TESTING GOODNESS-OF-FIT AND CONDITIONAL INDEPENDENCE WITH APPROXIMATE CO-SUFFICIENT SAMPLING

BY RINA FOYGEL BARBER[1,a] AND LUCAS JANSON[2,b]

[1]*Department of Statistics, University of Chicago,* [a]*rina@uchicago.edu*

[2]*Department of Statistics, Harvard University,* [b]*ljanson@fas.harvard.edu*

Goodness-of-fit (GoF) testing is ubiquitous in statistics, with direct ties to model selection, confidence interval construction, conditional independence testing, and multiple testing, just to name a few applications. While testing the GoF of a simple (point) null hypothesis provides an analyst great flexibility in the choice of test statistic while still ensuring validity, most GoF tests for composite null hypotheses are far more constrained, as the test statistic must have a tractable distribution over the entire null model space. A notable exception is *co-sufficient sampling* (CSS): resampling the data conditional on a sufficient statistic for the null model guarantees valid GoF testing using any test statistic the analyst chooses. But CSS testing requires the null model to have a compact (in an information-theoretic sense) sufficient statistic, which only holds for a very limited class of models; even for a null model as simple as logistic regression, CSS testing is powerless. In this paper, we leverage the concept of *approximate sufficiency* to generalize CSS testing to essentially any parametric model with an asymptotically efficient estimator; we call our extension "approximate CSS" (aCSS) testing. We quantify the finite-sample Type I error inflation of aCSS testing and show that it is vanishing under standard maximum likelihood asymptotics, for any choice of test statistic. We apply our proposed procedure both theoretically and in simulation to a number of models of interest to demonstrate its finite-sample Type I error and power.

1. Introduction. Suppose we observe data X belonging to some sample space $\mathcal{X}$, and would like to test whether it comes from some parametric null model $\{P_\theta : \theta \in \Theta\}$, where $\Theta \subseteq \mathbb{R}^d$, versus a more complex (usually higher-dimensional) model. This problem of so-called "goodness-of-fit" (GoF) testing is one of the most fundamental in statistics, with a vast literature exhibiting applications and theoretical and methodological development. We pause here to highlight a few of the many areas of statistics to which GoF testing is directly applicable, including some problems that are not obviously or commonly associated with GoF.

PROBLEM DOMAIN 1 (Standard goodness-of-fit testing). GoF testing is commonly used to test a postulated model or distributional property, often as a precursor to further statistical analysis that assumes the postulated model/property to be correct. Such null models/properties include standard distributional families, nonparametric properties such as symmetry or log-concavity, time-series properties such as stationarity, and relational properties such as independence.

PROBLEM DOMAIN 2 (Model selection). GoF testing can also be used to select a best-fitting model through simultaneously testing a family of models. For instance, this could be

choosing a sparse model in regression, selecting the number of clusters or principle components in unsupervised learning, or identifying change points in a time series.

PROBLEM DOMAIN 3 (Confidence interval construction). Suppose the data X is distributed according to a known model $\{P_{\gamma,\theta} : (\gamma, \theta) \in \Gamma \times \Theta\}$, where $\Gamma \subseteq \mathbb{R}^m$ and $\Theta \subseteq \mathbb{R}^d$, and the goal is to construct a confidence region for γ in the presence of the nuisance parameter θ. If for any γ_0, we can construct a GoF test for the null model $\{P_{\gamma_0,\theta} : \theta \in \Theta\}$, then the set of γ_0 at which we fail to reject the test constitutes a valid confidence region.

PROBLEM DOMAIN 4 (Conditional independence testing). In many regression and graphical modeling problems, the primary question of interest is whether a given triple of random variables (X, Y, Z) satisfies conditional independence, that is, $Y \perp\!\!\!\perp X|Z$. If $X|Z$ is distributed according to a known model $\{P_\theta(\cdot|Z) : \theta \in \Theta\}$, then testing conditional independence can be formulated as a GoF test where $\{P_\theta(\cdot|Z) : \theta \in \Theta\}$ is the null model for $X|Z, Y$. (Note that, when we apply a GoF test to the conditional independence problem in this way, we implicitly treat Y and Z as fixed, and check for goodness-of-fit of X's conditional model.)

PROBLEM DOMAIN 5 (Multiple testing). In multiple testing, the goal is to reject a subset of a fixed family of null hypotheses. When each hypothesis test is a GoF test (i.e., its null is lower-dimensional than its alternative), testing any intersection of null hypotheses (i.e., testing the "global null" for any subset of hypotheses) constitutes a GoF test as well, and combining such intersection GoF tests through a closed testing procedure [38] produces a subset to reject that controls the familywise error rate.

The key challenges of any GoF testing problem are to find a test that is valid, in that it controls the Type I error at a prespecified significance level, and that is powerful, in that it rejects the null model as often as possible when it does not hold. For parametric null models (the focus of this paper), there are many existing methods for testing GoF, with canonical choices including the popular score, likelihood ratio, and Wald tests. The standard approach for these tests and many others is to prescribe a test statistic (chosen to be powerful under a given alternative model) and derive a (often asymptotic) null distribution for it. Such tests require certain regularity conditions on the alternative model (in order to construct a well-behaved test statistic) and on the null model (in order to establish the validity of the null distribution for the test statistic) that are generally quite similar to those needed for the maximum likelihood estimator under both the null and alternative to be asymptotically normal. While these tests are extremely popular and have been fruitfully applied through much of the history of statistics to many problems, the regularity assumptions placed on the alternative distribution in particular limit the ability to fully leverage domain knowledge to maximize the statistical power. To elaborate, consider the following cases which often arise in practice when applying parametric GoF tests.

- *Some prior information is available about the relative plausibility of different regions of the alternative model.* Ideally we would like to incorporate this prior information into our test statistic in order to maximize power (e.g., through a test statistic derived from Bayesian inference), but standard GoF tests only provide the null distribution for a test statistic which is determined by the entire alternative space, and give little flexibility to incorporate prior knowledge into that test statistic while still retaining a valid null distribution. An extreme case would be when certain regions of the alternative are known to be completely implausible, that is, we would like to remove them from the alternative model entirely, yet removing them from the model would violate the regularity conditions required for the alternative model. For example, we may know that under the alternative some k-dimensional

parameter has at most $d < k$ nonzero entries, but we do not know which ones. Such a sparse alternative model would violate the usual assumption that the parameter space is convex, forcing one to ignore the sparsity and instead operate under (and hence, derive a test statistic from) the larger, mostly implausible, k-dimensional alternative model. As we see in the next scenario, if k is too large, even this route is not feasible.

- *The alternative model is high- or even infinite-dimensional.* Since standard GoF tests treat their prescribed test statistics as fixed in the asymptotic regime in which they prove validity, those test statistics can only be designed to be powerful against fixed- (and finite-) dimensional alternatives. If we have a high-dimensional alternative (i.e., whose dimension is not assumed negligible relative to the sample size, which includes any nonparametric alternative), we would ideally like to choose a test statistic which changes with the sample size to be powerful against a sequence of alternatives which changes as the sample size grows asymptotically. But standard GoF tests cannot accommodate such a choice, forcing one to instead operate under a fixed-dimensional alternative that may represent a vanishing fraction of plausible alternatives, or a very coarse approximation to the space of realistic alternatives.

The common thread in these cases is that the test statistic that would be most powerful to use given the domain knowledge at hand is not accompanied by theoretical guarantees or any known (exact or approximate) null distribution. In our simulations in Section 4, we will study some examples where standard tests can be applied (and will compare aCSS to the score test for those examples), and others where, as in the scenarios above, standard tests cannot be applied and thus a more flexible method like aCSS is necessary.

However, constructing a valid test around an arbitrary test statistic $T(X)$ is possible only in very limited settings. In particular, if Θ is simple, that is, it contains only a single point so that there are no unknown/nuisance parameters in the null model, then any test statistic $T(X)$'s null distribution can be arbitrarily well approximated computationally by repeatedly independently sampling copies $\widetilde{X}$ of X from the null distribution and recomputing the same test statistic on the copies. To be concrete, if the statistic $T(X)$ is chosen such that larger (positive) values are seen as evidence against the null, we can draw M i.i.d. copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ from the null distribution, and define a (discretized) p-value

$$(1.1) \qquad \text{pval} = \text{pval}_T\left(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}\right) = \frac{1}{M+1}\left(1 + \sum_{m=1}^{M} \mathbb{1}\{T(\widetilde{X}^{(m)}) \geq T(X)\}\right),$$

which is guaranteed to satisfy $\mathbb{P}(\text{pval} \leq \alpha) \leq \alpha$ under the null, for any predefined rejection level α.

More generally, when Θ is not a singleton set (i.e., the null hypothesis is composite), in principle we can still construct a p-value of the form (1.1) as long as we are able to sample a set of copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ of X so that $X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ are *exchangeable* under the null. We emphasize that this exchangeability property continues to enable the analyst to use any desired test statistic $T(X)$, as the validity of the p-value is unaffected. Of course, to achieve high power, we should aim to choose a function $T(X)$ that is likely to be large under the alternative hypothesis. Note that we absorb everything that is not X into the definition of the function T, for example, for testing conditional independence $X \perp\!\!\!\perp Y | Z$ as in Problem Domain 4, T can depend arbitrarily on Y and Z as well (since, after conditioning on Y and Z, they are treated as fixed and nonrandom).

To summarize, we have seen that

the problem of GoF testing with arbitrary test statistics can be reduced to

one of sampling copies of X that are exchangeable under the null.

These copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ then act as a "control group" for the real data X, and we can compare the real statistic $T(X)$ against the "control group" values $T(\widetilde{X}^{(m)})$ to test the null. Of course, aside from the setting of a simple null, sampling exchangeable copies is not necessarily a straightforward task. In particular, in order to have power, we must ensure our null-exchangeable copies do not remain exchangeable under the alternative; for instance sampling $\widetilde{X}^{(1)} = \cdots = \widetilde{X}^{(M)} = X$ trivially satisfies the exchangeability property under the null, but also under any alternative, and hence clearly equation (1.1) produces a useless p-value of 1 with probability 1 (for any choice of T).

One way to sample exchangeable copies when Θ is composite is by conditioning on a sufficient statistic $S(X)$ for θ, since then by definition the conditional distribution $X|S(X)$ does not depend on θ. By drawing the copies $\widetilde{X}^{(m)}$ from this conditional distribution, we achieve exchangeability of $X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ (more concretely, X and its copies are i.i.d. conditional on $S(X)$). This approach is known as *co-sufficient sampling* (CSS) [42]. However, this approach is viable in only a limited range of settings. In particular, many null models do not admit a compact (in an information-theoretic sense) sufficient statistic, meaning any sufficient statistic for the null model will remain sufficient for many alternative models as well. In such cases, which we term *degenerate*, CSS testing runs into the problem described at the end of the preceding paragraph—the copies $\widetilde{X}^{(m)}$ will still be exchangeable with X under the alternative, resulting in a powerless test. This situation arises quite often, and we will give a number of common examples shortly in Section 1.2.

As an alternative approach, we might consider the *parametric bootstrap* [19], where after estimating the true parameter θ via some estimator $\widehat{\theta}$ (e.g., the maximum likelihood estimate), the copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ are sampled from $P_{\widehat{\theta}}$. While this widely used approach often works well in practice, the parametric bootstrap does not create exchangeable copies of the data, and is not guaranteed to achieve the desired Type I error level when paired with an arbitrary test statistic T—in fact, it may even lead to substantial error inflation. To take a simple example, consider a Gaussian linear regression setting where P_θ is given by the distribution $X \sim \mathcal{N}(\theta \cdot Z, \mathbf{I}_d)$, where $Z \in \mathbb{R}^n$ is a fixed covariate. Suppose we are interested in testing whether X in fact has more dependence with another covariate $Y \in \mathbb{R}^n$, and so our test statistic is given by

$$T(X) = \frac{(X^\top Y)^2}{(X^\top Z)^2}.$$

Figure 1 compares the parametric bootstrap against co-sufficient sampling (full details for this simulation are given in Appendix E in the Supplementary Material [5]). The results show that CSS results in a uniform distribution of p-values, while the parametric bootstrap results in a highly nonuniform distribution, and could lead to substantially inflated Type I error. Therefore, we would instead prefer to extend the CSS framework in order to enjoy theoretically guaranteed error control.

1.1. *Our contribution.*

In this paper, we demonstrate how to escape the problem of zero power in the degenerate setting, by introducing a new generalization of CSS testing that conditions only on an approximately sufficient statistic [28, 29, 44]. We call such a test an "approximate co-sufficient sampling" test. This paper makes four main contributions:

1. We propose *approximate co-sufficient sampling* (aCSS), which samples approximately exchangeable copies of the data by conditioning on an approximately sufficient statistic and plugging in a consistent estimator for the unknown parameter.
2. Under weak conditions closely resembling those for standard maximum likelihood asymptotics, we provide a finite-sample upper-bound for the total variation (TV) distance from exchangeability of our aCSS samples.
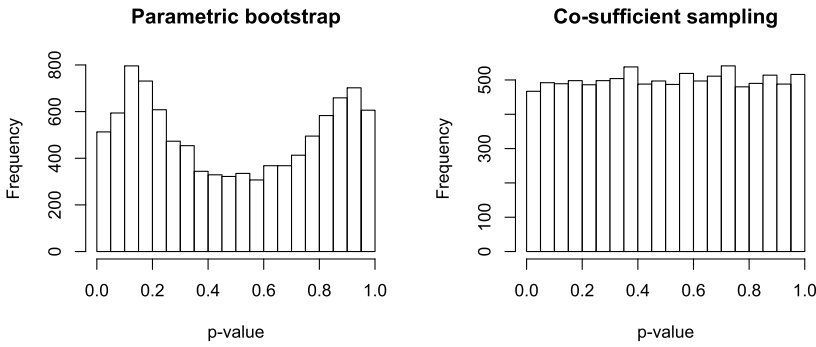
Parametric bootstrap **Co-sufficient sampling**

FIG. 1. *Comparison of the parametric bootstrap versus CSS, in a Gaussian linear model example where the null hypothesis is true. We can see that CSS yields uniformly distributed p-values, but the parametric bootstrap does not, resulting in an inflated Type I error rate. (See Section 1 for details.)*

3. We show that the aforementioned TV bound translates directly to a bound on the Type I error inflation of an aCSS test that holds uniformly over the choice of test statistic, and we apply this bound to a number of important models to prove the inflation vanishes asymptotically as the sample size approaches infinity.

4. We provide general algorithms for aCSS and demonstrate their use in a series of simulations that exhibit the validity and power of aCSS testing.

1.2. *Applications.* The problem of zero power for CSS testing arises surprisingly often— while we call such settings "degenerate", they are not extreme cases but rather constitute a large portion of common statistical models of interest. To illustrate this, we will consider the following settings in which CSS testing is powerless, while aCSS testing can still be quite powerful and remains asymptotically valid for any test statistic. (Our theoretical results later on will quantify its finite-sample Type I error.)

MODEL CLASS 1 (Data with associated covariates). Suppose $X = (X_1, \dots, X_n)$ where the X_i's are independent, and each X_i has an associated covariate Z_i (i.e., the null distribution of each X_i depends on this Z_i). In this setting, the distribution of X (i.e., the joint distribution of $X_1, \dots, X_n$) will often have the data X itself as a minimal sufficient statistic. This is even true when $X_i | Z_i$ follows logistic regression: for generic values of the Z_i's (e.g., if each $Z_i \in \mathbb{R}^d$ is drawn from some continuous distribution), the minimal sufficient statistic is equivalent to X itself because the value of $Z^\top X \in \mathbb{R}^d$ determines $X \in \{0, 1\}^n$ uniquely, and hence is sufficient under any alternative as well (and hence any CSS test must be powerless). This problem class applies not just for GoF testing for conditional models for X (including the logistic model), but also for conditional independence testing as described in Problem Domain 4.

