using this estimator along with OASIS results in a DP sample with optimal asymptotics. While we work with the beta distribution, this approach may be of value for other exponential family distributions as well.

Recall that $n^{-1} \sum_{i=1}^n \log(X_i)$ and $n^{-1} \sum_{i=1}^n \log(1 - X_i)$ are sufficient statistics for the beta distribution. We will add Laplace noise to each of these statistics to achieve differential

privacy. However, the sensitivity of these quantities is unbounded. First we will pre-process the data by setting $\tilde{X}_i = \min\{\max(X_i, t), 1 - t\}$, where t is a threshold that depends on n . Then the ℓ_1 -sensitivity of the pair of sufficient statistics is $\Delta(t) = 2n^{-1} |\log(t) - \log(1 - t)|$. We add independent noise to each of the statistics from the distribution $\text{Laplace}(\Delta(t)/\epsilon)$, which results in ϵ -DP versions of these statistics. Finally, we estimate $\theta = (\alpha, \beta)$ by plugging in the privatized sufficient statistics into the log-likelihood function and maximizing over θ . The resulting parameter estimate satisfies ϵ -DP by post-processing.

We must carefully choose the threshold t to ensure that the resulting estimate is efficient. The choice of t must satisfy $\Delta(t) = o(n^{-1/2})$ to ensure that the noise does not affect the asymptotics of the likelihood function. We also require that both $P(X_i < t) = o(n^{-1/2})$, and $P(X_i > 1 - t) = o(n^{-1/2})$ to ensure that $\tilde{X}_i = X_i + o_p(n^{-1/2})$, which limits the bias to $o_p(n^{-1/2})$. For the beta distribution, we can calculate that $P(X_i < t) = O(t^\alpha)$ and $P(X_i > 1 - t) = O(t^\beta)$. Since we assume that $\alpha, \beta \geq 1$, so long as $t = o(n^{-1/2})$ the probability bounds will hold. Taking $t = \min\{1/2, 10/(\log(n)\sqrt{n})\}$ satisfies $t = o(n^{-1/2})$, and we estimate the sensitivity as

$$\Delta(t) \leq 2n^{-1} \log(t^{-1}) \leq 2n^{-1} \log(\log(n)\sqrt{n}) = O(\log(n)/n) = o(n^{-1/2}),$$

which satisfies our requirement for Δ . While there are many choice of t which satisfy the requirements, our threshold (including the constant 10) was chosen to optimize the finite sample performance, so that the asymptotics could be illustrated with smaller sample sizes.

For the simulation, we sample $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(5, 3)$ for $n \in \{10^3, 10^4, 10^5, 10^6\}$. We estimate $\hat{\theta}(\underline{X})$ with the MLE. Using $\epsilon = 1$, we privatize the sufficient statistics as described above, and obtain $\hat{\theta}_{DP}$ from the privatized log-likelihood function. We sample $Z_1, \dots, Z_n \stackrel{\text{i.i.d.}}{\sim} f_{\hat{\theta}_{DP}}$ and estimate $\hat{\theta}(\underline{Z})$ using maximum likelihood. We produce $(Y_i)_{i=1}^n$ from

Algorithm 2 using $\hat{\theta}_{DP}$ in place of $\hat{\theta}(\underline{X})$. In Figure 1b, we plot the average squared ℓ_2 error between each estimate of θ from the true value $(5, 3)$. The errors are averaged over 200 replicates, and are plotted on the log-scale. We see that $\hat{\theta}_{DP}$ and $\hat{\theta}(\underline{Y})$ have the same asymptotic performance as the MLE, whereas $\hat{\theta}(\underline{Z})$ has inflated variance. See the discussion in Example 5.2 to understand this interpretation of the log-scale plot.

5.4 DP two sample proportion test

Suppose that we collect two independent samples of binary data, one from a “control population” and another from a “treatment population”. We denote $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \text{Bern}(\theta_X)$ as the control sample, and $Y_1, \dots, Y_m \stackrel{\text{i.i.d.}}{\sim} \text{Bern}(\theta_Y)$ as the treatment sample. To satisfy ϵ -differential privacy, the data curators release the following noisy statistics: $\tilde{X} = \sum_{i=1}^n X_i + N_1$ and $\tilde{Y} = \sum_{i=1}^m Y_i + N_2$, where $N_1, N_2 \stackrel{\text{i.i.d.}}{\sim} \text{Tulap}(0, \exp(-\epsilon), 0)$; the sample sizes m and n are released without modification. The Tulap distribution was proposed in Awan and Slavković [2018] and Awan and Slavković [2020] as the optimal mechanism for Bernoulli data in terms of generating uniformly most powerful hypothesis tests and uniformly most accurate confidence intervals. Recall that $N_1 \stackrel{d}{=} N_2 \stackrel{d}{=} G_1 - G_2 + U$, where $G_1, G_2 \stackrel{\text{i.i.d.}}{\sim} \text{Geom}(1 - \exp(-\epsilon))$ and $U \sim \text{Unif}(-1/2, 1/2)$. We can think of \tilde{X} and \tilde{Y} as noisy counts for the number of ones in the control and treatment groups.


Based on the privatized summary statistics \tilde{X} and \tilde{Y} , we are to test $H_0 : \theta_X = \theta_Y$ versus $H_1 : \theta_X \leq \theta_Y$. Without privacy, there exists a uniformly most powerful test, which is constructed by conditioning on the total number of ones: $\sum_{i=1}^n X_i + \sum_{j=1}^m Y_j$, a minimal sufficient statistic under the null hypothesis. However, with the noisy counts, it can be verified that there is not a low-dimensional sufficient statistic. On the other hand, an efficient estimator for $\theta_X = \theta_Y$ under the null hypothesis is $\hat{\theta}(\tilde{X}, \tilde{Y}) = \min\{\max\{(\tilde{X} + \tilde{Y})/(m + n), 0\}, 1\}$. Note that deriving the exact distribution of $(\tilde{X}, \tilde{Y}) \mid \hat{\theta}(\tilde{X}, \tilde{Y})$ is fairly complex,

involving the convolution of distributions. However, the indirect method can easily produce approximate samples from this conditional distribution. In what follows, we use the OASIS sampling algorithm, given in Algorithm 2 and investigate the properties of a hypothesis test based on this conditional distribution in comparison with a parametric bootstrap test.

Recall that without privacy, the uniformly most powerful test uses the test statistic Y , and threshold computed from the conditional distribution of $Y|X + Y$ under the null hypothesis. With privacy, we will use the test statistic \tilde{Y} , and compute the p -values based on $\tilde{Y} | \hat{\theta}(\tilde{X}, \tilde{Y})$. In particular, we compare the performance of this test versus the “parametric bootstrap” test, which uses the test statistic \tilde{Y} based on the approximate sampling distribution, which is the convolution of $\text{Binom}(\hat{\theta}(\tilde{X}, \tilde{Y}))$ and $\text{Tulap}(0, \exp(-\epsilon), 0)$.

For the simulation, we use the sample size $n = m = 200$, fix $\theta_X = 0.3$, set the privacy parameter to $\epsilon = 1$, and base the simulation on 10,000 replicates. Under the null hypothesis, where $\theta_Y = \theta_X$, we plot the empirical cumulative distribution function (CDF) for the p -values of the proposed test as well as for the parametric bootstrap test in Figure 2(a). Recall that a properly calibrated p -value will have the CDF of $U(0, 1)$. We see that the empirical CDF of the p -values for the OASIS test closely approximate the ideal CDF, whereas the p -values of the parametric bootstrap test are overly conservative.

Next, we study the power of the OASIS test versus the parametric bootstrap test in Figure 2(b). For this simulation, we set $n = m = 200$, fix $\theta_X = .3$, and set $\epsilon = 1$. We vary the value of θ_Y along the x -axis by increments of .001 and plot the empirical power of the two tests, averaged over 10,000 replicates for each value of θ_Y . We see that the OASIS test offers a considerable increase in power over the bootstrap test.