MODEL CLASS 2 (Curved exponential families). Consider a null model that is a curved exponential family, that is, a full-rank k-parameter exponential family with an added constraint that reduces the dimension of the parameter space to some $d < k$ and is nonlinear in the canonical parameters. In this setting, the minimal sufficient statistic is generally the same as that for the unconstrained (full-rank) exponential family. This means that any CSS test must be powerless against any alternative that lies in the larger exponential family, for example, if we want to test whether the parameter constraint holds or not. A classical example is the Behrens–Fisher problem of testing equality of (unknown) means between independent normal samples having different (unknown) variances: any CSS test will be powerless for every alternative pair of means and variances. The same issue arises in, for example, the

study of contingency tables (where the canonical family is multinomial and the null hypothesis imposes a constraint on its probabilities), and spatial and time-series models (where the null hypothesis imposes a spatial or temporal structure on the canonical parameters of an exponential family).

MODEL CLASS 3 (Heavy-tailed models). Many heavy-tailed models are not exponential families and do not admit compact sufficient statistics. For instance, the Cauchy location family's minimal sufficient statistic is given by the order statistics. Any CSS test of this null is therefore powerless against any i.i.d. alternative, since the order statistics are sufficient for this alternative as well. As another example, the Student's t scale family's minimal sufficient statistic is the order statistics of the absolute values, so any CSS test for it must be powerless against any i.i.d. symmetrical alternative.

MODEL CLASS 4 (Models with latent variables). Many popular models capture domain-specific properties through latent variables. In such models, if we condition on the latent variable, then the data often comes from a well-behaved distribution with compact minimal sufficient statistic. However, when the latent variable is unobserved, we are forced to perform inference unconditionally, and the unconditional model rarely has a compact minimal sufficient statistic. Examples include hidden Markov models, mixture distributions, data with missing values, errors-in-variables models, and factor models.

Later on in Section 4, we will return to Model Classes 1, 2, and 3 and give concrete examples of models where aCSS can be applied. We leave Model Class 4 for future work.

1.3. *Related work.* GoF testing dates back to the very early days of the field of statistics, and the literature even on just parametric GoF is far too numerous to cite. Instead, we focus our literature review on the subfield of CSS testing, which distinguishes itself within the broader field of parametric GoF testing by guaranteeing Type I error control with *any* test statistic under a parametric null model, the potential advantages of which have been described earlier in this section (though see [23] for a recent work on GoF testing in generalized linear models that allows a fair degree of test statistic flexibility). In fact, some of the most foundational nonparametric tests can be thought of as CSS tests, including the permutation test (conditions on the order statistics for an i.i.d. null). The formal idea of a CSS test seems to date back to at least [6], although the value of sufficient statistics for GoF testing in the presence of nuisance parameters has also been used in many other ways, for example, [8, 18, 27] decompose the data into a minimal sufficient statistic and an ancillary statistic and construct a GoF test based on the parameter-free distribution of the ancillary statistic.

CSS testing has gained substantial interest in the last 30 years, though with a focus on nondegenerate hypothesis testing settings [1, 2, 9, 13, 15, 20, 21, 33–37, 39, 41, 42]. Our work differs from the existing work in CSS testing by allowing for degenerate (and nondegenerate) models by conditioning only on an approximately sufficient statistic. Similar techniques have been used to obtain exact confidence intervals in the presence of nuisance parameters [32], again in nondegenerate settings. As an example, [25, 40] study conditional independence testing of $X \perp\!\!\!\perp Y | Z$ where $X | Z$ follows a logistic regression model and Z is discrete; logistic regression is degenerate when Z has a continuous distribution but nondegenerate when Z is discrete.

For conditional independence testing (Problem Class 4), in the setting where H_0 is a simple null hypothesis (i.e., $X | Z$ has a known distribution), [14] study procedures of the form (1.1) under the name "conditional randomization test"; their work also constructs the model-X knockoffs framework for simultaneously testing multiple conditional independence hypotheses (i.e., variable selection in a multivariate regression). This construction provides an exact

"swap-exchangeability" property that enables variable selection, that is, simultaneous conditional independence testing of many covariates when the multivariate covariates X come from a nondegenerate model, and leads to exact false discovery rate control [4]. Generalizing to the setting where H_0 is not simple, [22] construct model-X knockoffs [14] conditional on a sufficient statistic, retaining the exact "swap-exchangeability" property and exact false discovery rate control; we can think of this as a knockoffs-analogue of CSS testing.

Moving beyond CSS testing, we are only aware of a few works which take a similarly approximate approach to that of the present paper. First, the most related to our approach is the work of [31], where they mention the possibility of an aCSS-type test to solve the Behrens–Fisher problem (i.e., testing for a difference of means between two Gaussian samples, described above in Model Class 2), but conclude that such an approach would be computationally intractable; they instead propose a heuristic sampling procedure which they support with simulations but no theory. Second, [16, 24] focus on parametric likelihood-based testing in the presence of nuisance parameters, but study the case where the nuisance parameters are orthogonal or asymptotically orthogonal to the parameter of interest. Third, approximate Bayesian computation, also known as likelihood-free inference, conducts Bayesian inference conditional on a compact nonsufficient statistic, but for computational, as opposed to statistical reasons, since in the Bayesian framework there is no statistical downside to conditioning on as much as possible (see, e.g., [26] for such a paper that explicitly addresses the role of sufficiency). Finally, another related Bayesian approach is the *conditional predictive p-value* method of [7], where to sample new copies of X, they need to specify a prior on the unknown parameter θ; uncertainty about the correct prior motivates their proposed method, which essentially replaces the prior on θ with the posterior distribution of θ given some approximately sufficient statistic $S(X)$, to reduce the impact of a potentially misspecified prior.

1.4. *Notation.* We will write $\|\cdot\|$ to denote the usual Euclidean (ℓ_2) norm on vectors, and to denote the operator norm (i.e., spectral norm) on matrices. For a matrix M, $\lambda_{\max}(M)$ denotes its largest eigenvalue in the positive direction. We write $(x)_+$ to denote $\max\{x, 0\}$. We will write $\mathbb{E}_\theta$ and $\mathbb{P}_\theta$ to denote expectation or probability taken with respect to $X \sim P_\theta$, where the parametric family $\{P_\theta : \theta \in \Theta\}$ is our null model.

2. Method. The goal of approximate co-sufficient sampling is to generate copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ of the observed data X such that if the null hypothesis

$$(2.1) \qquad\qquad H_0 : X \sim P_\theta \quad \text{for some } \theta \in \Theta$$

is true, then the random variables $X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ are approximately exchangeable. Recalling the p-value defined in (1.1), we can then test the null hypothesis using any desired test statistic $T(X)$. The choice of statistic is completely unconstrained, and this flexibility enables us to design very powerful tests in many settings. Note that, although this flexibility allows us to design any form of function T, T itself (as a function, i.e., before seeing its argument) cannot depend on X. For example, if $T(X)$ uses X to tune parameters in a neural network and then computes a statistic of that neural network applied to X, then $T(\widetilde{X}^{(m)})$ cannot compute a statistic on the same (X-tuned) neural network applied to $\widetilde{X}^{(m)}$, but must use $\widetilde{X}^{(m)}$ to tune the parameters of a new neural network and compute a statistic of that ($\widetilde{X}^{(m)}$-tuned) neural network applied to $\widetilde{X}^{(m)}$.

To quantify our goal of generating approximately exchangeable copies of the data, we begin by defining a "distance to exchangeability":

DEFINITION 1. For any integer $k \geq 1$ and any set of random variables $A_1, \ldots, A_k$ with a joint distribution, define

$$\mathsf{d}_{\mathsf{exch}}(A_1, \ldots, A_k) = \inf\{\mathsf{d}_{\mathsf{TV}}((A_1, \ldots, A_k), (B_1, \ldots, B_k)) : B_1, \ldots, B_k \text{ are exchangeable}\}.$$

Here d_{TV} denotes the total variation distance, and the infimum is taken over all sets of k random variables $B_1, \ldots, B_k$ with an exchangeable joint distribution.

Of course, if $A_1, \ldots, A_k$ are exchangeable, then $\mathsf{d}_{\mathsf{exch}}(A_1, \ldots, A_k) = 0$. When we say informally that variables $A_1, \ldots, A_k$ are "approximately exchangeable", we mean that the distance to exchangeability is small.

Now we will see how this distance $\mathsf{d}_{\mathsf{exch}}$ relates to the problem of testing the null hypothesis (2.1) ([10] use a similar argument in a permutation test setting). Fix a threshold $\alpha \in [0, 1]$ and a function $T : \mathcal{X} \to \mathbb{R}$ (the test statistic). For any exchangeable random variables $(B_0, \ldots, B_M)$, by definition of exchangeability we have $\mathbb{P}(\mathrm{pval}_T(B_0, B_1, \ldots, B_M) \leq \alpha) \leq \alpha$, and therefore,

$$\mathbb{P}\big(\mathrm{pval}_T\big(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}\big) \leq \alpha\big) \leq \alpha + \mathsf{d}_{\mathsf{TV}}\big(\big(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}\big), (B_0, B_1 \ldots, B_M)\big).$$

Taking an infimum over all exchangeable distributions on $(B_0, B_1 \ldots, B_M)$, we have shown that

$$\mathbb{P}\big(\mathrm{pval}_T\big(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}\big) \leq \alpha\big) \leq \alpha + \mathsf{d}_{\mathsf{exch}}\big(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}\big)$$

under the null hypothesis H_0.

Therefore, we can see that, if we are able to construct copies of the data such that $\mathsf{d}_{\mathsf{exch}}(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)})$ is small, then we can construct an approximately valid test of H_0 using any desired test statistic T. From this point on, then, our task is to determine how we can use approximate sufficiency to generate such copies.

2.1. *Overview.* Consider any function $S = S(X)$ of the data, which is sufficient under the null hypothesis that $X \sim P_\theta$ for some $\theta \in \Theta$. Let $P(\cdot|s)$ be the conditional distribution of X given $S = s$ (sufficiency of $S(X)$ ensures that this distribution does not depend on θ). As described in Section 1, the co-sufficient sampling (CSS) method operates by drawing copies from this conditional distribution. That is, the joint distribution of the data and the copies, under the CSS method, is given by

$$\begin{cases} X \sim P_{\theta_0}, \\ S = S(X), \\ \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \mid X, S \overset{\mathrm{iid}}{\sim} P(\cdot|S), \end{cases}$$

where θ_0 is the unknown true parameter. Clearly, the real and fake data $X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ are i.i.d. conditional on S, and are therefore exchangeable, meaning that the $\widetilde{X}^{(m)}$'s provide a valid "control group" for the real data X regardless of the unknown θ_0.

As discussed above, this framework is limited to only certain specific models, since many common models are "degenerate" (such as the model classes described in Section 1.2), where any sufficient statistic $S = S(X)$ reveals so much information about X that it leads to a completely powerless procedure against the alternative hypothesis of interest. We can instead consider statistics $S = S(X)$ that are not sufficient, but are *approximately sufficient*, meaning that the distribution $P_\theta(\cdot|S)$ is *approximately* unaffected by the value of θ—more concretely, if we can estimate θ with a consistent estimator $\widehat{\theta}$, then we only need to ensure that $P_{\widehat{\theta}}(\cdot|S) \approx P_\theta(\cdot|S)$. In fact, for many settings, maximum likelihood estimation is known to provide an asymptotically sufficient statistic [28, 29, 44]. Thus, we can take $S = S(X)$ to simply be $\widehat{\theta}_{\mathrm{MLE}}(X)$, or more generally, any other estimator of θ_0 that is asymptotically sufficient.

In this setting, we write $P_{\theta_0}(\cdot|\widehat{\theta})$ to denote the conditional distribution of $X|\widehat{\theta}$ when the data is distributed as $X \sim P_{\theta_0}$ and we calculate $\widehat{\theta} = \widehat{\theta}_{\mathrm{MLE}}(X)$. Of course, we cannot draw

the copies from this distribution since θ_0 is unknown, but if $\widehat{\theta} = \widehat{\theta}_{\mathrm{MLE}}(X)$ is approximately sufficient, then the distribution $P_{\theta_0}(\cdot|\widehat{\theta})$ should depend only slightly on θ_0. In particular, we will use $\widehat{\theta}$ itself as a plug-in estimate for θ_0, leading to the joint model

$$\begin{cases} X \sim P_{\theta_0}, \\ \widehat{\theta} = \widehat{\theta}_{\mathrm{MLE}}(X), \\ \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \mid X, \widehat{\theta} \overset{\mathrm{iid}}{\sim} P_{\widehat{\theta}}(\cdot|\widehat{\theta}). \end{cases}$$

These copies form an approximately valid control group as long as $P_{\widehat{\theta}}(\cdot|\widehat{\theta}) \approx P_{\theta_0}(\cdot|\widehat{\theta})$.

In our aCSS algorithm, we will replace the deterministic step $\widehat{\theta} = \widehat{\theta}(X)$ with a randomized estimator (essentially, adding a small random perturbation into the likelihood maximization problem). Adding noise is beneficial for computational reasons, since the set of $x \in \mathcal{X}$ whose MLE is exactly equal to $\widehat{\theta}_{\mathrm{MLE}}(X)$ may be a challenging set to sample from. For certain examples, adding noise can also be beneficial from the statistical point of view, as for, for example, the logistic regression setting, described in Model Class 1, where conditioning on the exact MLE, $\widehat{\theta}_{\mathrm{MLE}}(X)$, may lead to a zero-power scenario. (We will discuss the role of σ further in Section 3.3 below.) In addition, we will also allow adding a twice-differentiable regularization function $\mathcal{R}(\theta)$ to the likelihood maximization problem, for instance $\mathcal{R}(\theta) \propto \|\theta\|^2$ for ridge regression, which may be beneficial in some applications.

Informally, our proposed aCSS algorithm takes the following form:

$$\begin{cases} X \sim P_{\theta_0}, \\ W \sim \mathcal{N}\left(0, \dfrac{1}{d}\mathbf{I}_d\right), \\ \widehat{\theta} = \widehat{\theta}(X, W) = \underset{\theta \in \Theta}{\arg\min}\{-\log f(X; \theta) + \mathcal{R}(\theta) + \sigma W^{\top}\theta\}, \\ \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \mid X, \widehat{\theta} \overset{\mathrm{iid}}{\sim} P_{\widehat{\theta}}(\cdot|\widehat{\theta}), \end{cases}$$

where again $P_{\theta_0}(\cdot|\widehat{\theta})$ denotes the conditional distribution of $X|\widehat{\theta}$ when the data is distributed as $X \sim P_{\theta_0}$, and $P_{\widehat{\theta}}(\cdot|\widehat{\theta})$ is a plug-in estimate.

However, in many settings the penalized negative log-likelihood may not be strongly convex, or might even be nonconvex, in which case we will need to modify this procedure—while it is the case that, in many statistical problems, many tools exist that are likely to find the (perturbed) MLE with high probability (e.g., by carefully choosing a good initialization point), we will need to account for the fact that finding the global optimum is not guaranteed. Furthermore, in order to construct the copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$, we are implicitly assuming that we are able to generate i.i.d. samples from the conditional distribution of $X|\widehat{\theta}$. In practice, sampling directly from this density may be impossible, so we may need to turn to techniques such as Markov chain Monte Carlo (MCMC), which can introduce dependence between the samples. Our next task, then, is to develop a more general and rigorous form of this simple algorithm, so that we can provide a practical method that can be deployed in a broad range of settings.

2.2. *Algorithm for approximate co-sufficient sampling.* In this section, we will formally define our aCSS algorithm. Below, we define our noisy estimator $\widehat{\theta}$ (Section 2.2.1), calculate the conditional distribution of $X|\widehat{\theta}$ (Section 2.2.2), and describe how to sample the copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ from the estimated conditional distribution (Section 2.2.3).

2.2.1. *Sampling the estimator.* Recall that the estimator $\widehat{\theta}$ is intended to be approximately equal to the MLE, even though it includes a regularization function and a random perturbation into the likelihood maximization problem. Writing

$$\mathcal{L}(\theta; x) = -\log f(x; \theta) + \mathcal{R}(\theta),$$

consider the optimization problem

$$(2.2) \qquad \underset{\theta \in \Theta}{\arg\min}\, \mathcal{L}(\theta; X, W) \quad \text{where } \mathcal{L}(\theta; x, w) = \mathcal{L}(\theta; x) + \sigma w^\top \theta,$$

where $W \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ is independent Gaussian noise, $\sigma > 0$ determines the noise level of the random perturbation, and $\mathcal{R} : \Theta \to \mathbb{R}$ is a twice-differentiable regularization function. In order to accommodate the penalized and unpenalized estimators with a single unified presentation, we can view the unpenalized version as a special case by simply taking $\mathcal{R}(\theta) \equiv 0$. (This type of randomly perturbed log-likelihood was previously studied by [43], with the different aim of enabling selective inference on a high-dimensional parameter θ. In their work, the object of interest is the distribution of $\widehat{\theta}$, to enable inference on θ_0, whereas in our setting θ_0 is essentially a nuisance parameter.)

In the general setting where the negative log-likelihood might be nonconvex, the optimization problem (2.2) may be challenging—in particular, in the presence of nonconvexity, how would we find a global minimizer, and is a global minimizer even guaranteed to exist? In many settings, any available algorithm would only be able to guarantee that we find a first-order stationary point to (2.2) (if it even converges at all). To address this, we modify our procedure to allow $\widehat{\theta}$ to only *usually* be a well-behaved *local* optimum of (2.2). This enables aCSS to draw on the vast literature on optimizing penalized maximum likelihoods. Although the random perturbation by W makes (2.2) slightly nonstandard for penalized maximum likelihood, the perturbation is linear in θ and hence has no impact on Hessians or convexity and only adds a fixed, trivially computable constant vector to the gradient. Thus, although large linear perturbations can "tip over" an otherwise well-behaved basin of attraction, our theory will ensure this never happens asymptotically and in practice one can always choose σ small enough to make this astronomically unlikely; see Appendix D.1 in the Supplementary Material for a more detailed discussion. In summary, we expect that any algorithm that empirically often (it need not be provably often) finds a local optimum for the unperturbed penalized maximum likelihood problem will suffice with almost no modification to solve (2.2) for the purposes required by the theory in this paper.

In particular, we will define $\widehat{\theta}$ to be any measurable function mapping a (data, noise) pair (x, w) to an estimate, that is,

$$\widehat{\theta} : \mathcal{X} \times \mathbb{R}^d \to \Theta,$$

and we will later assume this map is likely to return a strict second-order stationary point (SSOSP) of the minimization problem (2.2). Here we say that θ is a SSOSP of $\mathcal{L}(\theta; x, w)$ if two conditions are satisfied:

- θ is a first-order stationary point (FOSP) of $\mathcal{L}(\theta; x, w)$, meaning that $\nabla_\theta \mathcal{L}(\theta; x, w) = 0$ or equivalently $w = -\frac{\nabla_\theta \mathcal{L}(\theta; x)}{\sigma}$.
- The objective function is strictly convex at θ, that is, $\nabla_\theta^2 \mathcal{L}(\theta; x, w) \succ 0$ or equivalently $\nabla_\theta^2 \mathcal{L}(\theta; x) \succ 0$.

We should think of $\widehat{\theta}(x, w)$ as the output of some optimization algorithm, such as gradient descent, being run to convergence on the minimization problem (2.2).

From this point on, abusing notation, depending on context we may write $\widehat{\theta}$ to denote the map $\widehat{\theta} : \mathcal{X} \times \mathbb{R}^d \to \Theta$, or may also write $\widehat{\theta}$ to denote $\widehat{\theta}(X, W)$, the random variable obtained by applying this map to the data.

2.2.2. *Calculating the distribution conditioned on the estimator.* Our next step is to calculate the conditional distribution of $X | \widehat{\theta}$, where $\widehat{\theta} = \widehat{\theta}(X, W)$ for random Gaussian noise $W \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$. As it turns out, it is generally not possible to do this exactly—in the rare

degenerate case where $\widehat{\theta}(X, W)$ may fail to find a SSOSP of the optimization problem (2.2), we do not know the distribution of $\widehat{\theta}|X$ and therefore cannot calculate the distribution of $X|\widehat{\theta}$. We will avoid this degeneracy by conditioning on the event that $\widehat{\theta}(X, W)$ returns a SSOSP.

First, we assume some standard conditions on the parametric family, and a differentiability condition on the model and the regularization function (we will also assume implicitly that all the functions defined so far, namely, $\widehat{\theta}$, p, $\mathcal{L}$ and its derivatives, are measurable with respect to $\nu_\mathcal{X} \times$ Leb or $\nu_\mathcal{X}$ or Leb, as appropriate).

ASSUMPTION 1 (Regularity conditions).   The family $\{P_\theta : \theta \in \Theta\}$ and regularization function $\mathcal{R}(\theta)$ satisfy:

- $\Theta \subseteq \mathbb{R}^d$ is a convex and open subset;
- For each $\theta \in \Theta$, P_θ has density $f(x; \theta) > 0$ with respect to the base measure $\nu_\mathcal{X}$;
- For each $x \in \mathcal{X}$, the function $\theta \mapsto \mathcal{L}(\theta; x) = -\log f(x; \theta) + \mathcal{R}(\theta)$ is continuously twice differentiable.

We are now ready to calculate the conditional distribution of $X|\widehat{\theta}$.

LEMMA 1.   *Suppose Assumption 1 holds. Fix any $\theta_0 \in \Theta$, and let $(X, \widehat{\theta})$ be drawn from the joint model*

$$(2.3) \qquad \begin{cases} X \sim P_{\theta_0}, \\ W \sim \mathcal{N}\left(0, \frac{1}{d}\mathbf{I}_d\right), \\ \widehat{\theta} = \widehat{\theta}(X, W). \end{cases}$$

Suppose the event that $\widehat{\theta}$ is a SSOSP of $\mathcal{L}(\theta; X, W)$ has positive probability.
   Then, conditional on this event, the conditional distribution of $X|\widehat{\theta}$ has density

$$(2.4) \qquad p_{\theta_0}(\cdot|\widehat{\theta}) \propto f(x; \theta_0) \cdot \exp\left\{-\frac{\|\nabla_\theta \mathcal{L}(\widehat{\theta}; x)\|^2}{2\sigma^2/d}\right\} \cdot \det(\nabla_\theta^2 \mathcal{L}(\widehat{\theta}; x)) \cdot \mathbb{1}_{x \in \mathcal{X}_{\widehat{\theta}}}$$

with respect to the base measure $\nu_\mathcal{X}$, where

$$(2.5) \qquad \mathcal{X}_\theta = \{x \in \mathcal{X} : \text{for some } w \in \mathbb{R}^d, \theta = \widehat{\theta}(x, w) \text{ is a SSOSP of } \mathcal{L}(\theta; x, w)\}.$$

The proof of this lemma is given in Appendix A.2. For intuition, we can consider the terms appearing in the calculation (2.4): the first term $f(x; \theta_0)$ expresses the original distribution of X (before conditioning), the second term $\exp\{\dots\}$ comes from the density of the multivariate normal distribution of W, the third term $\det(\dots)$ arises from a change-of-variables calculation when we move from the joint distribution of (X, W) to that of $(X, \widehat{\theta})$, and the final term $\mathbb{1}_{x \in \mathcal{X}_{\widehat{\theta}}}$ handles potential technical issues such as failure to find a SSOSP. In particular, the form of the second term is due to our choice of the multivariate normal distribution for the noise W; if we instead chose a different noise distribution, the results of this lemma would still hold if we make the appropriate changes to this second term (and the method would yield the same types of theoretical results as long as the distribution of W is continuous, supported everywhere on $\mathbb{R}^d$, and has similar concentration properties for $\|W\|$). In this work, we choose a multivariate normal distribution since the outcome of the procedure will therefore be invariant to rotations of the parameter space Θ; in settings where the choice of the basis for Θ is meaningful (e.g., we expect sparsity), it may be interesting to instead consider a nonrotationally invariant noise distribution.

2.2.3. *Sampling the copies.* We next need to specify how to sample the copies $\widetilde{X}^{(1)}, \ldots,$ $\widetilde{X}^{(M)}$. Below we describe several different approaches—which one we use will depend on the computational complexity of the problem at hand.

2.2.3.1. *The i.i.d. sampling case.* In order to construct copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ that are exchangeable with the data X, we would like to sample the copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ i.i.d. from the density $p_{\theta_0}(\cdot|\widehat{\theta})$, which by Lemma 1 specifies the exact conditional distribution of $X|\widehat{\theta}$. Since θ_0 is unknown we will use $\widehat{\theta}$ as a plug-in estimator. Our procedure is the following: after observing the data X,

$$(2.6) \quad \begin{cases} \text{Draw } W \sim \mathcal{N}\left(0, \frac{1}{d}\mathbf{I}_d\right) \text{ and define } \widehat{\theta} = \widehat{\theta}(X, W). \\ \text{If } \widehat{\theta} \text{ is a SSOSP of } \mathcal{L}(\theta; X, W), \text{ then draw } \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \overset{\text{iid}}{\sim} p_{\widehat{\theta}}(\cdot|\widehat{\theta}), \\ \qquad\qquad\qquad\qquad\qquad \text{otherwise return } \widetilde{X}^{(1)} = \cdots \widetilde{X}^{(M)} = X. \end{cases}$$

Here our estimated density for the conditional distribution of $X|\widehat{\theta}$ is given by

$$(2.7) \qquad p_{\widehat{\theta}}(x|\widehat{\theta}) \propto f(x; \widehat{\theta}) \cdot \exp\left\{-\frac{\|\nabla_\theta \mathcal{L}(\widehat{\theta}; x)\|^2}{2\sigma^2/d}\right\} \cdot \det(\nabla_\theta^2 \mathcal{L}(\widehat{\theta}; x)) \cdot \mathbb{1}_{x \in \mathcal{X}_{\widehat{\theta}}}$$

with respect to the base measure $\nu_\mathcal{X}$. (Lemma 2, in Appendix B.3 in the Supplementary Material, will verify that this expression indeed defines a valid density.)

Of course, in order to implement the sampling algorithm given in (2.6), we are implicitly assuming that it is computationally feasible to generate i.i.d. samples from $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$. To avoid making this assumption, we next consider a more general framework.

2.2.3.2. *The MCMC sampling case.* In the general case where sampling directly from $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$ may not be possible, we can instead use MCMC or any other strategy that ensures exchangeability. To be concrete, we will consider two schemes from [11] for constructing the copies with MCMC sampling. Given $\widehat{\theta}$, let $\Pi(\cdot; x)$ be any collection of transition distributions, such that the density $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$ defines a stationary distribution. Assume that Π defines a reversible Markov chain. Given Π, we define two different schemes for generating the copies. (See Figure 2 for an illustration of these schemes.)

- *Hub-and-spoke sampler.* Given X and $\widehat{\theta}$, we sample the copies as follows:
  - Initialize at X, and run the Markov chain for L steps to define the "hub" $\widetilde{X}^*$.
  - Independently for $m = 1, \ldots, M$, initialize at $\widetilde{X}^*$ and run the Markov chain for L steps to define the "spoke" $\widetilde{X}^{(m)}$.
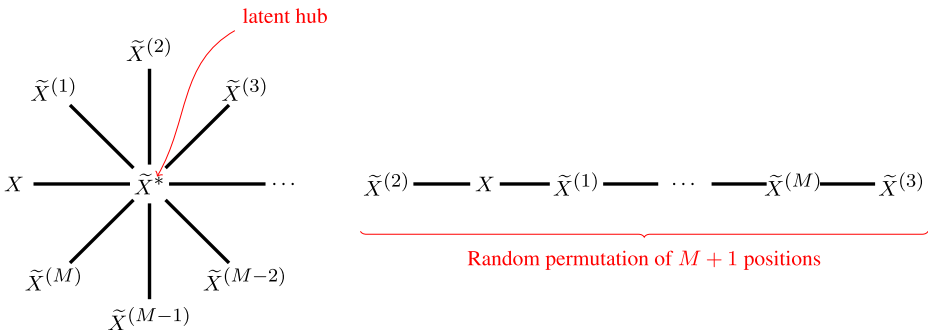

FIG. 2. *Left*: *the hub-and-spoke sampler. Right*: *the permuted serial sampler. In both diagrams, each thick black line represents running the reversible Markov chain for L steps.*

- *Permuted serial sampler*. Given X and $\widehat{\theta}$, we sample the copies as follows:
  - Draw a uniform permutation π on $\{0, \ldots, M\}$ and find $m^* \in \{0, \ldots, M\}$ such that $\pi(m^*) = 0$.
  - Initialize at X, and run the Markov chain for Lm^* steps, stopping every Lth step to define the copies $\widetilde{X}^{(\pi(m^*-1))}, \ldots, \widetilde{X}^{(\pi(0))}$.
  - Independently, initialize at X, and run the Markov chain for $L(M - m^*)$ steps, stopping every Lth step to define the copies $\widetilde{X}^{(\pi(m^*+1))}, \ldots, \widetilde{X}^{(\pi(M))}$.

Later on, we will give concrete examples of how to implement these sampling schemes for specific models.

2.2.3.3. *A unified definition.* To generalize our various options (i.i.d. sampling, hub-and-spoke MCMC sampling, and permuted serial MCMC sampling), we will write $\widetilde{P}_M(\cdot; X, \widehat{\theta})$ to denote the distribution of the collection of copies $(\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)})$ conditional on X and $\widehat{\theta}$. For all three cases, our aCSS procedure for sampling the copies is the following:

$$
(2.8) \quad
\begin{cases}
\text{Draw } W \sim \mathcal{N}\left(0, \dfrac{1}{d}\mathbf{I}_d\right) \text{ and define } \widehat{\theta} = \widehat{\theta}(X, W). \\
\text{If } \widehat{\theta} \text{ is a SSOSP of } \mathcal{L}(\theta; X, W), \text{ then draw } (\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \sim \widetilde{P}_M(\cdot; X, \widehat{\theta}), \\
\qquad\qquad\qquad\text{otherwise return } \widetilde{X}^{(1)} = \cdots \widetilde{X}^{(M)} = X.
\end{cases}
$$

In the i.i.d. sampling case, $\widetilde{P}_M(\cdot; X, \widehat{\theta})$ is simply equal to sampling from the product density $p_{\widehat{\theta}}(\cdot | \widehat{\theta}) \times \cdots \times p_{\widehat{\theta}}(\cdot | \widehat{\theta})$, and therefore depends on $\widehat{\theta}$ but not on X, while for the two MCMC samplers, there is dependence between the data and the copies even after conditioning on $\widehat{\theta}$ (although, if the chain length L is sufficiently long, we would expect this dependence to be weak). Despite this dependence, all three of these sampling schemes satisfy the following exchangeability condition: for all $\theta \in \Theta$ with $\nu_{\mathcal{X}}(\mathcal{X}_\theta) > 0$,

$$
(2.9) \quad
\begin{aligned}
&\text{If } X \sim p_\theta(\cdot | \theta) \text{ and } (\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \mid X \sim \widetilde{P}_M(\cdot; X, \theta), \text{ then} \\
&\quad \text{the random vector } (X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \text{ is exchangeable.}
\end{aligned}
$$

Note that $\widetilde{P}_M(\cdot; X, \theta)$ replaces *all* instances of $\widehat{\theta}$ in the definition of $\widetilde{P}_M(\cdot; X, \widehat{\theta})$ with θ's. Of course, it may be of interest to examine other sampling schemes, aside from the three described above. Our theoretical results below apply to any algorithm of the form (2.8) as long as the distribution $\widetilde{P}_M$ for drawing the copies is chosen to satisfy (2.9).

3. Theoretical results.
In this section, we present our main result, proving a bound on the excess Type I error of any aCSS testing procedure.

3.1. *Main result*: *Type I error bound.*
Before presenting the theorem, we will need a few more assumptions on the model and on the noisy estimator $\widehat{\theta}$. First, we need to assume that $\widehat{\theta}$ is (typically) an accurate estimator of the unknown true θ_0, and that $\widehat{\theta}$ will (typically) return a SSOSP for the optimization problem (2.2).