(a) Under the null hypothesis, the empirical cumulative distribution of both tests. $\theta_X = \theta_Y = .3$. Results based on 10,000 replicates.

(b) Empirical power at type I error .05. $\theta_X = .3$, θ_Y is varied along the x -axis. Results averaged from 10,000 replicates for each value of θ_Y .

Figure 2: Simulations for the DP two sample proportion test of Section 5.4. In red is the OASIS test, and in blue is the parametric bootstrap test. Sample sizes are $n = m = 200$, privacy parameter is $\epsilon = 1$, and type I error is .05.

5.5 Running time versus MCMC for linear regression

In this section, we compare the running time of Algorithm 2 to an MCMC method of sampling the conditional distribution, proposed by Barber and Janson [2020]. We will base the comparison on a linear regression problem.

We use a similar setting as Simulation 1 in Barber and Janson [2020]. We sample $Y_i \stackrel{\text{i.i.d.}}{\sim} N(Z_i\beta, 1)$, where β and Z_i are 5-dimensional, with the true (unknown) parameter $\beta^\top = (.2, .2, .2, .2, .2)$, and $Z_i^\top \stackrel{\text{i.i.d.}}{\sim} N(0, I)$. For synthetic data, we condition on the matrix Z , as well as on an efficient estimator for β . The approach of Barber and Janson [2020] requires


adding a small amount of independent normal noise $w \sim N(0, \sigma^2 I)$ to the score function before minimizing. To ensure that the estimator is still efficient, we set $\sigma = 1/n$. According to Barber and Janson [2020], the pdf of the conditional distribution is proportional to

$$f_{\hat{\beta}}(Y|Z) \exp\left(-\frac{\|\nabla\ell(\hat{\beta} | Y, Z)\|^2}{2\sigma^2 d}\right) \det\left(\nabla^2\ell(\hat{\beta} | Y, Z)\right), \quad (3)$$

where $\ell(\beta | Y, Z)$ is the log-likelihood function. This unnormalized density can then be used in an MCMC sampler. While this is a relatively simple model, sampling from the conditional distribution is still non-trivial as we are tasked with sampling from an n -dimensional space where the mass is concentrated around an $n - 5$ dimensional subspace.

For each $n \in \{2^6, 2^7, \dots, 2^{14}\}$, we simulate a sample (Y, Z) of the appropriate dimension. To do this, we run 100 rounds of a one-at-a-time MCMC method with Gaussian proposals, which consists of $100 \cdot n$ one-dimensional update steps, using Equation (3) and step-size of .1. We then run 1000 replicates of the OASIS method of Algorithm 2. In Figure 3, we average the replicates to get the average time for one round of the MCMC, as well as the average run time of Algorithm 2.

In the left plot of Figure 3, we see that even a single round of the MCMC is much more computationally expensive than a complete run of Algorithm 2. Through additional simulations at smaller sample sizes ($n \approx 100$), we found that approximately 35,000 rounds of the MCMC algorithm were required to produce a sample with less than .01 correlation with the initial state, when the step size was optimized to have between 10 – 20% acceptance rate. Thus, to produce a single sample at $n = 16384$, that is minimally correlated with the original Y , the MCMC method is expected to take $295610s \approx 82\text{hours}$. On the other hand, Algorithm 2 only takes $0.002538s$ to generate a sample. If these simulation methods are to be used for hypothesis testing p -value computations, then thousands of samples are needed.



(a) The ‘x’ represent the computed datapoints, whereas the lines are linear interpolations.

(b) Both axes are log-scale. In orange is a line with slope 2, and in green is a line of slope 1.

Figure 3: Linear regression simulation of Section 5.5. Running time for one round of a one-at-a-time MCMC in red, and average running time of one run of Algorithm 2 in blue.

Furthermore, in real-world datasets n and p may be far larger than those considered in this simulation, making the MCMC implementation even more impractical.

In the right plot of Figure 3, we have another plot of the running time of the MCMC procedure as well as Algorithm 2, but with both axes on the log-scale. As discussed in Section 5.2, this sort of plot allows us to infer the computational complexity. In green is a line of slope 1, which represents the computational complexity $O(n)$, and we see that the running time of Algorithm 2 has approximately the same slope as this line. In orange is a line of slope 2, which represents the computation complexity of $O(n^2)$, and we see that the running time of one round of the MCMC procedure has approximately the same slope as this line. Thus, not only is the observed running time of Algorithm 2 much smaller than for

the MCMC, but the computation complexity is reduced from $O(n^2)$ to $O(n)$.

Remark 5.1 (Computational Complexity). The computation complexity of Algorithm 2 is expected to be $O(n)$ for many common models, as it requires simulating samples of size n , and computing an efficient estimator for the parameter, which can often be done in $O(n)$ time. On the other hand, MCMC procedures are expected to take at least $O(n^2)$ time, as they require at least $O(n)$ steps to converge (consider for instance, the one-at-a-time MCMC) as well as evaluating the likelihood at each step, which often takes at least $O(n)$ time. This intuition is confirmed by the results in Figure 3.

This analysis assumes that the likelihood has a convenient expression that can be evaluated in $O(n)$ time. For latent variable and missing-data problems, the likelihood may be far more expensive, or even computationally intractable.

6 Discussion

In this paper, we extended the ICSS algorithm to models with efficient statistics, but which may not have low-dimensional sufficient statistics. We also proposed the OASIS sampling algorithm as an approximate solution to ICSS, which is easily implemented and highly computationally efficient. As compared to Barber and Janson [2020], our method offers similar theoretical guarantees, but with a much improved computational cost, and ease of implementation. A strength of our approach is that it only requires the ability to estimate parameters and sample from the model. This is great for usability as many practitioners can easily implement Algorithms 1 and 2, but may not have the expertise to implement a customized MCMC procedure. Furthermore, our method can be easily applied even to models with intractable likelihood functions, whereas MCMC approaches are inapplicable.

One area where likelihood-free inference is especially applicable is in the setting of

differential privacy. Suppose that the data analyst releases privatized summary statistics, we may wish to use these statistics to test certain hypotheses about the distribution that generated the private data. However, despite the fact that many DP estimators have asymptotically negligible noise, standard asymptotic tests often have unacceptably poor performance with DP data [Wang et al., 2018]. Furthermore, deriving the exact marginal likelihood of the DP output often requires evaluating a high dimensional and intractable integral [Williams and McSherry, 2010]. While parametric bootstrap tests are often applicable, as mentioned in Section 3.1 they do not guarantee accurate type I error rates. The OASIS algorithm is nearly as easy to implement as the parametric bootstrap, yet as we showed in Section 5.4 gives accurate p -values, even in moderate sample sizes.

The theory of this paper studied the case that the data was independent and identically distributed, but as demonstrated in the linear regression simulation, the method is applicable to a wider variety of models than just i.i.d. data. The results of Section 4 can likely be extended to account for more general models. We leave this for future work.

Co-sufficient sampling have been successfully used to produce valid confidence intervals in the presence of nuisance parameters [Lillegard and Engen, 1999]. However, the results of Lillegard and Engen [1999] are limited to the case of exact co-sufficient sampling, and where the nuisance parameters are location-scale. It would be interesting to investigate whether our method of approximate co-sufficient sampling can be used to generate confidence intervals in more complex models, when sufficient statistics are not available.

The distributional results of this paper are based on exact solutions to Equation (1). It would be worth investigating how to bridge the gap between the exact ICSS method and the OASIS algorithm. For instance, one could investigate conditions for an exact solution to Equation (1) as well as computationally efficient algorithms to solve this equation, or alternatively one could directly develop distributional results for OASIS sampling.

Acknowledgments

The authors are grateful to Matthew Reimherr, Aleksandra Slavković, and Roberto Molinari for discussions and feedback on early drafts.