ASSUMPTION 2. For any $\theta_0 \in \Theta$, the estimator $\widehat{\theta} : \mathcal{X} \times \mathbb{R}^d \to \Theta$ satisfies

$$
(3.1) \quad \mathbb{P}\big(\|\widehat{\theta}(X, W) - \theta_0\| \le r(\theta_0), \text{ and } \widehat{\theta}(X, W) \text{ is a SSOSP of } \mathcal{L}(\theta; X, W)\big) \ge 1 - \delta(\theta_0),
$$

where the probability is taken with respect to the distribution $(X, W) \sim P_{\theta_0} \times \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$.

For many parametric families, the maximum likelihood estimator (or a penalized MLE) is typically shown to satisfy this type of condition with $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$ (here $\widetilde{\mathcal{O}}(\cdot)$ denotes that the scaling holds up to powers of $\log n$). This assumption has essentially the same flavor,

except that our estimator $\widehat{\theta}$ is a random perturbation of the penalized MLE. We discuss this assumption in more detail in Appendix C in the Supplementary Material.

Next, we place some assumptions on the derivatives of the log-likelihood. Let $H(\theta; x) = -\nabla_\theta^2 \log f(x; \theta)$ and let $H(\theta) = \mathbb{E}_{\theta_0}[H(\theta; X)]$ (in particular, $H(\theta_0)$ is the Fisher information).

ASSUMPTION 3. For any $\theta_0 \in \Theta$, the expectation $H(\theta)$ exists for all $\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta$, and furthermore

$$(3.2) \qquad \mathbb{E}_{\theta_0}\left[\sup_{\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} r(\theta_0)^2 \cdot (\lambda_{\max}(H(\theta) - H(\theta; X)))_+\right] \le \varepsilon(\theta_0)$$

and

$$(3.3) \qquad \mathbb{E}_{\theta_0}\left[\exp\left\{\sup_{\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} r(\theta_0)^2 \cdot (\lambda_{\max}(H(\theta; X) - H(\theta)))_+\right\}\right] \le e^{\varepsilon(\theta_0)}.$$

Here $r(\theta_0)$ is the same constant as appears in Assumption 2 (which, as mentioned above, will scale as $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$ in many settings). To interpret our assumption, we note that assumptions of the form

$$\|H(\theta; X) - H(\theta)\| = \mathcal{O}_P(n^{1/2})$$

are standard for establishing classical results such as asymptotic normality of the MLE; even with a bound as weak as $r(\theta_0) = o(n^{-1/4})$, this type of assumption will immediately imply that the first bound (3.2) holds. However, this type of condition is not quite sufficient for the theoretical arguments we need to establish, and we instead need the condition (3.3), which implies the same rate of convergence but with stronger control of the tails.

With our assumptions in place, we state the main result, which bounds the distance to exchangeability—and therefore, the Type I error—of any aCSS procedure.

THEOREM 1. *Suppose Assumptions 1, 2, and 3 all hold. After observing the data X, suppose we run the aCSS algorithm (2.8), where the distribution $\widetilde{P}_M$ is chosen to satisfy (2.9). Then, if $X \sim P_{\theta_0}$ for some $\theta_0 \in \Theta$, the copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ are approximately exchangeable with X, satisfying*

$$\mathsf{d}_{\mathsf{exch}}(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \le 3\sigma \cdot r(\theta_0) + \delta(\theta_0) + \varepsilon(\theta_0).$$

In particular, this implies that for any predefined test statistic $T : \mathcal{X} \to \mathbb{R}$ and rejection threshold $\alpha \in [0, 1]$, the p-value defined in (1.1) satisfies

$$\mathbb{P}(\mathsf{pval}_T(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \le \alpha) \le \alpha + 3\sigma \cdot r(\theta_0) + \delta(\theta_0) + \varepsilon(\theta_0).$$

The proof of this theorem is given in Appendix A.1.

3.2. *The asymptotic view.* The theoretical guarantee given in Theorem 1 is nonasymptotic, but it typically implies asymptotic control of the Type I error. In particular, in many standard settings where the observed data arises from an independent sample of size n, the terms $r(\theta_0)$, $\delta(\theta_0)$, and $\varepsilon(\theta_0)$ are all vanishing, and in particular we will expect to see $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$. Thus, if we choose noise level $\sigma \asymp n^a$ for some $a < \frac{1}{2}$, this will lead to asymptotic Type I error control, that is, $\mathbb{P}(\mathsf{pval} \le \alpha) = \alpha + o(1)$.

Furthermore, the Type I error bound in Theorem 1 gives insight into the role of approximate (or asymptotic) sufficiency in the method— $\widehat{\theta}(X, W)$ is essentially a MLE (assuming $\sigma = o(n^{1/2})$ as before)—this is because the size of the perturbation of the negative log-likelihood, $\|\nabla_\theta \mathcal{L}(\theta; X, W) - \nabla_\theta \mathcal{L}(\theta; X)\| = \sigma\|W\| = o(n^{1/2})$, is vanishing relative to

$\|\nabla_\theta \mathcal{L}(\theta_0; X)\| \asymp n^{1/2}$. Thus under standard assumptions, $\widehat{\theta}(X, W)$ is asymptotically efficient, and inherits the asymptotic sufficiency properties of the MLE. At a high level, this means that the distributions $p_{\theta_0}(\cdot|\widehat{\theta})$ of $X|\widehat{\theta}$ and $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$ of $\widetilde{X}^{(m)}|\widehat{\theta}$ are asymptotically equal (i.e., the total variation distance between them is vanishing), leading to asymptotic exchangeability between X and its copies, and consequently an asymptotic Type I error bound at the nominal level α as shown in Theorem 1.

3.3. *Choosing* σ. It may seem odd that we have advocated for $\sigma > 0$ and yet the Type 1 error bound in our main result gets *worse* as σ increases. Indeed, increasing σ will generally degrade the Type 1 error of aCSS testing due to the fact that, as σ is increased, the method moves farther from conditioning on a sufficient statistic. And in fact, taking the limit as $\sigma \to 0$ in Theorem 1 gives the tightest possible Type 1 error bound (only Assumption 2 depends on σ, and in general we would expect it to be even more plausible for smaller σ). In addition, as discussed in Section 2.2.1 as well as in the Supplementary Material (see Appendix D.1), increasing σ can decrease the probability of finding an SSOSP for the optimization (2.2), which will not negatively impact the Type 1 error, but will decrease the power of the test by increasing the probability of returning a p-value of 1. However, despite these two downsides, there are two critical reasons why it is advantageous, and arguably necessary, to take $\sigma > 0$, and this is why we allow for it in Theorem 1.

First, note that if we took $\sigma = 0$, aCSS would need to sample from a distribution supported on a level set of the MLE function of X. This level set is a low-dimensional (and hence measure-zero) subset of $\mathcal{X}$, and thus it is generally computationally intractable to sample from exactly. There is some work on sampling a random variable conditional on the value of a function of it (e.g., [17]), but only in very limited settings. Thus in most applications of aCSS, we are not aware of a computationally tractable approach that does not take $\sigma > 0$. Once we accept that $\sigma > 0$ is computationally necessary, the choice of its value modulates a power–computation trade-off within the MCMC samplers we propose in this paper, that is, for any value of $\sigma > 0$, higher power comes at the cost of more computation, and the value of σ determines this trade-off curve. Increasing σ improves the entire trade-off curve, making it easier to achieve higher power with less computation, but comes at the potential cost of higher Type I error as reflected in the bound of Theorem 1. This trade-off is discussed more in Appendix D.2 in the Supplementary Material, but essentially as σ approaches zero, it will take increasingly many MCMC steps (and associated computation) for the sampler to move "away" from the original X towards conditional independence. The more the sampler can move "away" from X, the higher the power of aCSS testing will tend to be, since a small p-value is obtained exactly when X stands out among the sampled copies.

Second, for models in which the MLE is sufficient for θ as well as for the parameters in a higher-dimensional supermodel of $\{P_\theta : \theta \in \Theta\}$ (e.g., in the logistic regression example the MLE is equivalent to X itself and thus is sufficient for all the parameters in any model), taking $\sigma = 0$ would lead to a completely powerless test for all alternatives in that supermodel. Exactly how large σ needs to be to break this degeneracy will likely need to be worked out on a case-by-case basis, and we defer a general treatment to future work. However, we see in Section 4 that for the logistic regression setting, described in Model Class 1, we can easily achieve high power with a σ value that is still sufficiently small to have no visible impact on the Type 1 error.

4. Examples. To provide further insight into the generality and practicality of aCSS testing, we establish that the necessary assumptions hold for four specific models. Example 1 (generalized linear models with canonical parameters) is an example of a regression model containing data with associated covariates, as discussed in Model Class 1. Example 2 (the

Behrens–Fisher problem) and Example 3 (a Gaussian spatial process) are both examples of curved exponential families, discussed earlier in Model Class 2. Example 4 (a multivariate t distribution) is a heavy-tailed model, and is thus an instance of Model Class 3.

In each case, we will see that the assumptions of Theorem 1 are satisfied with $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$, and with vanishing $\delta(\theta_0)$ and $\varepsilon(\theta_0)$. In particular, choosing a noise level $\sigma \asymp n^a$ for any $a < \frac{1}{2}$ is sufficient to ensure that the Type I error is asymptotically bounded by the nominal level α. We will then show simulation results for each of the four examples in Section 4.5 below.

4.1. *Canonical generalized linear models (GLMs)*.

EXAMPLE 1. We begin with the setting of a generalized linear model (GLM) with canonical parameters. Consider a logistic regression model with covariates $Z_i \in \mathbb{R}^d$ associated with each $X_i \in \mathbb{R}$, so that

$$f(x; \theta) = \prod_{i=1}^{n} \left(\frac{e^{Z_i^\top \theta}}{1 + e^{Z_i^\top \theta}} \right)^{x_i} \cdot \left(\frac{1}{1 + e^{Z_i^\top \theta}} \right)^{1-x_i},$$

parametrized by $\theta \in \Theta = \mathbb{R}^d$. (We interpret $f(x; \theta)$ as a density with respect to the base measure $\nu_{\mathcal{X}}$ on $\mathcal{X} = \mathbb{R}^n$ that places mass 1 on each point $x \in \{0, 1\}^n$.) We can rewrite this in the notation of a generalized linear model (GLM),

$$f(x; \theta) = \exp\left\{ x^\top Z\theta - \sum_{i=1}^{n} \log(1 + e^{Z_i^\top \theta}) \right\},$$

where $Z \in \mathbb{R}^{n \times d}$ is the matrix with rows Z_i. As discussed above, for $X \sim P_\theta$, the random vector $S(X) = Z^\top X \in \mathbb{R}^d$ provides a sufficient statistic; however, if the rows Z_i are in general position, then $Z^\top X$ will determine $X \in \{0, 1\}^n$ uniquely, meaning that X is no longer random after we condition on $S(X) = Z^\top X$. In other words, co-sufficient sampling (CSS) would lead to zero power, and we therefore need to turn to aCSS testing.

More generally, we can consider any canonical GLM, of the form

$$f(x; \theta) = \exp\left\{ x^\top Z\theta - \sum_{i=1}^{n} a(Z_i^\top \theta) \right\},$$

with respect to some base measure $\nu_{\mathcal{X}} = \mu \times \cdots \times \mu$ on $\mathcal{X} = \mathbb{R}^d$, where μ is a measure on $\mathbb{R}$. The function a is known as the *partition function*, and is strictly convex on its domain, which must be an open subset of $\mathbb{R}$. As for logistic regression, $Z^\top X$ is a sufficient statistic for $X \sim P_\theta$, but in the case of a discrete distribution (e.g., Poisson), CSS will again lead to zero power and so we should instead consider aCSS.

Suppose that the sample size n tends to infinity, while the parameter θ_0 is held constant (in particular, this implies that dimension d is held constant—we leave the high-dimensional setting for future work). For this example, and all the others below, we will consider the unpenalized version of the method, that is, $\mathcal{R}(\theta) \equiv 0$. Assume the covariates are entrywise bounded, that is, $\max_{i,j} \|Z_{ij}\|_\infty$ is bounded by a constant, and $\frac{1}{n} Z^\top Z \succeq \lambda_0 \mathbf{I}_d$ for a positive constant λ_0. We treat the covariates as fixed (i.e., the theory holds conditional on the covariates). Then, as we will show in Appendix C in the Supplementary Material, for an appropriately chosen initial estimator this example satisfies Assumptions 1, 2, and 3 with $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$, $\delta(\theta_0) = \mathcal{O}(n^{-1})$, and $\varepsilon(\theta_0) = 0$.

4.2. *The Behrens–Fisher problem.*

EXAMPLE 2.   Next, we consider the classical example of the Behrens–Fisher problem. Consider data

$$X_1^{(0)}, \ldots, X_{n^{(0)}}^{(0)} \overset{\text{iid}}{\sim} \mathcal{N}(\mu^{(0)}, \gamma^{(0)}), \qquad X_1^{(1)}, \ldots, X_{n^{(1)}}^{(1)} \overset{\text{iid}}{\sim} \mathcal{N}(\mu^{(1)}, \gamma^{(1)}),$$

with the two samples drawn independently. We are interested in testing the null hypothesis $H_0 : \mu^{(0)} = \mu^{(1)}$, and therefore the family of distributions can be parameterized by $\theta = (\mu, \gamma^{(0)}, \gamma^{(1)}) \in \Theta = \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}_+ \subseteq \mathbb{R}^3$, yielding a family $\{P_\theta : \theta \in \Theta\}$ where P_θ has density

$$f(x; \theta) = f(x; (\mu, \gamma^{(0)}, \gamma^{(1)}))$$

$$= \prod_{i=1}^{n^{(0)}} \frac{1}{\sqrt{2\pi\gamma^{(0)}}} e^{-(X_i^{(0)} - \mu)^2 / 2\gamma^{(0)}} \cdot \prod_{i=1}^{n^{(1)}} \frac{1}{\sqrt{2\pi\gamma^{(1)}}} e^{-(X_i^{(1)} - \mu)^2 / 2\gamma^{(1)}}$$

with respect to the Lebesgue measure on $\mathcal{X} = \mathbb{R}^{n^{(0)} + n^{(1)}}$.

This problem is an example of a curved exponential family (Problem Domain 2), for which the larger model is parametrized by $(\mu^{(0)}, \gamma^{(0)}, \mu^{(1)}, \gamma^{(1)})$—note that the constraint $\mu^{(0)} = \mu^{(1)}$ is a nonlinear constraint once we transform to the canonical parameters, which are given by $(\gamma^{(\ell)})^{-1}\mu^{(\ell)}, (\gamma^{(\ell)})^{-1}$ for each $\ell \in \{0, 1\}$. For this problem, under the null model (i.e., parametrized by $\theta = (\mu, \gamma^{(0)}, \gamma^{(1)})$), the minimal sufficient statistic is nonetheless four-dimensional—for example, the sample means and sample standard deviations of $\{X_i^{(0)}\}$ and of $\{X_i^{(1)}\}$ form a minimal sufficient statistic. Of course, this statistic is also sufficient for the larger alternative model (where $\mu^{(0)} \neq \mu^{(1)}$); once we condition on this sufficient statistic, the remaining randomness in the data carries no information about the parameters $\mu^{(0)}$ and $\mu^{(1)}$. Therefore, CSS would lead to a completely powerless procedure, and we instead turn to aCSS. (As mentioned earlier in Section 1.3, [31] mention the possibility of, but do not pursue, an aCSS-like procedure for this specific example.)

Suppose that the sample size n tends to infinity, while the parameter θ_0 is held constant and the ratio $\frac{\max\{n^{(0)}, n^{(1)}\}}{\min\{n^{(0)}, n^{(1)}\}}$ is bounded by a constant. Then, as we will show in Appendix C in the Supplementary Material, for an appropriately chosen initial estimator this example satisfies Assumptions 1, 2, and 3 with $r(\theta_0) \asymp \tilde{\mathcal{O}}(n^{-1/2})$, $\delta(\theta_0) \asymp \mathcal{O}(n^{-1})$, and $\varepsilon(\theta_0) = \tilde{\mathcal{O}}(n^{-1})$.

4.3. *A Gaussian spatial process.*

EXAMPLE 3.   For our next example, we will work in a dependent data setting—unlike the other three examples, we do not have independent observations. Our model is a Gaussian spatial process. Suppose that $X \in \mathbb{R}^n$ is distributed according to a multivariate Gaussian,

$$X \sim \mathcal{N}(0, \Sigma_\theta),$$

where the covariance matrix Σ_θ is parametrized by a scalar $\theta \in \mathbb{R}$. Specifically, we will consider a spatial Gaussian process where

$$(\Sigma_\theta)_{ij} = \exp\{-\theta \cdot D_{ij}\},$$

where $(D_{ij}) \in \mathbb{R}^{n \times n}$ is a pairwise distance matrix among n spatial points. In other words, we can think of the observation X_i as corresponding to a location $z_i \in \mathbb{R}^k$ for some ambient dimension k, and the correlation between X_i and X_j is a decaying function of the distance between locations z_i and z_j, that is, $D_{ij} = \|z_i - z_j\|$. We assume that the distances D_{ij} are

known, and the parameter $\theta \in \Theta = (0, \infty) \subseteq \mathbb{R}$ is the only unknown. This example, like Example 2, is an instance of a curved exponential family. In this case, the larger model is given by $X \sim \mathcal{N}(0, \Omega^{-1})$, where the inverse covariance Ω is the canonical parameter. The nonlinear constraints introduced by the spatial model take the form

$$(D_{ij})^{-1} \log(\Omega^{-1})_{ij} = (D_{k\ell})^{-1} \log(\Omega^{-1})_{kl}$$

for all indices i, j, k, ℓ (since the expression on each side of this equation should be equal to the same value θ). As in Example 4, the minimal sufficient statistic for our curved exponential null model is the same as that for the larger exponential family—in this case, it is given by the (uncentered) sample covariance—and therefore CSS would result in a powerless procedure for testing against any mean-zero multivariate Gaussian alternative.

Now we turn to aCSS for this example. In this setting, the distribution P_θ has density

$$f(x; \theta) = \frac{1}{(2\pi)^{n/2} \det(\Sigma_\theta)^{1/2}} e^{-x^\top \Sigma_\theta^{-1} x/2},$$

with respect to the Lebesgue measure on $\mathbb{R}^n$. The negative log-likelihood $\theta \mapsto -\log f(x; \theta)$ is therefore nonconvex, due to the nature of the map $\theta \mapsto \Sigma_\theta$.

It is known, however, that in the special case where the locations z_i are on a regular integer lattice, standard results such as asymptotic normality of the MLE can be obtained [3], and so we will work in this setting. Consider the integer grid $\{z_1, \ldots, z_n\} = \{1, \ldots, N\}^k$, where $n = N^k$. As above, the distances D_{ij} are given by $\|z_i - z_j\|$. Suppose that the grid size N tends to infinity, while the dimension k and the parameter θ_0 are held constant. Then, as we will show in Appendix C in the Supplementary Material, for an appropriately chosen initial estimator this example satisfies Assumptions 1, 2, and 3 with $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$, $\delta(\theta_0) = \mathcal{O}(n^{-1})$, and $\varepsilon(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$.

4.4. *The multivariate t distribution with unknown covariance.*

EXAMPLE 4. Our last example will demonstrate that our methodology can be applied even in settings where the data is extremely heavy-tailed—specifically, the multivariate t distribution. We consider a setting with n i.i.d. draws from a zero-mean multivariate t distribution,

$$X_i \overset{\text{iid}}{\sim} t_\gamma(0, \theta^{-1}),$$

where $\theta^{-1} \in \mathbb{R}^{k \times k}$ is an unknown covariance matrix while $\gamma > 0$ is the known degrees-of-freedom parameter. (Breaking with standard notation, we will use a lowercase θ to denote a matrix parameter, to agree with our notation throughout this paper.) Our family of distributions is therefore given by $\{P_\theta : \theta \in \Theta\}$, where $\Theta \subseteq \mathbb{R}^{k \times k}$ is the set of positive definite $k \times k$ matrices. We can view Θ as a convex open subset of $\mathbb{R}^d$ with $d = \frac{k(k+1)}{2}$, by considering the upper triangle of a positive definite matrix θ. The density is

$$f(x; \theta) = \prod_{i=1}^{n} c_{k,\gamma} \det(\theta)^{1/2} (\gamma + x_i^\top \theta x_i)^{-\frac{\gamma+k}{2}},$$

with respect to the Lebesgue measure on $\mathcal{X} = (\mathbb{R}^k)^n$, where $c_{k,\gamma}$ depends only on the dimension k and the degrees-of-freedom parameter γ, and not on the unknown parameter θ. Unlike a GLM, we cannot write the log-density $\log f(x; \theta)$ in the form (function of x)·(function of θ). In fact, we can see that, up to permutation and/or multiplication by -1 of the data points $i = 1, \ldots, n$, the data X itself is a minimal sufficient statistic for θ, so there is no sufficient statistic that would not essentially fully specify the data. Thus for instance, CSS

testing would be powerless against any i.i.d. alternative that is invariant to reflection through the origin. However, the approximate sufficiency framework is well suited for this example.

Suppose that the sample size n tends to infinity, while the degrees-of-freedom parameter γ and the unknown matrix parameter θ_0 are held constant (in particular, this implies that the dimension k is held constant—we leave the high-dimensional setting for future work). Then, as we will show in Appendix C in the Supplementary Material, for an appropriately chosen initial estimator this example satisfies Assumptions 1, 2, and 3 with $r(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$, $\delta(\theta_0) = \mathcal{O}(n^{-1})$, and $\varepsilon(\theta_0) = \widetilde{\mathcal{O}}(n^{-1/2})$.

4.5. *Simulations.* We now demonstrate the performance of aCSS for each of the four examples described above; code to reproduce the simulations is available at https://rinafb. github.io/code/aCSS.zip. We will first show two examples in Section 4.5.1 with relatively simple parametric alternative models, for which competing methods exist; in these examples, we will see aCSS testing is as powerful as the most powerful established method, namely, the score test. Then, in Section 4.5.2, we will consider two more complex examples exhibiting alternative models which elude standard approaches, and for which we are unaware of any existing test that would be powerful; we will see that aCSS testing can be powerful in such settings through the choice of a relatively sophisticated test statistic that fully leverages the particular alternative model.

For both types of examples, we will also see that the aCSS test is empirically valid (the rejection probability is almost exactly the nominal level $\alpha = 0.05$ under the null hypothesis) and that it has only slightly less power than an oracle method—this oracle method is given extra information about the distribution that reduces the composite null to a simple null, and computes a p-value (1.1) by applying the same statistic function T as aCSS to M copies $\widetilde{X}^{(m)}$ drawn independently (unconditionally) from that simple null.

4.5.1. *Simulations with a parametric alternative.* We use Examples 2 (Behrens–Fisher) and 4 (multivariate t) to demonstrate similar power between the aCSS test and the score test under parametric alternatives. The results, plotted in Figure 3, show the aCSS tests have very similar power to both the oracle and score tests. The simulation setups for the two examples are summarized below; the choice of the proposal distributions for the MCMC samplers, and chain lengths L, are described in detail in Appendix D in the Supplementary Material.
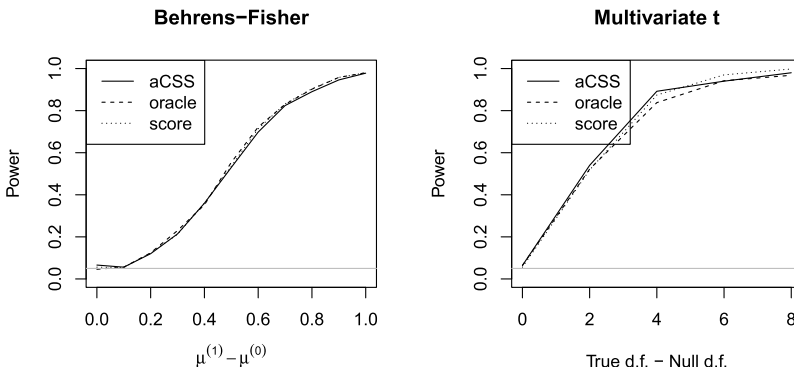

FIG. 3. *Power of the aCSS test compared to an unconditional oracle that knows the (simple) null hypothesis, and compared also to the score test, for the two examples discussed in Section 4.5.1. The aCSS test controls the Type I error at the nominal 5% level (dotted line) under the null (represented by 0 on the x-axis in each plot), and has very similar power to the oracle and score test under the alternatives. Each point represents 500 independent replications, with the maximum standard error ≈ 2% and the standard error at the left edge of each plot (under the null) below 1%.*

4.5.1.1. *Example* 2 (*Behrens–Fisher*).   For the Behrens–Fisher example, the alternative model is as described in Section 4.2 but with $(\mu^{(0)}, \mu^{(1)}, \gamma^{(0)}, \gamma^{(1)})$ unconstrained in $\mathbb{R} \times \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}_+$.

- To generate the data, we take $n^{(0)} = n^{(1)} = 50$, $\mu^{(0)} = 0$, $\gamma^{(0)} = 1$, $\gamma^{(1)} = 2$, and $\mu^{(1)} \in \{0, 0.1, 0.2, \ldots, 1\}$ (with $\mu^{(1)} = 0$ corresponding to the case where the null hypothesis holds).
- The test statistic T (used both for aCSS and for the oracle) is given by the absolute difference in sample means between the two halves of the data.
- aCSS is run with the hub-and-spoke sampler with parameters $\sigma^2 = 1$ and $M = 500$. The oracle method is given all parameter values except for $\mu^{(1)}$, so that the null $\mu^{(1)} = 0$ is simple.

4.5.1.2. *Example* 4 (*multivariate t*).   For the multivariate t example, the alternative model is as described in Section 4.4 but with the degrees-of-freedom parameter γ unknown and unconstrained (aside from being positive).

- To generate the data, we let $n = 100$, $\theta_0 = \begin{pmatrix} 1 & -0.5 \\ -0.5 & 2 \end{pmatrix}$, and $\gamma = 2$ be the assumed degrees of freedom under the null hypothesis ("d.f.$_{\text{null}}$"). The distribution of the data is given by $t_{\text{d.f.}}(0, \theta_0^{-1})$, where the degrees of freedom "d.f." is taken from $\{2, 4, 6, 8, 10\}$. Therefore d.f. $= 2$ represents the case where the null is true, and d.f. $-$ d.f.$_{\text{null}}$ measures the deviation from the null hypothesis.
- The test statistic T (used both for aCSS and for the oracle) is chosen to be the same as for the score test.
- aCSS is run with the hub-and-spoke sampler with parameters $\sigma^2 = 1$ and $M = 100$. The oracle method is given all parameter values except for γ, so that the null $\gamma = 2$ is simple.

4.5.2. *Simulations without a parametric alternative.*   We use Examples 1 and 3 to demonstrate the power of aCSS testing under more complex alternative models for which no existing methods (including the score test) are suitable. The results, plotted in Figure 4, show the aCSS tests have very similar power to the oracle. For the four examples, the settings of the simulation are as follows. In each case, the choice of the proposal distribution for the MCMC sampler, and chain length L, are described in detail in Appendix D in the Supplementary Material.
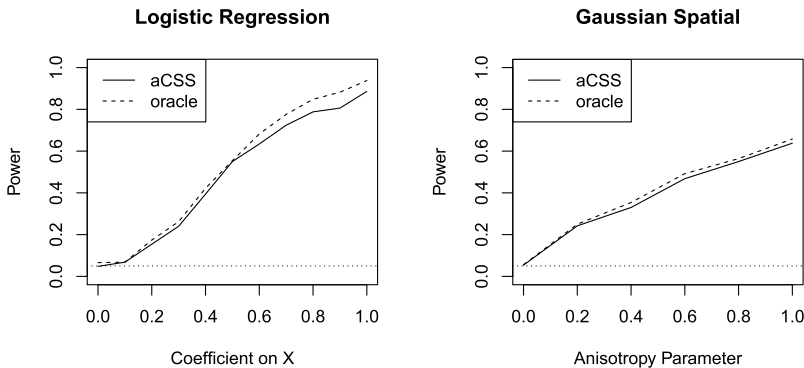

FIG. 4.   *Power of the aCSS test compared to an unconditional oracle that knows the (simple) null hypothesis, for the two examples discussed in Section 4.5.1. The aCSS test controls the Type I error at the nominal 5% level (dotted line) under the null (represented by 0 on the x-axis in each plot), and has very similar power to the oracle and score test under the alternatives. Each point represents 500 independent replications, with the maximum standard error $\approx 2\%$ and the standard error at the left edge of each plot (under the null) below 1%.*

4.5.2.1. *Example* 1 (*logistic regression*).   For the logistic regression example, we use aCSS to test a conditional independence hypothesis, so there is a response variable Y that, under the alternative, changes the conditional distribution of $X|Z$ given in Section 4.1. Y is drawn from a nonparametric model which is well approximated by a single index model, but does not exactly follow this model.

- To generate the data, we take $n = 100$, and $X|Z$ follows 5-dimensional logistic regression with coefficient vector $\theta_0 = 0.2 \cdot \mathbf{1}$. Y's conditional distribution is given by $Y|(Z, X = 0) = f_0(g_0(Z) + \beta_0^\top Z) + \mathcal{N}(0, 1)$ and $Y|(Z, X = 1) = f_1(g_1(Z) + \beta_1^\top Z) + \mathcal{N}(0, 1)$. We choose $f_0 = f_1 = t \mapsto t + 0.5t^3$, $g_0 = g_1 = z \mapsto 0.5 \sum_{j=1}^5 (z_j)_+$, $\beta_0 = c \cdot \mathbf{e}_1$, and $\beta_1 = c \cdot \mathbf{e}_5$, where $c \in \{0, 0.1, 0.2, \ldots, 1\}$ indicates the signal strength (with $c = 0$ corresponding to the null hypothesis). The nonlinearity of g_0 and g_1 means that the single index model does not exactly describe the conditional distribution of Y.
- The test statistic T (used both for aCSS and for the oracle) is computed by estimating the coefficient vector on Z in a single index model via sliced inverse regression [30] separately on the data sets $\{(Y_i, Z_i) : X_i = 0\}$ and $\{(Y_i, Z_i) : X_i = 1\}$, respectively (though recall that the single index model does not strictly hold for either data set), and then computing the angle between these estimated coefficient vectors.
- aCSS is run with the hub-and-spoke sampler with parameters $\sigma^2 = 10$ and $M = 500$. To implement the oracle method in this example, the oracle is given the distribution of $X|Z$, that is, the true coefficient vector θ_0 for the logistic regression model; under the null hypothesis, $X|Z, Y$ follows the same distribution as $X|Z$, and thus the oracle is given full knowledge of the distribution of $X|Z, Y$ under the null.

4.5.2.2. *Example* 3 (*Gaussian spatial*).   For the Gaussian spatial process example, we take a 2-dimensional 10×10 integer lattice $\{1, \ldots, 10\}^2$ for the spatial points.