References

- John M Abowd. The US Census Bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2867–2867, 2018.
- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. *Advances in Neural Information Processing Systems*, 31:6878–6891, 2018.
- Alan Agresti. Exact inference for categorical data: recent advances and continuing controversies. *Statistics in medicine*, 20(17-18):2709–2722, 2001.
- Alan Agresti. *Categorical data analysis*, volume 482. John Wiley & Sons, 2003.
- Alan Agresti et al. A survey of exact inference for contingency tables. *Statistical science*, 7(1):131–153, 1992.
- Jordan Awan and Aleksandra Slavković. Differentially private uniformly most powerful tests for binomial data. In *Advances in Neural Information Processing Systems*, pages 4208–4218, 2018.
- Jordan Alexander Awan and Aleksandra Slavković. Differentially private inference for binomial data. *Journal of Privacy and Confidentiality*, 10(1), 2020.

- Rina Foygel Barber and Lucas Janson. Testing goodness-of-fit and conditional independence with approximate co-sufficient sampling. *arXiv preprint arXiv:2007.09851*, 2020.
- Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473. IEEE, 2014.
- Dennis D Boos et al. Introduction to the bootstrap world. *Statistical Science*, 18(2):168–174, 2003.
- Michel Broniatowski and Virgile Caron. Conditional inference in parametric models. *Journal de la Société Française de Statistique*, 160(2):48–66, 2019.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.
- Mark Bun, Cynthia Dwork, Guy N Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 74–86, 2018.
- Jim Burrridge. Information preserving statistical obfuscation. *Statistics and Computing*, 13(4):321–327, 2003.
- Clément L Canonne, Gautam Kamath, Audra McMillan, Adam Smith, and Jonathan Ullman. The structure of optimal private tests for simple hypotheses. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 310–321, 2019.

- Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. *Advances in Neural Information Processing Systems*, 21:289–296, 2008.
- Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(3), 2011.
- Yuguo Chen, Ian H Dinwoodie, Seth Sullivant, et al. Sequential importance sampling for multiway tables. *The Annals of Statistics*, 34(1):523–545, 2006.
- Christopher C Drovandi, Anthony N Pettitt, and Malcolm J Faddy. Approximate bayesian computation using indirect inference. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 60(3):317–337, 2011.
- John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE, 2013.
- Cynthia Dwork. Differential privacy and the us census. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 1–1, 2019.
- Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006a.

- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006b.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Steinar Engen and Magnar Lillegård. Stochastic simulations conditioned on sufficient statistics. *Biometrika*, 84(1):235–240, 1997.
- Xianjin Fang, Fangchao Yu, Gaoming Yang, and Youyang Qu. Regression analysis with differential privacy preserving. *IEEE access*, 7:129353–129361, 2019.
- Marco Gaboardi, Hyun Lim, Ryan Rogers, and Salil Vadhan. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In *International Conference on Machine Learning*, pages 2111–2120. PMLR, 2016.
- Marc G Genton and Elvezio Ronchetti. Robust indirect inference. *Journal of the American Statistical Association*, 98(461):67–76, 2003.
- Christian Gourieroux, Alain Monfort, and Eric Renault. Indirect inference. *Journal of applied econometrics*, 8(S1):S85–S118, 1993.
- Stéphane Guerrier, Elise Dupuis-Lozeron, Yanyuan Ma, and Maria-Pia Victoria-Feser. Simulation-based bias correction methods for complex models. *Journal of the American Statistical Association*, 114(525):146–157, 2019.
- Knut Heggland and Arnaldo Frigessi. Estimating functions in indirect inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(2):447–462, 2004.

- Anco Hundepool, Josep Domingo-Ferrer, Luisa Franconi, Sarah Giessing, Eric Schulte Nordholt, Keith Spicer, and Peter-Paul De Wolf. *Statistical disclosure control*. John Wiley & Sons, 2012.
- Vishesh Karwa and Aleksandra Slavković. Conditional inference given partial information in contingency tables using Markov bases. *Wiley Interdisciplinary Reviews: Computational Statistics*, 5(3):207–218, 2013.
- Vishesh Karwa and Aleksandra B Slavković. Differentially private graphical degree sequences and synthetic graphs. In *International Conference on Privacy in Statistical Databases*, pages 273–285. Springer, 2012.
- Vishesh Karwa, Dan Kifer, and Aleksandra Slavković. Private posterior distributions from variational approximations. *NIPS 2015 Workshop on Learning and Privacy with Incomplete Data and Weak Supervision*, 2016.
- Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- Daniel Kifer and Bing-Rong Lin. An axiomatic view of statistical privacy and utility. *Journal of Privacy and Confidentiality*, 4(1), 2012.
- Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, pages 25–1, 2012.
- Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum. *Privacy, big data, and the public good: Frameworks for engagement*. Cambridge University Press, 2014.

- Lucien Le Cam et al. On the asymptotic theory of estimation and testing hypotheses. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California, 1956.
- Erich Leo Lehmann. *Elements of large-sample theory*. Springer Science & Business Media, 2004.
- Chong K Liew, Uinam J Choi, and Chung J Liew. A data distortion by probability distribution. *ACM Transactions on Database Systems (TODS)*, 10(3):395–411, 1985.
- Magnar Lillegard and Steinar Engen. Exact confidence intervals generated by conditional parametric bootstrapping. *Journal of Applied Statistics*, 26(4):447–459, 1999.
- Bo H Lindqvist and Bjarte Rannestad. Monte carlo exact goodness-of-fit tests for nonhomogeneous poisson processes. *Applied Stochastic Models in Business and Industry*, 27(3):329–341, 2011.
- Bo Henry Lindqvist and Gunnar Taraldsen. Monte carlo conditioning on a sufficient statistic. *Biometrika*, 92(2):451–464, 2005.
- Bo Henry Lindqvist and Gunnar Taraldsen. Conditional monte carlo based on sufficient statistics with applications. In *Advances In Statistical Modeling And Inference: Essays in Honor of Kjell A Doksum*, pages 545–561. World Scientific, 2007.
- Bo Henry Lindqvist and Gunnar Taraldsen. Exact statistical inference for some parametric nonhomogeneous poisson processes. *Journal of The Iranian Statistical Society*, 12(1):113–126, 2013.

- Fang Liu. Model-based differentially private data synthesis. *arXiv preprint arXiv:1606.08052*, 2016.
- Richard A Lockhart, Federico J O'Reilly, and Michael A Stephens. Use of the gibbs sampler to obtain conditional tests, with applications. *Biometrika*, 94(4):992–998, 2007.
- Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *2008 IEEE 24th International Conference on Data Engineering*, pages 277–286. IEEE, 2008.
- Josep Maria Mateo-Sanz, Antoni Martínez-Ballesté, and Josep Domingo-Ferrer. Fast generation of accurate synthetic microdata. In *International Workshop on Privacy in Statistical Databases*, pages 298–306. Springer, 2004.
- David McClure and Jerome P Reiter. Differential privacy and statistical disclosure risk measures: An investigation with binary synthetic data. *Transactions on Data Privacy*, 5(3):535–552, 2012.
- James B McDonald. Some generalized functions for the size distribution of income. In *Modeling Income Distributions and Lorenz Curves*, pages 37–55. Springer, 2008.
- Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.
- Krishnamurty Muralidhar and Rathindra Sarathy. A theoretical basis for perturbation methods. *Statistics and Computing*, 13(4):329–335, 2003.
- Federico O'Reilly and Leticia Gracia-Medrano. On the conditional distribution of goodness-of-fit tests. *Communications in Statistics—Theory and Methods*, 35(3):541–549, 2006.