- The distribution of the data is as described in Example 3 with the exception that there exists a line $\mathcal{L}$ bisecting the lattice, and for two points i and j whose positions (z_i and z_j, respectively) are on opposite sides of $\mathcal{L}$, instead of their covariance being given by $e^{-\theta_0 \|z_i - z_j\|}$, it is instead given by $(1 - c)e^{-\theta_0 \|z_i - z_j\|}$. For instance, the data points could come from soil samples, and $\mathcal{L}$ might be a possible geological ridge reducing the dependence between points on either side of it. In our experiments, $\theta_0 = 0.2$, $\mathcal{L}$ is horizontal with intercept 5.5 so that 50 of the lattice points lie below it and the other 50 lie above it, and $c \in \{0, 0.2, \ldots, 1\}$ is an anisotropy parameter, with $c = 0$ indicating an isotropic spatial process so that the null hypothesis holds.
- The test statistic T (used both for aCSS and for the oracle) is computed as follows. We first compute a threshold kernel matrix $\Delta \in \mathbb{R}^{n \times n}$ with entries $\Delta_{i,j} = e^{-|X_i - X_j|} \mathbb{1}_{\|z_i - z_j\| = 1}$ and then use Δ as the kernel matrix for spectral clustering with two clusters. Denoting the two clusters as S and S^c, the value of T is then computed as the normalized negative sum of kernel distances between the two groups:

$$-\left(\frac{1}{\sum_{i \in S^c, j \in S \cup S^c} \Delta_{i,j}} + \frac{1}{\sum_{i \in S, j \in S \cup S^c} \Delta_{i,j}} \right) \sum_{i \in S, j \in S^c} \Delta_{i,j}.$$

- aCSS is run with the hub-and-spoke sampler with parameters $\sigma^2 = 1$ and $M = 100$. The oracle method is given θ_0, $\mathcal{L}$, and the functional form for Σ in terms of c, so that the null $c = 0$ is simple.

5. Discussion.   Approximate co-sufficient sampling offers a new framework for inference on goodness-of-fit and related problems such as conditional independence testing and inference on target parameters, under mild assumptions on a composite null model. In this

section, we will first revisit the construction of aCSS to develop a deeper intuition for the ideas behind the method, and will then examine some open questions and directions that remain.

5.1. *The importance of conditioning*: *Comparison to the parametric bootstrap.* Here we return to the construction of the aCSS method, with new insights obtained from the proof of our main result, Theorem 1. In particular, why is it important to condition on $\widehat{\theta}$ when we sample the copies?

In the construction of aCSS, after conditioning on $\widehat{\theta}$, we sample copies $\widetilde{X}^{(m)}$ that are approximately exchangeable with X as long as it holds that $p_{\widehat{\theta}}(\cdot|\widehat{\theta}) \approx p_{\theta_0}(\cdot|\widehat{\theta})$. This is because, conditional on $\widehat{\theta}$, the copies are sampled from the density $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$, while the unknown true null density of $X|\widehat{\theta}$ is instead $p_{\theta_0}(\cdot|\widehat{\theta})$; we simply use $\widehat{\theta}$ as a plug-in estimator of θ_0 to define the distribution from which we sample the copies. In our proofs, we saw that aCSS leads to asymptotically valid tests as long as $d_{TV}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))$ is vanishing.

It is tempting to ask whether the same idea can be used without conditioning on $\widehat{\theta}$. That is, since the true data is distributed as $X \sim P_{\theta_0}$ under the null, can we plug in $\widehat{\theta}$ for θ_0 and sample the copies $\widetilde{X}^{(m)}$ from $P_{\widehat{\theta}}$? In fact, this nonconditional version of the procedure is simply recovering the parametric bootstrap—and, as we observed in Section 1, the parametric bootstrap may result in inflated Type I error rates in certain settings, depending on the test statistic T that we use. This is because, in general, it will not be the case that $d_{TV}(P_{\theta_0}, P_{\widehat{\theta}})$ is vanishing, even for $\widehat{\theta}$ chosen to be the MLE, and therefore, if we define the copies $\widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}$ by sampling (unconditionally) from $P_{\widehat{\theta}}$, rather than from the conditional distribution estimate $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$, it will generally be the case that, for some adversarially chosen test statistic $T(X)$, we may have Type I error that exceeds the nominal level α by a nonvanishing amount.

5.2. *Can we condition on less information?.* More generally, what if we consider conditioning on a different statistic $S = S(X)$ (or a perturbed version $S = S(X, W)$), which contains strictly less information about the data X than the (perturbed) MLE $\widehat{\theta}$? Of course, the above unconditional distribution is simply the extreme case of this idea, since it conditions on no information at all. Can we choose S so that it reveals less information about X and thus yields potentially higher power against the alternative, while still retaining approximate validity of our test? To run such a test, we would need to sample the copies from the plug-in estimated distribution $P_{\widehat{\theta}}(\cdot|S)$ rather than the true conditional null distribution $P_{\theta_0}(\cdot|S)$ of $X|S$, and in order for the copies to be approximately exchangeable with X under the null, we will need this plug-in estimate to be accurate, that is, $P_{\widehat{\theta}}(\cdot|S) \approx P_{\theta_0}(\cdot|S)$—in other words, S needs to be (approximately) sufficient. As discussed earlier in Section 3.2, the perturbed MLE $\widehat{\theta}$ is asymptotically sufficient under standard conditions; since $\widehat{\theta}$ has the same dimension d as the true parameter θ_0, it is clear that it is also (asymptotically) a *minimal* sufficient statistic. Therefore, if we choose to condition on any other statistic S, if S contains strictly less information about the data X than $\widehat{\theta}$, the approximate validity of aCSS would no longer hold.

5.3. *Open questions.* Given our new framework for inference via approximate co-sufficient sampling, many open questions remain regarding the properties of this framework, and the settings in which it can be applied.

1. *Power.* How does the choice of statistic T interact with the aCSS framework, to offer the best possible power? In particular, might it be the case that the choice of T that is most powerful under an aCSS test is not the same as the T that is most powerful for an oracle test (with a known point null hypothesis, i.e., θ_0 known)?

2. *Computation.* Are there particular algorithms that enable efficient sampling of the copies $\widetilde{X}^{(m)}$, or are there statistics T that allow us to calculate $T(\widetilde{X}^{(m)})$ without needing to fully observe $\widetilde{X}^{(m)}$—for example, through leveraging symmetries in the model and the conditional distribution?

3. *Additional models.* In addition to the examples described in this paper, can the aCSS framework be applied to similar problems such as noncanonical generalized linear models, low-rank regression, or rank-based data? Moving to more challenging settings, does the aCSS framework extend to latent variable models, errors-in-variables models, or models with missing data?

4. *Broader settings.* Can aCSS be applied in a nonparametric setting (perhaps with constraints on the statistics T allowed)? Is aCSS robust to model misspecification?

5. *Relaxing regularity conditions and extending to high dimensions.* Can aCSS be applied in settings where the null model is d-dimensional, but cannot be represented as a convex and open subset of $\mathbb{R}^d$? For instance, we may have sparsity constraints (with the parameter space given by all s-sparse vectors in $\mathbb{R}^p$) or rank constraints (with the parameter space consisting of all matrices with rank at most r in $\mathbb{R}^{a \times b}$). It would appear that any extension of aCSS testing to high dimensions would require incorporating some such low-dimensional structure, in order to ensure the existence of a nondegenerate approximately sufficient statistic, as well as a consistent estimator $\widehat{\theta}$.

APPENDIX A: PROOFS OF MAIN RESULTS

Before presenting the proofs of our theoretical results, we first establish some notation that we will use throughout these proofs. Let

$$\Omega_{\text{SSOSP}} = \{(x, w) \in \mathcal{X} \times \mathbb{R}^d : \widehat{\theta}(x, w) \text{ is a SSOSP of } \mathcal{L}(\theta; x, w)\},$$

and let

$$\Psi_{\text{SSOSP}} = \{(x, \theta) \in \mathcal{X} \times \Theta : x \in \mathcal{X}_\theta\},$$

where $\mathcal{X}_\theta$ is defined as in (2.5). The following lemma (proved in Appendix B.1 in the Supplementary Material) establishes a bijection between these sets:

LEMMA 2. *Under Assumption* 1, *the map*

$$\psi : (x, w) \mapsto (x, \widehat{\theta}(x, w))$$

defines a bijection between Ω_{SSOSP} *and* Ψ_{SSOSP}, *with inverse*

$$\psi^{-1} : (x, \theta) \mapsto \left(x, -\frac{\nabla_\theta \mathcal{L}(\theta; x)}{\sigma}\right).$$

A.1. Proof of Theorem 1. Define $P_{\theta_0}^*$ to be the distribution of $(X, W) \sim P_{\theta_0} \times \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ conditional on the event $(X, W) \in \Omega_{\text{SSOSP}}$. (If this event has probability 0 then the theorem holds trivially, so we can ignore this case.) Consider the joint distribution

$$\text{Distrib. (a):} \quad \begin{cases} (X, W) \sim P_{\theta_0}^*, \\ \widehat{\theta} = \widehat{\theta}(X, W), \\ \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} | X, \widehat{\theta} \sim \widetilde{P}_M(\cdot; X, \widehat{\theta}), \end{cases}$$

which is clearly equivalent to the aCSS procedure (2.8) if we condition on the event $(X, W) \in \Omega_{\text{SSOSP}}$. On the other hand, on the event that $(X, W) \notin \Omega_{\text{SSOSP}}$, then by definition we set

$\widetilde{X}^{(1)} = \cdots = \widetilde{X}^{(M)} = X$, and so exchangeability can only be violated on the event Ω_{SSOSP}. Therefore, we have

(A.1)
$$\mathsf{d}_{\mathsf{exch}}(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)})$$
$$\leq \mathsf{d}_{\mathsf{exch}}(\text{Distribution of } X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \text{ under Distrib. (a)}).$$

(We formalize this intuition in Lemma 1 in Appendix B.2 in the Supplementary Material.)

Next, let $Q_{\theta_0}^*$ be the marginal distribution of $\widehat{\theta}(X, W)$ under $(X, W) \sim P_{\theta_0}^*$, and define

$$\text{Distrib. (b):} \quad \begin{cases} \widehat{\theta} \sim Q_{\theta_0}^*, \\ X|\widehat{\theta} \sim p_{\theta_0}(\cdot|\widehat{\theta}), \\ \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}|X, \widehat{\theta} \sim \widetilde{P}_M(\cdot; X, \widehat{\theta}), \end{cases}$$

where $p_{\theta_0}(\cdot|\widehat{\theta})$ is defined as in Lemma 1. By definition of $Q_{\theta_0}^*$, together with Lemma 1, we can see that the joint distribution of $(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)})$ under Distrib. (b), is equal to its joint distribution under Distrib. (a), and therefore

$$\mathsf{d}_{\mathsf{exch}}(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \leq \mathsf{d}_{\mathsf{exch}}(\text{Distribution of } X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \text{ under Distrib. (b)}).$$

Finally, we define another distribution,

$$\text{Distrib. (c):} \quad \begin{cases} \widehat{\theta} \sim Q_{\theta_0}^*, \\ X|\widehat{\theta} \sim p_{\widehat{\theta}}(\cdot|\widehat{\theta}), \\ \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}|X, \widehat{\theta} \sim \widetilde{P}_M(\cdot; X, \widehat{\theta}). \end{cases}$$

(As mentioned earlier, Lemma 2 in Appendix B.3 in the Supplementary Material will verify that the density $p_{\widehat{\theta}}(\cdot|\widehat{\theta})$ exists almost surely over $\widehat{\theta}$.) Since $\widetilde{P}_M(\cdot; X, \theta)$ was constructed to satisfy (2.9), it holds that under Distrib. (c), the random variables $(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)})$ are exchangeable (in fact, they are exchangeable conditional on $\widehat{\theta}$). Therefore, by definition of $\mathsf{d}_{\mathsf{exch}}$, we have

$$\mathsf{d}_{\mathsf{exch}}(\text{Distribution of } X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)} \text{ under Distrib. (b)}) \leq \mathsf{d}_{\mathsf{TV}}(\text{Distrib. (b), Distrib. (c)}),$$

and comparing the definitions of Distrib. (b) and Distrib. (c), it is easy to verify that

$$\mathsf{d}_{\mathsf{TV}}(\text{Distrib. (b), Distrib. (c)}) = \mathbb{E}_{Q_{\theta_0}^*}[\mathsf{d}_{\mathsf{TV}}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))].$$

Combining everything, we have shown that the aCSS procedure (2.8) satisfies

(A.2)
$$\mathsf{d}_{\mathsf{exch}}(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \leq \mathbb{E}_{Q_{\theta_0}^*}[\mathsf{d}_{\mathsf{TV}}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))].$$

We next need to bound this expected total variation.

We begin with the well-known expression for total variation distance between two densities g, h, which is given by $\mathsf{d}_{\mathsf{TV}}(g, h) = \mathbb{E}_g[(1 - \frac{h(X)}{g(X)})_+]$. Therefore,

(A.3)
$$\mathbb{E}_{Q_{\theta_0}^*}[\mathsf{d}_{\mathsf{TV}}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))] = \mathbb{E}_{Q_{\theta_0}^*}\left[\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}\left[\left(1 - \frac{p_{\widehat{\theta}}(X|\widehat{\theta})}{p_{\theta_0}(X|\widehat{\theta})}\right)_+\right]\right].$$

Recalling the definitions (2.4) and (2.7) (and noting in particular that these two densities have the same support by definition), after calculating normalizing constants we can verify that

(A.4)
$$\frac{p_{\widehat{\theta}}(x|\widehat{\theta})}{p_{\theta_0}(x|\widehat{\theta})} = \frac{\frac{f(x;\widehat{\theta})}{f(x;\theta_0)}}{\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}[\frac{f(X;\widehat{\theta})}{f(X;\theta_0)}]}.$$

Next, we take a Taylor series for the function $\theta \mapsto \log f(X; \theta)$. For any x, θ we can calculate

$$\log\left(\frac{f(x; \theta_0)}{f(x; \theta)}\right) = (\theta_0 - \theta)^\top \nabla_\theta \log f(x; \theta) + \int_{t=0}^1 (1 - t) \cdot (\theta_0 - \theta)^\top \nabla_\theta^2 \log f(x; \theta_t)(\theta_0 - \theta) \, dt,$$

where we write $\theta_t = (1 - t)\theta_0 + t\theta$. Therefore, for any x, x' we have

$$\frac{\frac{f(x'; \theta)}{f(x'; \theta_0)}}{\frac{f(x; \theta)}{f(x; \theta_0)}} = \exp\left\{\log\left(\frac{f(x; \theta_0)}{f(x; \theta)}\right) - \log\left(\frac{f(x'; \theta_0)}{f(x'; \theta)}\right)\right\}$$

$$= \exp\left\{-(\theta_0 - \theta)^\top (\nabla_\theta \log f(x'; \theta) - \nabla_\theta \log f(x; \theta))\right.$$

$$\left. - \int_{t=0}^1 (1 - t) \cdot (\theta_0 - \theta)^\top (\nabla_\theta^2 \log f(x'; \theta_t) - \nabla_\theta^2 \log f(x; \theta_t))(\theta_0 - \theta) \, dt\right\}$$

$$= \exp\left\{(\theta_0 - \theta)^\top (\nabla_\theta \mathcal{L}(\theta; x') - \nabla_\theta \mathcal{L}(\theta; x))\right.$$

$$\left. + \int_{t=0}^1 (1 - t) \cdot (\theta_0 - \theta)^\top (H(\theta_t; x') - H(\theta_t; x))(\theta_0 - \theta) \, dt\right\}$$

$$\leq \exp\left\{(\theta_0 - \theta)^\top (\nabla_\theta \mathcal{L}(\theta; x') - \nabla_\theta \mathcal{L}(\theta; x))\right.$$

$$\left. + \frac{1}{2} \sup_{t \in [0,1]} (\theta_0 - \theta)^\top (H(\theta_t; x') - H(\theta_t; x))(\theta_0 - \theta)\right\},$$

where the inequality holds since $\int_{t=0}^1 (1 - t) \cdot h(t) \, dt \leq \int_{t=0}^1 (1 - t) \, dt \cdot \sup_{t \in [0,1]} h(t) = \frac{1}{2} \sup_{t \in [0,1]} h(t)$ for any function $h : \mathbb{R} \to \mathbb{R}$. For any $\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta$, it therefore holds that, for all x, x',

$$\frac{\frac{f(x'; \theta)}{f(x'; \theta_0)}}{\frac{f(x; \theta)}{f(x; \theta_0)}} \leq \exp\left\{r(\theta_0)(\|\nabla_\theta \mathcal{L}(\theta; x')\| + \|\nabla_\theta \mathcal{L}(\theta; x)\|)\right.$$

$$\left. + \frac{r(\theta_0)^2}{2} \sup_{\theta' \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} \lambda_{\max}(H(\theta'; x') - H(\theta'; x))\right\}$$

$$\leq \exp\{\Delta_1(x, \theta) + \Delta_1'(x', \theta)\},$$

where we define

$$\Delta_1(x, \theta) = r(\theta_0) \|\nabla_\theta \mathcal{L}(\theta; x)\| + \frac{r(\theta_0)^2}{2} \sup_{\theta' \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} (\lambda_{\max}(H(\theta') - H(\theta'; x)))_+,$$

and

$$\Delta_1'(x, \theta) = r(\theta_0) \|\nabla_\theta \mathcal{L}(\theta; x)\| + \frac{r(\theta_0)^2}{2} \sup_{\theta' \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} (\lambda_{\max}(H(\theta'; x) - H(\theta')))_+.$$

Applying this calculation with $x' = X$, we obtain

$$\frac{\frac{f(x; \theta)}{f(x; \theta_0)}}{\mathbb{E}_{p_{\theta_0}(\cdot|\theta)}[\frac{f(X; \theta)}{f(X; \theta_0)}]} = \left(\mathbb{E}_{p_{\theta_0}(\cdot|\theta)}\left[\frac{\frac{f(X; \theta)}{f(X; \theta_0)}}{\frac{f(x; \theta)}{f(x; \theta_0)}}\right]\right)^{-1} \geq \frac{1}{\mathbb{E}_{p_{\theta_0}(\cdot|\theta)}[e^{\Delta_1'(X, \theta)}] \cdot e^{\Delta_1(x, \theta)}}$$

for all x and for all $\theta \in \Theta$ such that $\|\theta - \theta_0\| \le r(\theta_0)$. Returning to (A.3) and (A.4) above, and defining $\mathcal{E}_{\text{ball}}$ to be the event that $\|\widehat{\theta} - \theta_0\| \le r(\theta_0)$, we therefore have

$$\mathbb{E}_{Q_{\theta_0}^*}\big[\mathsf{d}_{\mathsf{TV}}\big(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta})\big)\big]$$

$$= \mathbb{E}_{Q_{\theta_0}^*}\bigg[\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}\bigg[\bigg(1 - \frac{p_{\widehat{\theta}}(X|\widehat{\theta})}{p_{\theta_0}(X|\widehat{\theta})}\bigg)_+\bigg]\bigg]$$

$$\le \mathbb{P}_{Q_{\theta_0}^*}\big(\mathcal{E}_{\text{ball}}^c\big) + \mathbb{E}_{Q_{\theta_0}^*}\bigg[\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}\bigg[\mathbb{1}_{\mathcal{E}_{\text{ball}}} \cdot \bigg(1 - \frac{p_{\widehat{\theta}}(X|\widehat{\theta})}{p_{\theta_0}(X|\widehat{\theta})}\bigg)_+\bigg]\bigg]$$

$$\le \mathbb{P}_{Q_{\theta_0}^*}\big(\mathcal{E}_{\text{ball}}^c\big) + \mathbb{E}_{Q_{\theta_0}^*}\bigg[\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}\bigg[1 - \frac{1}{\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}[e^{\Delta_1'(X,\widehat{\theta})}] \cdot e^{\Delta_1(X,\widehat{\theta})}}\bigg]\bigg]$$

$$\le \mathbb{P}_{Q_{\theta_0}^*}\big(\mathcal{E}_{\text{ball}}^c\big) + \mathbb{E}_{Q_{\theta_0}^*}\bigg[\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}[\Delta_1(X,\widehat{\theta})] + 1 - \frac{1}{\mathbb{E}_{p_{\theta_0}(\cdot|\widehat{\theta})}[e^{\Delta_1'(X,\widehat{\theta})}]}\bigg],$$

where the last step holds since $1 - ab \le (1 - a) + (1 - b) \le \log(1/a) + (1 - b)$ for any $a, b \in (0, 1]$. (Note that, in the next-to-last line, the two random variables X appearing in the denominator are different—they are sampled independently conditional on $\widehat{\theta}$ from the distribution $p_{\theta_0}(\cdot|\widehat{\theta})$.)

Next, recall that by Lemma 1 together with the definition of $Q_{\theta_0}^*$, the joint distribution of $(X, \widehat{\theta})$ in this calculation above (i.e., $\widehat{\theta} \sim Q_{\theta_0}^*$ and $X|\widehat{\theta} \sim p_{\theta_0}(\cdot|\widehat{\theta})$), is equivalent to the joint distribution of $(X, \widehat{\theta}(X, W))$ when $(X, W) \sim P_{\theta_0}^*$. Therefore, our calculation above can be rewritten as follows (where we also apply Jensen's inequality to the last term):

$$\mathbb{E}_{Q_{\theta_0}^*}\big[\mathsf{d}_{\mathsf{TV}}\big(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta})\big)\big] \le \mathbb{P}_{P_{\theta_0}^*}\big(\mathcal{E}_{\text{ball}}^c\big) + \mathbb{E}_{P_{\theta_0}^*}\big[\Delta_1(X, \widehat{\theta}(X, W))\big]$$

$$+ \bigg(1 - \frac{1}{\mathbb{E}_{P_{\theta_0}^*}[e^{\Delta_1'(X, \widehat{\theta}(X, W))}]}\bigg).$$

Next, let

$$\Delta_2(x, w) = r(\theta_0)\sigma\|w\| + \frac{r(\theta_0)^2}{2} \sup_{\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} \big(\lambda_{\max}\big(H(\theta) - H(\theta; x)\big)\big)_+,$$

and

$$\Delta_2'(x, w) = r(\theta_0)\sigma\|w\| + \frac{r(\theta_0)^2}{2} \sup_{\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} \big(\lambda_{\max}\big(H(\theta; x) - H(\theta)\big)\big)_+,$$

and observe that $\Delta_1(x, \widehat{\theta}(x, w)) = \Delta_2(x, w)$ and $\Delta_1'(x, \widehat{\theta}(x, w)) = \Delta_2'(x, w)$ for all $(x, w) \in \Omega_{\text{SSOSP}}$, since $0 = \nabla_\theta \mathcal{L}(\widehat{\theta}(x, w); x, w) = \nabla_\theta \mathcal{L}(\widehat{\theta}(x, w); x) + \sigma w$ for all (x, w) in this set by definition. Therefore, since $(X, W) \in \Omega_{\text{SSOSP}}$ almost surely under $P_{\theta_0}^*$ by definition, we have

$$\mathbb{E}_{Q_{\theta_0}^*}\big[\mathsf{d}_{\mathsf{TV}}\big(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta})\big)\big] \le \mathbb{P}_{P_{\theta_0}^*}\big(\mathcal{E}_{\text{ball}}^c\big) + \mathbb{E}_{P_{\theta_0}^*}\big[\Delta_2(X, W)\big]$$

$$+ \bigg(1 - \frac{1}{\mathbb{E}_{P_{\theta_0}^*}[e^{\Delta_2'(X, W)}]}\bigg).$$

Now let $\mathcal{E}_{\text{SSOSP}}$ be the event that $(X, W) \in \Omega_{\text{SSOSP}}$. Recall that $P_{\theta_0}^*$ is the joint distribution of $(X, W) \sim P_{\theta_0} \times \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ conditional on $\mathcal{E}_{\text{SSOSP}}$. Therefore, we can write this as

follows where we now take all probabilities and expectations with respect to $(X, W) \sim P_{\theta_0} \times \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$:

$$\mathbb{E}_{Q_{\theta_0}^*}\big[\mathsf{d}_{\mathsf{TV}}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))\big]$$

$$\leq \mathbb{P}(\mathcal{E}_{\text{ball}}^c|\mathcal{E}_{\text{SSOSP}}) + \mathbb{E}[\Delta_2(X, W)|\mathcal{E}_{\text{SSOSP}}] + \left(1 - \frac{1}{\mathbb{E}[e^{\Delta_2'(X, W)}|\mathcal{E}_{\text{SSOSP}}]}\right)$$

$$\leq \frac{\mathbb{P}(\mathcal{E}_{\text{ball}}^c \cap \mathcal{E}_{\text{SSOSP}}) + \mathbb{E}[\Delta_2(X, W)]}{\mathbb{P}(\mathcal{E}_{\text{SSOSP}})} + \left(1 - \frac{\mathbb{P}(\mathcal{E}_{\text{SSOSP}})}{\mathbb{E}[e^{\Delta_2'(X, W)} \cdot \mathbb{1}_{\mathcal{E}_{\text{SSOSP}}}]}\right)$$

$$\leq \frac{\mathbb{P}(\mathcal{E}_{\text{ball}}^c \cap \mathcal{E}_{\text{SSOSP}}) + \mathbb{E}[\Delta_2(X, W)]}{\mathbb{P}(\mathcal{E}_{\text{SSOSP}})} + \left(1 - \frac{1 - \mathbb{P}(\mathcal{E}_{\text{SSOSP}}^c)}{\mathbb{E}[e^{\Delta_2'(X, W)}] - \mathbb{P}(\mathcal{E}_{\text{SSOSP}}^c)}\right)$$

$$\leq \frac{\mathbb{P}(\mathcal{E}_{\text{ball}}^c \cap \mathcal{E}_{\text{SSOSP}}) + \mathbb{E}[\Delta_2(X, W)] + \log \mathbb{E}[e^{\Delta_2'(X, W)}]}{\mathbb{P}(\mathcal{E}_{\text{SSOSP}})},$$

where the next-to-last step holds since $\Delta_2'(X, W) \geq 0$ by definition, and the last step holds since $1 - \frac{1-a}{b-a} \leq \frac{1-1/b}{1-a} \leq \frac{\log(b)}{1-a}$ for all $a \in [0, 1)$ and $b \geq 1$. Finally, we apply our assumptions. By Assumption 2, we have $\mathbb{P}(\mathcal{E}_{\text{ball}} \cap \mathcal{E}_{\text{SSOSP}}) \geq 1 - \delta(\theta_0)$, and so

$$\mathbb{P}(\mathcal{E}_{\text{ball}}^c \cap \mathcal{E}_{\text{SSOSP}}) \leq \delta(\theta_0) - \mathbb{P}(\mathcal{E}_{\text{SSOSP}}^c).$$

Next,

$$\mathbb{E}[\Delta_2(X, W)] = \mathbb{E}[r(\theta_0)\sigma\|W\|] + \mathbb{E}\left[\frac{r(\theta_0)^2}{2} \sup_{\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta} (\lambda_{\max}(H(\theta) - H(\theta; X)))_+\right]$$

$$\leq \frac{1}{2}\log \mathbb{E}[e^{2r(\theta_0)\sigma\|W\|}] + \frac{\varepsilon(\theta_0)}{2},$$

where the last step holds by Jensen's inequality for the first term and by the bound (3.2) in Assumption 3 for the second term. And, by Cauchy–Schwarz,

$$\log \mathbb{E}[e^{\Delta_2'(X, W)}] \leq \frac{1}{2}\log \mathbb{E}[e^{2r(\theta_0)\sigma\|W\|}] + \frac{1}{2}\log \mathbb{E}[e^{r(\theta_0)^2 \sup_{\theta \in \mathbb{B}(\theta_0, r(\theta_0)) \cap \Theta}(\lambda_{\max}(H(\theta; X) - H(\theta)))_+}]$$

$$\leq \frac{1}{2}\log \mathbb{E}[e^{2r(\theta_0)\sigma\|W\|}] + \frac{\varepsilon(\theta_0)}{2},$$

where the last step holds by the bound (3.3) in Assumption 3. Finally, since $W \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ we know that $\mathbb{E}[e^{t\|W\|}] \leq e^{t+t^2/2d}$ for any $t > 0$ (see, e.g., [12], Theorem 5.5). Therefore,

$$\log \mathbb{E}[e^{2r(\theta_0)\sigma\|W\|}] \leq 2\sigma \cdot r(\theta_0) + \frac{2\sigma^2 \cdot r(\theta_0)^2}{d} \leq 3\sigma \cdot r(\theta_0),$$

where the last step holds since $d \geq 1$ and we can assume $2\sigma \cdot r(\theta_0) \leq 1$ (as otherwise, the result of the theorem holds trivially). Combining everything, we have

$$\mathbb{E}_{Q_{\theta_0}^*}\big[\mathsf{d}_{\mathsf{TV}}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))\big] \leq \frac{3\sigma \cdot r(\theta_0) + \delta(\theta_0) + \varepsilon(\theta_0) - \mathbb{P}(\mathcal{E}_{\text{SSOSP}}^c)}{1 - \mathbb{P}(\mathcal{E}_{\text{SSOSP}}^c)}.$$

Since total variation distance is bounded by 1, trivially we can relax this to

$$\mathbb{E}_{Q_{\theta_0}^*}\big[\mathsf{d}_{\mathsf{TV}}(p_{\theta_0}(\cdot|\widehat{\theta}), p_{\widehat{\theta}}(\cdot|\widehat{\theta}))\big] \leq 3\sigma \cdot r(\theta_0) + \delta(\theta_0) + \varepsilon(\theta_0).$$

Returning to (A.2), we see that the aCSS procedure (2.8) satisfies

$$\mathsf{d}_{\text{exch}}(X, \widetilde{X}^{(1)}, \ldots, \widetilde{X}^{(M)}) \leq 3\sigma \cdot r(\theta_0) + \delta(\theta_0) + \varepsilon(\theta_0),$$

as desired.

A.2. Proof of Lemma 1. Consider the joint distribution $(X, W) \sim P_{\theta_0} \times \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ conditioned on the event that $(X, W) \in \Omega_{\text{SSOSP}}$, which is assumed to occur with positive probability. The joint density of (X, W), after conditioning on this event, is therefore proportional to the function

$$(A.5) \qquad g_{\theta_0}(x, w) = f(x; \theta_0) \cdot \exp\left\{-\frac{d}{2}\|w\|^2\right\} \cdot \mathbb{1}_{(x,w)\in\Omega_{\text{SSOSP}}},$$

with respect to the measure $\nu_{\mathcal{X}} \times$ Leb. We will consider the induced joint distribution of $(X, \widehat{\theta}(X, W))$, and will calculate its joint density.

Define ψ and ψ^{-1} as in Lemma 2. Fix any measurable subset $A \subseteq \Psi_{\text{SSOSP}}$. Then, writing $\psi^{-1}(A) = \{(x, w) \in \Omega_{\text{SSOSP}} : \psi(x, w) \in A\} \subseteq \Omega_{\text{SSOSP}}$,

$$\mathbb{P}((X, \widehat{\theta}(X, W)) \in A) = \mathbb{P}((X, W) \in \psi^{-1}(A)) = \frac{\int_{\psi^{-1}(A)} g_{\theta_0}(x, w)\, d\nu_{\mathcal{X}}(x)dw}{\int_{\mathcal{X}\times\mathbb{R}^d} g_{\theta_0}(x', w')\, d\nu_{\mathcal{X}}(x')dw'},$$

where the probability is taken with respect to $(X, W) \sim P_{\theta_0} \times \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ conditioned on the event that $(X, W) \in \Omega_{\text{SSOSP}}$. (Note that, since g_{θ_0} is proportional to a density on (X, W) with respect to $\nu_{\mathcal{X}} \times$ Leb, this implies that the denominator $\int_{\mathcal{X}\times\mathbb{R}^d} g_{\theta_0}(x', w')\, d\nu_{\mathcal{X}}(x')dw'$ in the last expression above must be finite and positive.)