- Leandro Pardo. *Statistical inference based on divergence measures*. CRC press, 2018.
- Matthew Reimherr and Jordan Awan. Kng: The k-norm gradient mechanism. In *Advances in Neural Information Processing Systems*, pages 10208–10219, 2019.
- James M Robins, Aad van der Vaart, and Valérie Ventura. Asymptotic distribution of p values in composite null models. *Journal of the American Statistical Association*, 95(452): 1143–1156, 2000.
- James D Santos and Nelson L Souza Filho. A metropolis algorithm to obtain co-sufficient samples with applications in conditional tests. *Communications in Statistics-Simulation and Computation*, 48(9):2655–2659, 2019.
- Mark J Schervish. *Theory of statistics*. Springer Science & Business Media, 2012.
- R. L. Serfling. *Approximation Theorems in Mathematical Statistics*. New York: Wiley, 1980.
- Aleksandra B Slavković and Juyoun Lee. Synthetic two-way contingency tables that preserve conditional frequencies. *Statistical Methodology*, 7(3):225–239, 2010.
- Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, pages 813–822, 2011.
- Michael A Stephens. Goodness-of-fit and sufficiency: Exact and approximate tests. *Methodology and Computing in Applied Probability*, 14(3):785–791, 2012.
- Gunnar Taraldsen and Bo Henry Lindqvist. Conditional fiducial models. *Journal of Statistical Planning and Inference*, 195:141–152, 2018.

Daniel Ting, Stephen Fienberg, and Mario Trottini. Romm methodology for microdata release. *Monographs of Official Statistics*, page 89, 2005.

Aad W Van der Vaart. *Asymptotic statistics*. Cambridge University Press, 2000.

Yue Wang, Daniel Kifer, Jaewoo Lee, and Vishesh Karwa. Statistical approximating distributions under differential privacy. *Journal of Privacy and Confidentiality*, 8(1), 2018.

Oliver Williams and Frank McSherry. Probabilistic inference and differential privacy. *Advances in Neural Information Processing Systems*, 23:2451–2459, 2010.

Ram Zamir. A proof of the fisher information inequality via a data processing argument. *IEEE Transactions on Information Theory*, 44(3):1246–1250, 1998.

Jun Zhang, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, and Marianne Winslett. Functional mechanism: regression analysis under differential privacy. *Proceedings of the VLDB Endowment*, 5(11):1364–1375, 2012.

Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. Privbayes: Private data release via bayesian networks. *ACM Transactions on Database Systems (TODS)*, 42(4):1–41, 2017.

7 Supplementary materials

7.1 Background on differential privacy

In this section, we review the basics of differential privacy (DP), which was originally proposed in Dwork et al. [2006b] as a framework to develop methods of preserving privacy,

with mathematical guarantees. Intuitively, the constraint of differential privacy requires that for all possible databases, the change in one person’s data does not significantly change the distribution of outputs. Consequently, having observed the DP output, an adversary cannot accurately determine the input value of any single person in the database. Definition 7.1 gives a formal definition of DP. In Definition 7.1, $h : \mathcal{X}^n \times \mathcal{X}^n \rightarrow \mathbb{Z}^{\geq 0}$ represents the *Hamming metric*, defined by $h(\underline{x}, \underline{x}') = \#\{i \mid x_i \neq x'_i\}$.

Definition 7.1 (Differential privacy: Dwork et al. [2006b]). Let $\epsilon > 0$ and $n \in \{1, 2, \dots\}$ be given. Let \mathcal{X} be any set, and $(\mathcal{Y}, \mathcal{S})$ a measurable space. Let $\mathcal{M} = \{M_{\underline{x}} \mid \underline{x} \in \mathcal{X}^n\}$ be a set of probability measures on $(\mathcal{Y}, \mathcal{S})$, which we call a *mechanism*. We say that \mathcal{M} satisfies *ϵ -differential privacy* (ϵ -DP) if $M_{\underline{x}}(S) \leq e^\epsilon M_{\underline{x}'}(S)$ for all $S \in \mathcal{S}$ and all $\underline{x}, \underline{x}' \in \mathcal{X}^n$ such that $h(\underline{x}, \underline{x}') = 1$.

An important property of differential privacy is that it is invariant to post-processing. Applying any data-independent procedure to the output of a DP mechanism preserves ϵ -DP [Dwork et al., 2014, Proposition 2.1]. Furthermore, Smith [2011] demonstrated that under conditions similar to (R1)-(R3), there exist efficient DP estimators for parametric models. Using these techniques, our synthetic data procedure can produce DP synthetic data by using a DP efficient statistic.

Remark 7.2. Besides Definition 7.1, there are many other variations of differential privacy, the majority of which are relaxations of Definition 7.1 which also allow for efficient estimators. For instance, approximate DP [Dwork et al., 2006a], concentrated DP [Dwork and Rothblum, 2016, Bun and Steinke, 2016], truncated-concentrated DP [Bun et al., 2018], and Renyi DP [Mironov, 2017] all allow for efficient estimators. On the other hand, local differential privacy [Kasiviswanathan et al., 2011, Duchi et al., 2013] in general does not permit efficient estimators and would not fit in our framework. For an axiomatic treatment of formal

privacy, see Kifer and Lin [2012].

While there are some general methods of producing efficient DP parameter estimates, such as in Smith [2011], often these approaches do not perform well in practical sample sizes. We demonstrate our approach using a modification of the standard *Laplace mechanism*. Given a statistic T , the Laplace mechanism adds independent Laplace noise to each entry of the statistic, with scale parameter proportional to the *sensitivity* of the statistic. Informally, the sensitivity of T is the largest amount that T changes, when one person's data is changed in the dataset.

Proposition 7.3 (Sensitivity and Laplace mechanism: Dwork et al. [2006b]). *Let $\epsilon > 0$ be given, and let $T : \mathcal{X}^n \rightarrow \mathbb{R}^p$ be a statistic. The ℓ_1 -sensitivity of T is $\Delta_n(T) = \sup \|T(\underline{x}) - T(\underline{x}')\|_1$, where the supremum is over all $\underline{x}, \underline{x}' \in \mathcal{X}^n$ such that $h(\underline{x}, \underline{x}') = 1$. Provided that $\Delta_n(T)$ is finite, releasing the vector $(T_j(\underline{x}) + L_j)_{j=1}^p$ satisfies ϵ -DP, where $L_1, \dots, L_p \stackrel{i.i.d.}{\sim} \text{Laplace}(\Delta_n(T)/\epsilon)$.*

7.2 Proofs and technical lemmas

Proof of Proposition 3.6. First we will establish the asymptotic distribution of $\hat{\theta}(\underline{Z})$. Recall that by efficiency, we know that $\sqrt{n}(\hat{\theta}(\underline{X}) - \theta) \xrightarrow{d} N(0, I^{-1}(\theta))$ and $\sqrt{n}(\hat{\theta}(\underline{Z}) - \mathbb{E}[\hat{\theta}(\underline{Z}) \mid \hat{\theta}(\underline{X})]) \mid \hat{\theta}(\underline{X}) \xrightarrow{d} N(0, I^{-1}(\hat{\theta}(\underline{X})))$. Then by Slutsky's theorem, we have that $\sqrt{n}(\hat{\theta}(\underline{Z}) - \mathbb{E}[\hat{\theta}(\underline{Z}) \mid \hat{\theta}(\underline{X})]) \xrightarrow{d} N(0, I^{-1}(\theta))$. We can easily compute that $\text{Cov}(\hat{\theta}(\underline{Z}) - \mathbb{E}[\hat{\theta}(\underline{Z}) \mid \hat{\theta}(\underline{X})], \hat{\theta}(\underline{X})) = 0$ using the law of total covariance. So, we have that $\sqrt{n}(\hat{\theta}(\underline{Z}) - \mathbb{E}[\hat{\theta}(\underline{Z}) \mid \hat{\theta}(\underline{X})] + \hat{\theta}(\underline{X}) - \theta) \xrightarrow{d} N(0, 2I^{-1}(\theta))$. We also know that $\sqrt{n}(\mathbb{E}[\hat{\theta}(\underline{Z}) \mid \hat{\theta}(\underline{X})] - \hat{\theta}(\underline{X})) = o_p(1)$, since $\mathbb{E}[\hat{\theta}(\underline{Z}) \mid \hat{\theta}(\underline{X})] = \hat{\theta}(\underline{X}) + o_p(n^{-1/2})$. Together, we have that $\sqrt{n}(\hat{\theta}(\underline{Z}) - \theta) \xrightarrow{d} N(0, 2I^{-1}(\theta))$.

Next, we lower bound the total variation using the data processing inequality as well as

a lower bound on total variation with the Hellinger distance: $H^2(P||Q) \leq \text{TV}(P||Q)$.

$$\text{TV}(\underline{X}||\underline{Z}) \geq \text{TV}(\hat{\theta}(\underline{X})||\hat{\theta}(\underline{Z})) \quad (4)$$

$$= \text{TV}(N(0, I^{-1}(\theta))||N(0, 2I^{-1}(\theta))) + \text{TV}(\hat{\theta}(\underline{X})||N(0, I^{-1}(\theta))) \quad (5)$$

$$+ \text{TV}(\hat{\theta}(\underline{Z})||N(0, 2I^{-1}(\theta))) \quad (6)$$

$$= \text{TV}(N(0, I^{-1}(\theta))||N(0, 2I^{-1}(\theta))) + o(1) \quad (7)$$

$$\geq H^2(N(0, I^{-1}(\theta))||N(0, 2I^{-1}(\theta))) + o(1) \quad (8)$$

$$\geq (1 - (.971)^k)^2 + o(1) \quad (9)$$

$$\geq .00084 + o(1), \quad (10)$$

where (4) is by the data processing inequality, the (6) applies two triangle inequalities, (7) uses the asymptotic distributions of $\hat{\theta}(\underline{X})$ and $\hat{\theta}(\underline{Z})$. For (9), we compute the Hellinger distance between two normal distributions: by Pardo [2018], we have the following formula for the Hellinger distance between two normal distributions with the same mean, but difference covariance matrices Σ_1 and Σ_2 :

$$H(N(0, \Sigma_1)||N(0, \Sigma_2)) = 1 - \frac{\det \Sigma_1^{1/4} \det \Sigma_2^{1/4}}{\det(\frac{\Sigma_1 + \Sigma_2}{2})^{1/2}}.$$

Plugging in $\Sigma_1 = I^{-1}(\theta)$ and $\Sigma_2 = 2I^{-1}(\theta)$, we get

$$\begin{aligned}
H(N(0, I^{-1}(\theta)) || N(0, 2I^{-1}(\theta))) &= 1 - \frac{\det(I^{-1}(\theta))^{1/4} \det(2I^{-1}(\theta))^{1/4}}{\det(3/2 I^{-1}(\theta))^{1/2}} \\
&= 1 - \det(I^{-1/4}(\theta) 2^{1/4} I^{-1/4}(\theta) (3/2)^{-1/2} I^{1/2}(\theta)) \\
&= 1 - \det(2^{3/4} 3^{-1/2} I) \\
&= 1 - (2^{3/4} 3^{-1/2})^k \\
&\geq 1 - (.971)^k.
\end{aligned}$$

□

Proof of Theorem 4.2. We will consider for the purposes of this proof that $\hat{\theta}(\underline{X})$ is a randomized statistic, so we can also write $g(\hat{\theta} | \underline{X})$ to represent the distribution of $\hat{\theta}(\underline{X})$ given \underline{X} (which we assume does not depend on θ , since $\hat{\theta}(\underline{X})$ is a statistic). Any deterministic statistic can be expressed as a limit of randomized statistics, where the noise due to randomness goes to zero. For example, Barber and Janson [2020] consider statistics which are solutions to noisy score equations, where the noise is normally distributed and $o_p(n^{-1/2})$. While the distributions g depend on n , we will suppress this dependence for notational simplicity. We can then express the desired KL divergence as follows:

First, by the data processing inequality, we can add in the random variable $\hat{\theta}(\underline{X}) = \hat{\theta}(\underline{Y})$ to get an upper bound on the KL divergence. We then have closed formulas for the joint

distributions $(X_1, \dots, X_n, \hat{\theta}(\underline{X}))$ and $(Y_1, \dots, Y_n, \hat{\theta}(\underline{X}))$.

$$\text{KL}(X_1, \dots, X_n || Y_1, \dots, Y_n) \tag{11}$$

$$\leq \text{KL}(X_1, \dots, X_n, \hat{\theta}(\underline{X}) || Y_1, \dots, Y_n, \hat{\theta}(\underline{X})) \tag{12}$$

$$= \text{KL}\left(f_{\hat{\theta}}^n(\underline{x} | \hat{\theta}(\underline{x}))g_{\theta}(\hat{\theta}(x)) || f_{\theta_n}^n(\underline{x} | \hat{\theta}(\underline{x}))g_{\theta}(\hat{\theta}(\underline{x}))\right) \tag{13}$$

$$= \mathbb{E}_{\hat{\theta} \sim g(\cdot | \underline{X})} \mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{f_{\theta}(\underline{X} | \hat{\theta})g_{\theta}(\hat{\theta})}{f_{\theta_n}(\underline{X} | \hat{\theta})g_{\theta}(\hat{\theta})} \right) \tag{14}$$

$$= \mathbb{E}_{\hat{\theta} \sim g(\cdot | \underline{X})} \mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{f_{\theta}(\underline{X})g(\hat{\theta} | \underline{X})}{f_{\theta_n}(\underline{X} | \hat{\theta})g_{\theta}(\hat{\theta})} \right) \tag{15}$$

$$= \mathbb{E}_{\hat{\theta} \sim g(\cdot | \underline{X})} \mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{f_{\theta}(\underline{X})g(\hat{\theta} | \underline{X})}{[f_{\theta_n}(\underline{X})g(\hat{\theta} | \underline{X})/g_{\theta_n}(\hat{\theta})]g_{\theta}(\hat{\theta})} \right) \tag{16}$$

$$= \mathbb{E}_{\hat{\theta} \sim g(\cdot | \underline{X})} \mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{f_{\theta}(\underline{X})g(\hat{\theta} | \underline{X})}{f_{\theta_n}(\underline{X})g(\hat{\theta} | \underline{X})} \right) + \mathbb{E}_{\hat{\theta} \sim g(\cdot | \underline{X})} \mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{g_{\theta_n}(\hat{\theta})}{g_{\theta}(\hat{\theta})} \right) \tag{17}$$

$$= -\mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{f_{\theta_n}(\underline{X})}{f_{\theta}(\underline{X})} \right) + \mathbb{E}_{\hat{\theta} \sim g_{\theta}} \log \left(\frac{g_{\theta_n}(\hat{\theta})}{g_{\theta}(\hat{\theta})} \right), \tag{18}$$

Above, line (14) simply applies the definition of KL divergence, and line (16) uses the definition of conditional distribution.

At this point, we need to compute the two expectations of line (18), and show that everything cancels except for an $o_p(1)$ term.

We write $\ell(\theta | \underline{x}) = \sum_{i=1}^n \log f_{\theta}(x_i)$. Using our assumptions, we can expand $\ell(\theta_n | \underline{x})$:

$$\begin{aligned} \ell(\theta_n | \underline{x}) &= \ell(\theta | \underline{x}) + (\theta_n - \theta)^{\top} \nabla \ell(\theta | \underline{x}) + \frac{1}{2} (\theta_n - \theta)^{\top} \nabla^2 \ell(\theta | \underline{x}) (\theta_n - \theta) \\ &\quad + \frac{1}{6} \xi^* \sum_{i,j,k} (\theta_n - \theta)_i (\theta_n - \theta)_j (\theta_n - \theta)_k \sum_{s=1}^n g_{ijk}(x_s), \end{aligned}$$

where $|\xi^*| < 1$ and $g_{ijk}(x)$ is an upper bound for $\left| \frac{\partial^3 \ell(\theta|x)}{\partial \theta_i \partial \theta_j \partial \theta_k} \right|$ for a ball $\theta \in N(\theta)$ which exists by (R3). These expansions are based on those from Serfling [1980]. Applying $\mathbb{E}_{\underline{X} \sim f_\theta}$ to this derivation gives