From this point on, the result essentially follows from a change-of-variables calculation, under the transformation $\theta = \widehat{\theta}(x, w)$. However, with our weak assumptions, we cannot assume standard conditions (such as, e.g., the support of $\widehat{\theta}|X$ being an open subset of $\mathbb{R}^d$— it may even be the case that this set does not contain any open subset), so we will need to be careful. Fixing any $x \in \mathcal{X}$, a change-of-variables calculation, proved formally in Appendix B.4 in the Supplementary Material, establishes that

$$\int_\Theta \exp\left\{-\frac{d}{2\sigma^2}\|\nabla_\theta \mathcal{L}(\theta; x)\|^2\right\} \cdot \det(\nabla_\theta^2 \mathcal{L}(\theta; x)) \cdot \mathbb{1}_{(x,\theta)\in A\cap\Psi_{\text{SSOSP}}}\, d\theta$$

$$(A.6) \qquad = \sigma^d \int_{\mathbb{R}^d} \exp\left\{-\frac{d}{2\sigma^2}\|\nabla_\theta \mathcal{L}(\widehat{\theta}(x, w); x)\|^2\right\} \cdot \mathbb{1}_{(x,\widehat{\theta}(x,w))\in A} \cdot \mathbb{1}_{(x,w)\in\Omega_{\text{SSOSP}}}\, dw$$

$$= \sigma^d \int_{\mathbb{R}^d} \exp\left\{-\frac{d}{2}\|w\|^2\right\} \cdot \mathbb{1}_{(x,\widehat{\theta}(x,w))\in A} \cdot \mathbb{1}_{(x,w)\in\Omega_{\text{SSOSP}}}\, dw$$

$$(A.7) \qquad = \sigma^d \int_{\mathbb{R}^d} \exp\left\{-\frac{d}{2}\|w\|^2\right\} \cdot \mathbb{1}_{(x,w)\in\psi^{-1}(A)\cap\Omega_{\text{SSOSP}}}\, dw,$$

where the second step uses the fact that $w = -\frac{\nabla_\theta \mathcal{L}(\widehat{\theta}(x,w);x)}{\sigma}$ for any $(x, w) \in \Omega_{\text{SSOSP}}$ by the SSOSP conditions, and the last step applies the definition of ψ as in Lemma 2. Now define the function

$$h_{\theta_0}(x, \theta) := \frac{f(x; \theta_0) \exp\{-\frac{d}{2\sigma^2}\|\nabla_\theta \mathcal{L}(\theta; x)\|^2\} \cdot \det(\nabla_\theta^2 \mathcal{L}(\theta; x)) \cdot \mathbb{1}_{x\in\mathcal{X}_\theta}}{\sigma^d \int_{\mathcal{X}\times\mathbb{R}^d} g_{\theta_0}(x', w')\, d\nu_{\mathcal{X}}(x')dw'}$$

on $(x, \theta) \in \mathcal{X} \times \Theta$. We then have

$$\mathbb{P}((X, \widehat{\theta}(X, W)) \in A)$$

$$= \frac{\int_{\psi^{-1}(A)} g_{\theta_0}(x, w)\, d\nu_{\mathcal{X}}(x)dw}{\int_{\mathcal{X}\times\mathbb{R}^d} g_{\theta_0}(x', w')\, d\nu_{\mathcal{X}}(x')dw'}$$

$$= \frac{\int_{\mathcal{X}} \int_{\mathbb{R}^d} f(x; \theta_0) \cdot \exp\{-\frac{d}{2}\|w\|^2\} \cdot \mathbb{1}_{(x,w)\in\psi^{-1}(A)\cap\Omega_{\text{SSOSP}}}\, dw\, d\nu_{\mathcal{X}}(x)}{\int_{\mathcal{X}\times\mathbb{R}^d} g_{\theta_0}(x', w')\, d\nu_{\mathcal{X}}(x')dw'} \qquad \text{by (A.5)}$$

$$= \frac{\int_{\mathcal{X}} f(x; \theta_0) \int_{\Theta} \exp\{-\frac{d}{2\sigma^2} \|\nabla_\theta \mathcal{L}(\theta; x)\|^2\} \cdot \det(\nabla_\theta^2 \mathcal{L}(\theta; x)) \cdot \mathbb{1}_{(x,\theta) \in A \cap \Psi_{\mathrm{SSOSP}}} \, d\theta \, d\nu_{\mathcal{X}}(x)}{\sigma^d \int_{\mathcal{X} \times \mathbb{R}^d} g_{\theta_0}(x', w') \, d\nu_{\mathcal{X}}(x') dw'}$$

by (A.7)

$$= \int_A h_{\theta_0}(x, \theta) \, d\nu_{\mathcal{X}}(x) \, d\theta,$$

where the last step holds since $\mathbb{1}_{x \in \mathcal{X}_\theta} = \mathbb{1}_{(x,\theta) \in \Psi_{\mathrm{SSOSP}}}$ for all (x, θ), by definition of Ψ_{SSOSP}. Therefore, this calculation establishes that, conditional on the event that $\widehat{\theta}(X, W)$ is a SSOSP of $\mathcal{L}(\theta; X, W)$, the joint distribution of $(X, \widehat{\theta}(X, W))$ has density $h_{\theta_0}(x, \theta)$ with respect to the base measure $\nu_{\mathcal{X}} \times \mathrm{Leb}$.

Finally, since $h_{\theta_0}(x, \theta)$ is the joint density of $(X, \widehat{\theta}) = (X, \widehat{\theta}(X, W))$, we therefore see that $X|\widehat{\theta}$ has conditional density equal to

$$\frac{h_{\theta_0}(x, \widehat{\theta})}{\int_{x'} h_{\theta_0}(x', \widehat{\theta}) \, d\nu_{\mathcal{X}}(x')} \propto f(x; \theta_0) \exp\left\{-\frac{d}{2\sigma^2} \|\nabla_\theta \mathcal{L}(\widehat{\theta}; x)\|^2\right\} \cdot \det(\nabla_\theta^2 \mathcal{L}(\widehat{\theta}; x)) \cdot \mathbb{1}_{x \in \mathcal{X}_{\widehat{\theta}}},$$

which verifies the desired expression (2.4).

SUPPLEMENTARY MATERIAL

Additional proofs and computational details (DOI: 10.1214/22-AOS2187SUPP; .pdf). The Supplementary Material contains Appendices B–E. Appendix B contains additional proofs supporting our main results, Appendix C contains proofs for the examples, Appendix D provides computational and implementation details for the aCSS algorithm and for our examples in particular, and Appendix E provides details for Figure 1 comparing CSS with the parametric bootstrap.

REFERENCES

[1] AGRESTI, A. (1992). A survey of exact inference for contingency tables. *Statist. Sci.* **7** 131–177. MR1173420

[2] AGRESTI, A. (2001). Exact inference for categorical data: Recent advances and continuing controversies. *Stat. Med.* **20** 2709–2722.

[3] BACHOC, F. (2014). Asymptotic analysis of the role of spatial sampling for covariance parameter estimation of Gaussian processes. *J. Multivariate Anal.* **125** 1–35. MR3163828 https://doi.org/10.1016/j.jmva.2013.11.015

[4] BARBER, R. F. and CANDÈS, E. J. (2015). Controlling the false discovery rate via knockoffs. *Ann. Statist.* **43** 2055–2085. MR3375876 https://doi.org/10.1214/15-AOS1337

[5] BARBER, R. F. and JANSON, L. (2022). Supplement to "Testing goodness-of-fit and conditional independence with approximate co-sufficient sampling." https://doi.org/10.1214/22-AOS2187SUPP

[6] BARTLETT, M. S. (1937). Properties of sufficiency and statistical tests. *Proc. R. Soc. Lond. Ser. A, Math. Phys. Sci.* **160** 268–282.

[7] BAYARRI, M. J. and BERGER, J. O. (1999). Quantifying surprise in the data and model verification. In *Bayesian Statistics, 6 (Alcoceber, 1998)* 53–82. Oxford Univ. Press, New York. MR1723493

[8] BELL, C. B. (1984). Inference for goodness-of-fit problems with nuisance parameters (applications to signal detection). *J. Statist. Plann. Inference* **9** 273–284. MR0748257 https://doi.org/10.1016/0378-3758(84)90003-X

[9] BELTRÁN-BELTRÁN, J. I. and O'REILLY, F. J. (2019). On goodness of fit tests for the Poisson, negative binomial and binomial distributions. *Statist. Papers* **60** 1–18. MR3905436 https://doi.org/10.1007/s00362-016-0820-5

[10] BERRETT, T. B., WANG, Y., BARBER, R. F. and SAMWORTH, R. J. (2020). The conditional permutation test for independence while controlling for confounders. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **82** 175–197. MR4060981

[11] BESAG, J. and CLIFFORD, P. (1989). Generalized Monte Carlo significance tests. *Biometrika* **76** 633–642. MR1041408 https://doi.org/10.1093/biomet/76.4.633

[12] BOUCHERON, S., LUGOSI, G. and MASSART, P. (2013). *Concentration Inequalities: A nonasymptotic theory of independence*. Oxford Univ. Press, Oxford. MR3185193 https://doi.org/10.1093/acprof:oso/9780199535255.001.0001

[13] BRONIATOWSKI, M. and CARON, V. (2019). Conditional inference in parametric models. *J. SFdS* **160** 48–66. MR3987789

[14] CANDÈS, E., FAN, Y., JANSON, L. and LV, J. (2018). Panning for gold: 'model-X' knockoffs for high dimensional controlled variable selection. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **80** 551–577. MR3798878 https://doi.org/10.1111/rssb.12265

[15] CONTRERAS-CRISTÁN, A., LOCKHART, R. A., STEPHENS, M. A. and SUN, S. Z. (2019). On the use of priors in goodness-of-fit tests. *Canad. J. Statist.* **47** 560–579. MR4035789 https://doi.org/10.1002/cjs.11512

[16] COX, D. R. and REID, N. (1987). Parameter orthogonality and approximate conditional inference. *J. Roy. Statist. Soc. Ser. B* **49** 1–39. MR0893334

[17] DIACONIS, P., HOLMES, S. and SHAHSHAHANI, M. (2013). Sampling from a manifold. In *Advances in Modern Statistical Theory and Applications: A Festschrift in Honor of Morris L. Eaton. Inst. Math. Stat. (IMS) Collect.* **10** 102–125. IMS, Beachwood, OH. MR3586941

[18] DURBIN, J. (1961). Some methods of constructing exact tests. *Biometrika* **48** 41–55. MR0126313 https://doi.org/10.1093/biomet/48.1-2.41

[19] EFRON, B. and TIBSHIRANI, R. J. (1993). *An Introduction to the Bootstrap. Monographs on Statistics and Applied Probability* **57**. CRC Press, New York. MR1270903 https://doi.org/10.1007/978-1-4899-4541-9

[20] ENGEN, S. and LILLEGÅRD, M. (1997). Stochastic simulations conditioned on sufficient statistics. *Biometrika* **84** 235–240. MR1450203 https://doi.org/10.1093/biomet/84.1.235

[21] HAZRA, A. (2013). An exact Kolmogorov–Smirnov test for the negative Binomial distribution with unknown probability of success. *Res. Rev.: J. Stat.* **2** 1–13.

[22] HUANG, D. and JANSON, L. (2020). Relaxing the assumptions of knockoffs by conditioning. *Ann. Statist.* **48** 3021–3042. MR4152633 https://doi.org/10.1214/19-AOS1920

[23] JANKOVÁ, J., SHAH, R. D., BÜHLMANN, P. and SAMWORTH, R. J. (2020). Goodness-of-fit testing in high dimensional generalized linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **82** 773–795. MR4112784 https://doi.org/10.1111/rssb.12371

[24] KALBFLEISCH, J. D. and SPROTT, D. A. (1970). Application of likelihood methods to models involving large numbers of parameters. *J. Roy. Statist. Soc. Ser. B* **32** 175–208. MR0270474

[25] KOLASSA, J. E. (2003). Algorithms for approximate conditional inference. *Stat. Comput.* **13** 121–126. MR1963328 https://doi.org/10.1023/A:1023252308207

[26] KOUSATHANAS, A., LEUENBERGER, C., HELFER, J., QUINODOZ, M., FOLL, M. and WEGMANN, D. (2016). Likelihood-free inference in high-dimensional models. *Genetics* **203** 893–904. https://doi.org/10.1534/genetics.116.187567

[27] KUMAR, A. and PATHAK, P. K. (1977). Sufficiency and tests of goodness of fit. *Scand. J. Stat.* **4** 39–43. MR0431506

[28] LE CAM, L. (1960). Locally asymptotically normal families of distributions. Certain approximations to families of distributions and their use in the theory of estimation and testing hypotheses. *Univ. California Publ. Statist.* **3** 37–98. MR0126903

[29] LE CAM, L. (2012). *Asymptotic Methods in Statistical Decision Theory. Springer Series in Statistics*. Springer, New York.

[30] LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *J. Amer. Statist. Assoc.* **86** 316–342. MR1137117

[31] LILLEGÅRD, M. (2001). Tests based on Monte Carlo simulations conditioned on maximum likelihood estimates of nuisance parameters. *J. Stat. Comput. Simul.* **71** 1–10. MR1888292 https://doi.org/10.1080/00949650108812130

[32] LILLEGÅRD, M. and ENGEN, S. (1999). Exact confidence intervals generated by conditional parametric bootstrapping. *J. Appl. Stat.* **26** 447–459. MR1703739 https://doi.org/10.1080/02664769922331

[33] LINDQVIST, B. H. and RANNESTAD, B. (2011). Monte Carlo exact goodness-of-fit tests for nonhomogeneous Poisson processes. *Appl. Stoch. Models Bus. Ind.* **27** 329–341. MR2849567 https://doi.org/10.1002/asmb.841

[34] LINDQVIST, B. H. and TARALDSEN, G. (2013). Exact statistical inference for some parametric nonhomogeneous Poisson processes. *J. Iran. Stat. Soc.* (*JIRSS*) **12** 113–126. MR3048729 https://doi.org/10.1214/13-aos1083

[35] LOCKHART, R. A. (2012). Conditional limit laws for goodness-of-fit tests. *Bernoulli* **18** 857–882. MR2948905 https://doi.org/10.3150/11-BEJ366

[36] LOCKHART, R. A., O'REILLY, F. and STEPHENS, M. (2009). Exact conditional tests and approximate bootstrap tests for the von Mises distribution. *J. Stat. Theory Pract.* **3** 543–554. MR2750456 https://doi.org/10.1080/15598608.2009.10411945

[37] LOCKHART, R. A., O'REILLY, F. J. and STEPHENS, M. A. (2007). Use of the Gibbs sampler to obtain conditional tests, with applications. *Biometrika* **94** 992–998. MR2416805 https://doi.org/10.1093/biomet/asm065

[38] MARCUS, R., PERITZ, E. and GABRIEL, K. R. (1976). On closed testing procedures with special reference to ordered analysis of variance. *Biometrika* **63** 655–660. MR0468056 https://doi.org/10.1093/biomet/63.3.655

[39] O'REILLY, F. and GRACIA-MEDRANO, L. (2006). On the conditional distribution of goodness-of-fit tests. *Comm. Statist. Theory Methods* **35** 541–549. MR2274070 https://doi.org/10.1080/03610920500476622

[40] ROSENBAUM, P. R. (1984). Conditional permutation tests and the propensity score in observational studies. *J. Amer. Statist. Assoc.* **79** 565–574. MR0763575

[41] SANTOS, J. D. and FILHO, N. L. S. (2019). A Metropolis algorithm to obtain co-sufficient samples with applications in conditional tests. *Comm. Statist. Simulation Comput.* **48** 2655–2659. MR4001227 https://doi.org/10.1080/03610918.2018.1458140

[42] STEPHENS, M. A. (2012). Goodness-of-fit and sufficiency: Exact and approximate tests. *Methodol. Comput. Appl. Probab.* **14** 785–791. MR2966320 https://doi.org/10.1007/s11009-011-9267-2

[43] TIAN, X. and TAYLOR, J. (2018). Selective inference with a randomized response. *Ann. Statist.* **46** 679–710. MR3782381 https://doi.org/10.1214/17-AOS1564

[44] VAN DER VAART, A. W. (1998). *Asymptotic Statistics. Cambridge Series in Statistical and Probabilistic Mathematics* **3**. Cambridge Univ. Press, Cambridge. MR1652247 https://doi.org/10.1017/CBO9780511802256