$$\begin{aligned} \mathbb{E}_{\underline{X} \sim f_\theta} \log \left(\frac{f_{\theta_n}(\underline{X})}{f_\theta(\underline{X})} \right) &= 0 - \frac{n}{2} (\theta_n - \theta)^\top I(\theta) (\theta_n - \theta) \\ &\quad + O(1) \frac{n}{6} \sum_{i,j,k} [\mathbb{E} g_{i,j,k}(x)] (\theta_n - \theta)_i (\theta_n - \theta)_j (\theta_n - \theta)_k, \quad (19) \\ &= \frac{-n}{2} (\theta_n - \theta)^\top I(\theta) (\theta_n - \theta) + O(n) \|\theta_n - \theta\|^3 \end{aligned}$$

where the first term is zero as the expected value of the score function is zero by (R3), the second term uses Lehmann [2004, Theorem 7.2.1] and (R3). The $O(1)$ factor in the third term is based on the fact that $|\xi^*| \leq 1$. Finally, note that $\sum_{i,j,k} [\mathbb{E} g_{i,j,k}(x)] (\theta_n - \theta)_i (\theta_n - \theta)_j (\theta_n - \theta)_k \leq p^3 \sup_{i,j,k} [\mathbb{E} g_{i,j,k}(x)] \|\theta_n - \theta\|_\infty^3 = O(1) \|\theta_n - \theta\|^3$. Note that all norms are equivalent in \mathbb{R}^p , so they can be interchanged up to a factor of $O(1)$.

Next, we will derive a similar formula for $\log g_{\theta^*}(\hat{\theta})$:

$$\begin{aligned} \log g_{\theta_n}(\hat{\theta}) &= \log g_\theta(\hat{\theta}) + \nabla \log g_\theta(\hat{\theta}) (\theta_n - \theta) + \frac{1}{2} (\theta_n - \theta)^\top \nabla^2 \log g_\theta(\hat{\theta}) (\theta_n - \theta) \\ &\quad + \frac{n}{6} \xi_2^* \sum_{i,j,k} (\theta_n - \theta)_i (\theta_n - \theta)_j (\theta_n - \theta)_k G_{i,j,k}(\hat{\theta}), \quad (20) \end{aligned}$$

where $|\xi_2^*| \leq 1$. In order to apply the expectation $\mathbb{E}_{\hat{\theta} \sim \theta}$ to this equation, we will first show $\mathbb{E}_{\hat{\theta} \sim \theta} \nabla \log g_\theta(\hat{\theta}) = 0$ and $\mathbb{E}_{\hat{\theta} \sim \theta} \nabla^2 \log g_\theta(\hat{\theta}) = -nI(\theta) + o(n)$.

$$\begin{aligned}
\left[\mathbb{E}_{\hat{\theta} \sim \theta} \nabla \log g_{\theta}(\hat{\theta})\right]_j &= \int \left[\frac{\partial}{\partial \theta_j} \log g_{\theta}(\hat{\theta})\right] g_{\theta}(\hat{\theta}) d\hat{\theta} \\
&= \int \frac{\partial}{\partial \theta_j} g_{\theta}(\hat{\theta}) d\hat{\theta} \\
&= \int \frac{\partial}{\partial \theta_j} \int_x f_{\theta}(x) g(\hat{\theta} | x) dx d\hat{\theta} \\
&= \int_{\hat{\theta}} \int_x \frac{\partial}{\partial \theta_j} f_{\theta}(x) g(\hat{\theta} | x) dx d\hat{\theta} \\
&= \frac{\partial}{\partial \theta_j} \int \int f_{\theta}(x) g(\hat{\theta} | x) dx d\hat{\theta} \\
&= 0,
\end{aligned}$$

where we use the assumption (R3) that $\left|\frac{\partial}{\partial \theta_j} f_{\theta}(x)\right|$ is bounded above by an integrable function and the dominated convergence theorem to interchange the derivative and the integral.

Next we work on the second derivative:

$$\begin{aligned}
\left[\mathbb{E}_{\hat{\theta} \sim \theta} \nabla^2 \log g_{\theta}(\hat{\theta})\right]_{j,k} &= \int \left[\frac{\partial^2}{\partial \theta_j \partial \theta_k} \log g_{\theta}(\hat{\theta})\right] g_{\theta}(\hat{\theta}) d\hat{\theta} \\
&= \int \frac{g_{\theta}(\hat{\theta}) \frac{\partial^2}{\partial \theta_j \partial \theta_k} g_{\theta}(\hat{\theta}) - \frac{\partial}{\partial \theta_j} g_{\theta}(\hat{\theta}) \left[\frac{\partial}{\partial \theta_k} g_{\theta}(\hat{\theta})\right]^{\top}}{g_{\theta}^2(\hat{\theta})} g_{\theta}(\hat{\theta}) d\hat{\theta} \\
&= 0 - \left[\mathbb{E}_{\hat{\theta} \sim \theta} \left(\nabla \log g_{\theta}(\hat{\theta}) \nabla^{\top} \log g_{\theta}(\hat{\theta})\right)\right]_{j,k},
\end{aligned}$$

where we used (R3) along with the dominated convergence theorem to set the first term equal to zero. We see that $\mathbb{E} \nabla^2 \log g_{\theta}(\hat{\theta}) = -I_{\hat{\theta}(X)}(\theta)$, where $I_{\hat{\theta}(X)}$ represents the Fisher information

of the random variable $\hat{\theta}(\underline{X}) \sim g$. It is our current goal to show that $I_{\hat{\theta}(\underline{X})}(\theta) = nI(\theta) + o(n)$, where $I(\theta)$ is the Fisher information for one sample $X \sim f_\theta$. First note that by the data processing inequality [Zamir, 1998], $I_{\hat{\theta}(\underline{X})}(\theta) \leq I_{X_1, \dots, X_n}(\theta) = nI(\theta)$, where the inequality represents the positive-definite ordering of matrices. Next, we need to find a matching lower bound. By the Cramer Rao lower bound, we have that

$$[I_{\hat{\theta}(\underline{X})}(\theta)]^{-1} \leq \text{Var}(\hat{\theta}(\underline{X})) + o(1/n),$$

where $\text{Var}(\hat{\theta}(\underline{X}))$ is the covariance matrix of the random variable $\hat{\theta}(\underline{X})$, and we used the fact that $\hat{\theta}(\underline{X})$ is asymptotically unbiased. By the efficiency of $\hat{\theta}(\underline{X})$, we have that

$$\text{Var}(\hat{\theta}(\underline{X})) = n^{-1}I^{-1}(\theta) + o(1/n).$$

We then have

$$I_{\hat{\theta}(\underline{X})}(\theta) \geq (n^{-1}I^{-1}(\theta) + o(1/n))^{-1} = n[I^{-1}(\theta) + o(1)]^{-1} = n[I(\theta) + o(1)],$$

where for the last equality, we use the following matrix identity:

$$(A + B)^{-1} = A^{-1} - A^{-1}B(A + B)^{-1},$$

where we set $A = I^{-1}(\theta)$ and $B = o(1)$.

Combining our results, we have that

$$\mathbb{E}_{\hat{\theta} \sim \theta} \nabla^2 \log g_\theta(\hat{\theta}) = -I_{\hat{\theta}(\underline{X})}(\theta) = n[-I(\theta) + o(1)].$$

Finally, applying the expectation to equation (20), we have

$$\begin{aligned}
\mathbb{E}_{\hat{\theta} \sim \theta} \log \left(\frac{g_{\theta_n}(\hat{\theta})}{g_{\theta}(\hat{\theta})} \right) &= 0 - \frac{n}{2}(\theta_n - \theta)(I(\theta) + o(1))(\theta_n - \theta) \\
&\quad + O(1) \frac{n}{6} \sum_{i,j,k} [\mathbb{E} G_{i,j,k}(\hat{\theta})] (\theta_n - \theta)_i (\theta_n - \theta)_j (\theta_n - \theta)_k \\
&= \frac{-n}{2} (\theta_n - \theta)^\top I(\theta) (\theta_n - \theta) + o(n) \|\theta_n - \theta\|^2 + O(n) \|\theta_n - \theta\|^3.
\end{aligned} \tag{21}$$

Combining equations (19) and (21), we have

$$\begin{aligned}
&\text{KL}(X_1, \dots, X_n \| Y_1, \dots, Y_n) \\
&\leq \text{KL} \left(f_{\theta}^n(\underline{x} | \hat{\theta}(\underline{x})) g_{\theta}(\hat{\theta}(\underline{x})) \middle\| \middle\| f_{\theta_n}^n(\underline{x} | \hat{\theta}(\underline{x})) g_{\theta}(\hat{\theta}(\underline{x})) \right) \\
&= -\mathbb{E}_{\underline{X} \sim f_{\theta}} \log \left(\frac{f_{\theta_n}(\underline{X})}{f_{\theta}(\underline{X})} \right) + \mathbb{E}_{\hat{\theta} \sim g_{\theta}(\hat{\theta})} \log \left(\frac{g_{\theta_n}(\hat{\theta})}{g_{\theta}(\hat{\theta})} \right) \\
&= o(n) \|\theta_n - \theta\|^2 + O(n) \|\theta_n - \theta\|^3.
\end{aligned}$$

□

Proof of Lemma 4.3. For a fixed $\theta \in \Theta$, for $\omega \sim P$, the random variable $Y = X_{\theta}(\omega)$ is distributed with probability measure PX_{θ}^{-1} : for any measurable set E , $P(Y \in E) = PX_{\theta}^{-1}(E)$. We denote by P_{Ω}^n the joint probability measure on Ω^n , and $(PX_{\theta}^{-1})^n$ the joint probability measure on \mathcal{X}^n .

Given $\theta^* \in \Theta$, our goal is to derive the probability distribution of the random variables $X_{\theta^*}(\omega_1), \dots, X_{\theta^*}(\omega_n)$ conditioned on the event that $\{\omega_1, \dots, \omega_n \mid \hat{\theta}(X_{\theta^*}(\omega_i)) = \hat{\theta}\}$. However, this event may have zero probability. Instead, we will condition on $S_{\hat{\theta}, \theta^*}^{\delta} = \{\omega_1, \dots, \omega_n \mid \hat{\theta}(X_{\theta^*}(\omega)) \in B_{\delta}(\hat{\theta})\}$, where $B_{\delta}(\hat{\theta}) = \{\theta \mid \|\hat{\theta} - \theta\| \leq \delta\}$, which has positive probability. At the end, we will take the limit as $\delta \rightarrow 0$ to derive the desired distribution.

Let $E \subset \mathcal{X}^n$ be a measurable set. Then

$$\begin{aligned}
& P(X_{\theta^*}(\omega_1), \dots, X_{\theta^*}(\omega_n) \in E \mid \omega_1, \dots, \omega_n \in S_{\theta^*, \hat{\theta}}^\delta) \\
&= P(\omega_1, \dots, \omega_n \in X_{\theta^*}^{-1}E \mid \omega_1, \dots, \omega_n \in S_{\theta^*, \hat{\theta}}^\delta) \\
&= \frac{P^n(X_{\theta^*}^{-1}E \cap S_{\theta^*, \hat{\theta}}^\delta)}{P^n(S_{\theta^*, \hat{\theta}}^\delta)} \\
&= \frac{(PX_{\theta^*}^{-1})^n(E \cap X_{\theta^*}S_{\theta^*, \hat{\theta}}^\delta)}{(PX_{\theta^*}^{-1})^n(X_{\theta^*}S_{\theta^*, \hat{\theta}}^\delta)},
\end{aligned}$$

where we used the definition of conditional probability and the fact that $X_{\theta^*}^{-1}X_{\theta^*}S_{\theta^*, \hat{\theta}}^\delta = S_{\theta^*, \hat{\theta}}^\delta$.

This last expression shows that $X_{\theta^*}(\omega_1), \dots, X_{\theta^*}(\omega_n)$ conditioned on $\underline{\omega} \in S_{\theta^*, \hat{\theta}}^\delta$ is distributed as $f_{\theta^*}^n(y_1, \dots, y_n \mid \hat{\theta}(\underline{y}) \in B_\delta(\hat{\theta}))$. This derivation is valid for all $\delta > 0$. Taking the limit as $\delta \rightarrow 0$ gives the desired formula:

$$Y_1^{\theta^*}, \dots, Y_n^{\theta^*} \mid \hat{\theta}(\underline{Y}^{\theta^*}) = \hat{\theta}(\underline{X}) \sim f_{\theta^*}^n(y_1, \dots, y_n \mid \hat{\theta}(\underline{y}) = \hat{\theta}(\underline{X})).$$

□

Proof of Lemma 4.4. First note that for all $\theta \in \Theta$, we have that $\hat{\theta}(X_\theta(\underline{\omega})) = \theta + O_p(n^{-1/2})$, by the efficiency of $\hat{\theta}$. In particular, we have that $\hat{\theta}(X_{\theta^*}(\underline{\omega})) = \theta^* + O_p(n^{-1/2})$. We also have that $\hat{\theta}(X_{\theta^*}(\underline{\omega})) = \hat{\theta}(X) = \theta + O_p(n^{-1/2})$. Combining these facts, we have that $\theta^* = \theta + O_p(n^{-1/2})$. □

Parts 1 and 2 of Lemma 7.4 can be rephrased as the following: $\hat{\theta}$ is efficient if and only if it is consistent and $n^{-1} \sum_{i=1}^n S(\hat{\theta}, X_i) = o_p(n^{-1/2})$. The third property of Lemma 7.4 is similar to many standard expansions used in asymptotics, for example in Van der Vaart [2000]. However, we require the expansion for arbitrary efficient estimators, and include a

proof for completeness.

Lemma 7.4. *Suppose $X_1, \dots, X_n \stackrel{i.i.d.}{\sim} f_{\theta_0}$, and assume that (R1)-(R3) hold. Let $\hat{\theta}$ be an efficient estimator, which is a sequence of zeros of the score equations. Suppose that $\tilde{\theta}$ is a \sqrt{n} -consistent estimator of θ_0 . Then*

1. *If $n^{-1} \sum_{i=1}^n S(\tilde{\theta}, X_i) = o_p(n^{-1/2})$, then $\tilde{\theta} - \hat{\theta} = o_p(n^{-1/2})$.*
2. *If $\tilde{\theta}$ is efficient, then $n^{-1} \sum_{i=1}^n S(\tilde{\theta}, X_i) = o_p(n^{-1/2})$.*
3. *If $\tilde{\theta}$ is efficient, then $\tilde{\theta} = \theta_0 + I^{-1}(\theta_0)n^{-1} \sum_{i=1}^n S(\theta_0, X_i) + o_p(n^{-1/2})$.*

Proof. As $\tilde{\theta}$ and $\hat{\theta}$ are both \sqrt{n} -consistent, we know that $\tilde{\theta} - \hat{\theta} = O_p(n^{-1/2})$. So, we may consider a Taylor expansion of the score function about $\tilde{\theta} = \hat{\theta}$.

$$\begin{aligned}
n^{-1} \sum_{i=1}^n S(\tilde{\theta}, X_i) &= n^{-1} \sum_{i=1}^n S(\hat{\theta}, X_i) + \left(\frac{d}{d\theta} n^{-1} \sum_{i=1}^n S(\hat{\theta}, X_i) \right) (\tilde{\theta} - \hat{\theta}) + O_p(n^{-1}) \\
&= 0 + \left[\frac{d}{d\theta} n^{-1} \sum_{i=1}^n S(\hat{\theta}, X_i) + O_p(n^{-1/2}) \right] (\tilde{\theta} - \hat{\theta}) \\
&= [-I(\theta_0) + o_p(1)] (\tilde{\theta} - \hat{\theta}),
\end{aligned} \tag{22}$$

where we used assumptions (R1)-(R3) to justify that 1) the second derivative is bounded in a neighborhood about θ_0 (as both $\hat{\theta}$ and $\tilde{\theta}$ converge to θ_0), 2) the derivative of the score converges to $-I(\theta_0)$ by Lehmann [2004, Theorem 7.2.1] along with the Law of Large Numbers, and 3) that $I(\theta_0)$ is finite, by (R3).

To establish property 1, note that the left hand side of Equation (22) is $o_p(n^{-1/2})$ implying that $(\tilde{\theta} - \hat{\theta}) = o_p(n^{-1/2})$. For property 2, recall that by Lehmann [2004, Page 479], if $\tilde{\theta}$ and $\hat{\theta}$ are both efficient, then $(\tilde{\theta} - \hat{\theta}) = o_p(n^{-1/2})$. Plugging this into the right hand side of Equation (22) gives $n^{-1} \sum_{i=1}^n S(\tilde{\theta}, X_i) = o_p(n^{-1/2})$, establishing property 2.

For property 3, we consider a slightly different expansion:

$$\begin{aligned}
o_p(n^{-1/2}) &= n^{-1} \sum_{i=1}^n S(\tilde{\theta}, X_i) \\
&= n^{-1} \sum_{i=1}^n S(\theta_0, X_i) + \frac{d}{d\theta_0} n^{-1} \sum_{i=1}^n S(\theta_0, X_i) (\tilde{\theta} - \theta_0) + O_p(n^{-1}), \\
&= n^{-1} \sum_{i=1}^n S(\theta_0, X_i) + (-I(\theta_0) + o_p(1)) (\tilde{\theta} - \theta_0) + O_p(n^{-1})
\end{aligned}$$

where we used property 2 for the first equality, expanded the score about $\hat{\theta} = \theta_0$ for the second, and justify the $O_p(n^{-1})$ by (R2). By (R1)-(R2) and Law of Large Numbers along with Lehmann [2004, Theorem 7.2.1], we have the convergence of the derivative of score to $-I(\theta_0)$. By (R3), $I(\theta_0)$ is invertible. Solving the equation for $\tilde{\theta}$ gives the desired result. \square

Lemma 7.5. *Assume that (R0)-(R3) hold, and let $\omega_1, \dots, \omega_n \stackrel{i.i.d.}{\sim} P$. Then*

$$n^{-1} \sum_{i=1}^n \frac{d}{d\theta} S(\theta, X_{\theta}(\omega_i)) = o_p(1).$$

Proof. First we can express the derivative as

$$n^{-1} \sum_{i=1}^n \frac{d}{d\theta} S(\theta, X_{\theta}(\omega_i)) = n^{-1} \sum_{i=1}^n \left(\frac{d}{d\alpha} S(\alpha, X_{\theta}(\omega_i)) + \frac{d}{d\alpha} S(\theta, X_{\alpha}(\omega_i)) \right) \Big|_{\alpha=\theta}.$$

The result follows from the Law of Large Numbers, provided that

$$\mathbb{E}_{\omega \sim P} \left(\frac{d}{d\alpha} S(\alpha, X_{\theta}(\omega)) + \frac{d}{d\alpha} S(\theta, X_{\alpha}(\omega)) \right) \Big|_{\alpha=\theta} = 0.$$

The expectation of the first term is $-I(\theta)$, by Lehmann [2004, Theorem 7.2.1]. For the

second term, we compute

$$\begin{aligned}
\mathbb{E}_{\omega \sim P} \frac{d}{d\alpha} S(\theta, X_\alpha(\omega)) \Big|_{\alpha=\theta} &= \int_{\Omega} \frac{d}{d\alpha} S(\theta, X_\alpha(\omega)) \Big|_{\alpha=\theta} dP(\omega) \\
&= \int_{\mathcal{X}} \frac{d}{d\alpha} S(\theta, x) f_\alpha(x) \Big|_{\alpha=\theta} d\mu(x) \\
&= \int_{\mathcal{X}} S(\theta, x) \left(\frac{d}{d\alpha} f_\alpha(x) \Big|_{\alpha=\theta} \right)^\top d\mu(x) \\
&= \int_{\mathcal{X}} S(\theta, x) \left(\frac{\frac{d}{d\theta} f_\theta(x)}{f_\theta(x)} \right)^\top f_\theta(x) d\mu(x) \\
&= \int_{\mathcal{X}} S(\theta, x) S^\top(\theta, x) f_\theta(x) d\mu(x) \\
&= \mathbb{E}_{X \sim \theta} [S(\theta, X) S^\top(\theta, X)] \\
&= I(\theta). \quad \square
\end{aligned}$$

Proof of Theorem 4.7. We expand $\hat{\theta}(\underline{Z})$ about $\hat{\theta}(\underline{X})$ using part 3 of Lemma 7.5:

$$\hat{\theta}(\underline{Z}) = \hat{\theta}(\underline{X}) + I^{-1}(\hat{\theta}(\underline{X})) n^{-1} \sum_{i=1}^n S(\hat{\theta}(\underline{X}), X_{\hat{\theta}(\underline{X})}(\omega_i)) + o_p(n^{-1/2}) \quad (23)$$

The score can be expanded about $\hat{\theta}(\underline{X}) = \theta_0$:

$$\begin{aligned}
&n^{-1} \sum_{i=1}^n S(\hat{\theta}(\underline{X}), X_{\hat{\theta}(\underline{X})}(\omega_i)) \\
&= n^{-1} \sum_{i=1}^n S(\theta_0, X_{\theta_0}(\omega_i)) + \left(\frac{d}{d\theta^*} n^{-1} \sum_{i=1}^n S(\theta^*, X_{\theta^*}(\omega_i)) \right) (\hat{\theta}(\underline{X}) - \theta_0) \\
&= n^{-1} \sum_{i=1}^n S(\theta_0, X_{\theta_0}(\omega_i)) + o_p(1) O_p(n^{-1/2}),
\end{aligned}$$

where θ^* is between $\hat{\theta}(\underline{X})$ and θ_0 ; by Lemma 7.5, we justify that the derivative is $o_p(1)$.

Combining this derivation along with the fact that $I^{-1}(\hat{\theta}(\underline{X})) = I^{-1}(\theta_0) + o_p(1)$ by the continuous mapping theorem, we have the following equation:

$$\hat{\theta}(\underline{Z}) = \hat{\theta}(\underline{X}) + I^{-1}(\theta_0)n^{-1} \sum_{i=1}^n S(\theta_0, X_{\theta_0}(\omega_i)) + o_p(n^{-1/2}). \quad (24)$$

Using the same techniques, we do an expansion for $\hat{\theta}(\underline{Y})$ about $\theta_1^* = 2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z})$:

$$\hat{\theta}(\underline{Y}) = \theta_1^* + I^{-1}(\theta_1^*)n^{-1} \sum_{i=1}^n S(\theta_1^*, X_{\theta_1^*}(\omega_i)) + o_p(n^{-1/2}) \quad (25)$$

$$= \theta_1^* + I^{-1}(\theta_0)n^{-1} \sum_{i=1}^n S(\theta_0, X_{\theta_0}(\omega_i)) + o_p(n^{-1/2}) \quad (26)$$

$$= \theta_1^* + [\hat{\theta}(\underline{Z}) - \hat{\theta}(\underline{X})] + o_p(n^{-1/2}) \quad (27)$$

$$= \hat{\theta}(\underline{X}) + o_p(n^{-1/2}), \quad (28)$$

where line (26) is a similar expansion as used for equation (23), in line (27) we substituted the expression from (24), and line (28) uses the fact that as $n \rightarrow \infty$, $\theta_1^* = 2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z})$ with probability tending to one. Indeed, since $2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z})$ is a consistent estimator of θ_0 , we have that as $n \rightarrow \infty$, $P(2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z}) \in \Theta) \geq P(2\hat{\theta}(\underline{X}) - \hat{\theta}(\underline{Z}) \in B(\theta_0)) \rightarrow 1$. \square