

# Rates of convergence for random forests via generalized U-statistics\*

Wei Peng, Tim Coleman and Lucas Mentch

*Department of Statistics, University of Pittsburgh, Pittsburgh, PA 15260, USA*  
*e-mail: [wep15@pitt.edu](mailto:wep15@pitt.edu); [tsc35@pitt.edu](mailto:tsc35@pitt.edu); [lkm31@pitt.edu](mailto:lkm31@pitt.edu)*

**Abstract:** Random forests are among the most popular off-the-shelf supervised learning algorithms. Despite their well-documented empirical success, however, until recently, few theoretical results were available to describe their performance and behavior. In this work we push beyond recent work on consistency and asymptotic normality by establishing rates of convergence for random forests and other supervised learning ensembles. We develop the notion of generalized U-statistics and show that within this framework, random forest predictions can remain asymptotically normal for larger subsample sizes and under weaker conditions than previously established. Moreover, we provide Berry-Esseen bounds in order to quantify the rate at which this convergence occurs, making explicit the roles of the subsample size and the number of trees in determining the distribution of random forest predictions. When these generalized estimators are reduced to their classical U-statistic form, the rates we establish are faster than any available in the existing literature.

**MSC2020 subject classifications:** Primary 62E20, 62E17; secondary 60F05.

**Keywords and phrases:** Berry-Esseen, random forests, U-statistics, bagging, CART.

Received May 2021.

## Contents

1	Introduction . . . . .	233
2	Background . . . . .	235
3	Asymptotic normality . . . . .	239
3.1	Variance ratio behavior . . . . .	241
4	Berry-Esseen bounds . . . . .	246
4.1	Bounds for generalized U-statistics . . . . .	247
4.2	A tighter bound . . . . .	250
5	Discussion . . . . .	251
A	Proofs in Section 2 . . . . .	251
B	Proofs in Section 3 . . . . .	254
B.1	H-decomposition . . . . .	254
B.2	Proofs of asymptotic normality . . . . .	255
B.3	Simple base learner variance ratios . . . . .	262

---

\*This work was partially supported by NSF DMS-1712041.

B.4	Variance ratios for RP trees . . . . .	265
B.5	Variance ratio for non-adaptive $k$ -trees . . . . .	267
B.6	Simulation . . . . .	269
C	Proofs in Section 4 . . . . .	274
C.1	Introduction to Lemma 4 . . . . .	274
C.2	Berry-Esseen bounds for generalized U-statistics . . . . .	278
C.3	Discussion on a tighter bound . . . . .	288
	Acknowledgments . . . . .	288
	References . . . . .	289

## 1. Introduction

The random forest algorithm is a supervised learning tool introduced by [1] that constructs many independently randomized decision trees and aggregates their predictions by averaging in the case of regression or taking a majority vote for classification. Random forests have been shown to successfully handle high-dimensional and correlated data while exhibiting appealing properties such as fast and accurate off-the-shelf fitting without the overfitting issues that often plague related methods. They have been successfully applied in a variety of scientific fields including remote sensing [2], computational biology [3], stock price forecasting [4], and forecasting bird migration [5]. In a recent large-scale empirical study comparing 179 classifiers across the 121 datasets comprising the entire UCI machine learning repository, [6] found that random forests performed extremely well with 3 of the top 5 algorithms being some variant of the standard procedure.

Despite their wide-ranging applicability and well-documented history of empirical success, establishing formal mathematical and statistical properties for random forests has proved quite difficult, due in large part to the complex, data-dependent nature of the CART-splitting criterion [7] traditionally used to construct individual trees. [1] provided the first such result, demonstrating that the expected misclassification rate is a function of the accuracy of the individual classifiers and the correlation between them. The bound on the misclassification rate postulated in the work is loose but suggestive in the sense that interplay between these two sets the foundation for understanding the inner-workings of the procedure. [8] established a limit law for the split location in a regression tree context with independent Gaussian noise. Further analysis of the behavior of CART-style splitting was conducted by [9] who demonstrated an end-cut preference, whereby splits along non-informative variables are more likely to occur near the edges of the feature space.

A variety of other work has focused on analyzing other properties of random forest ensembles or extending the methodology to related problem types. [10] developed the idea of potential nearest neighbors and demonstrated their relationship to tree-based ensembles. More recently [11] analyzed the tradeoff between the size of the ensemble and the classification accuracy. [12], [13], and

[14] studied various idealized versions of random forests and investigated consistency while [15] proved consistency for a particular type of online forest. [16] developed the idea of random survival forests and the consistency of such models is investigated in [17] and [18]. [19] extended random forest estimates to the context of quantile regression and [20] experimented with reinforcement learning trees. For a more detailed accounting of random-forest-related research, we refer readers to an excellent recent review by [21].

In recent years, many promising developments have come by considering individual trees built with subsamples rather than the more traditional bootstrap samples. [22] extended the infinitesimal jackknife estimates of variance introduced by [23] to produce confidence intervals for subsampled random forest predictions. [24] provided the first consistency result for Breiman's original forests, establishing the  $L^2$  consistency whenever the underlying regression function is additive. [25] made the connection to infinite-order U-statistics and provided the asymptotic distributions of random forest predictions. [26] showed that for large ensembles, subsampled random forests are both asymptotically unbiased and Gaussian whenever individual trees are built according to honesty and regularity conditions.

In this paper, we continue the trend of establishing mathematical properties of random forests by building on the U-statistic connection made in [25]. As in other recent theoretical analyses on the topic (e.g. [12, 13, 14, 25, 26]), we adopt a general notion of random forests, viewing this class of estimators as those producing predictions of the form

$$\text{RF}(x) = \frac{1}{N} \sum_{i=1}^N h(x; Z_{i1}, \dots, Z_{is}; \omega)$$

where each  $Z_{i1}, \dots, Z_{is}$  denotes a subsample taken without replacement from the available training data and  $\omega$  denotes additional randomness injected into the base learner  $h$ . In particular, we do not require that base learners be trees constructed via the CART methodology as originally proposed in [1]. We establish central limit theorems for such estimators, that, to our knowledge, cover a broader range of estimators and also allow for faster subsampling rates than established in existing literature. More notably, we take a step forward in the theoretical analysis of random forests by providing Berry-Esseen Theorems governing the rate at which this convergence takes place by bounding the maximal error of approximation between the Gaussian distribution and that of the random forest predictions. In establishing these, we develop the notion of *generalized* U-statistics which allow for kernels to be incomplete, randomized, and infinite-order. Importantly, when these estimators are simplified to their classical U-statistic form, the resulting bounds we provide are faster than any in the existing literature.

The remainder of this paper is organized as follows. In Section 2, we provide additional background on the random forest algorithm and introduce the notion of generalized U-statistics. In Section 3 we provide a theorem that describes the asymptotic distribution of these statistics when the rank of the kernel is allowed

to grow with  $n$ . These distributional results rely on the behavior of a variance ratio and we conclude Section 3 by discussing its behavior for a variety of base learners. Building on these preliminary results, in Section 4, we provide Berry-Esseen bounds for both complete and incomplete generalized U-statistics.

## 2. Background

Suppose that we have data of the form  $Z_1, \dots, Z_n$  assumed to be independent and identically distributed (i.i.d.) from some distribution  $F_Z$  and let  $\theta$  be some parameter of interest. Suppose further that there exists an unbiased estimator  $h$  of  $\theta$  that is a function of  $s \leq n$  arguments and without loss of generality, assume that  $h$  is permutation symmetric in those arguments. The minimum variance unbiased estimator for  $\theta$  given by

$$U_{n,s} = \binom{n}{s}^{-1} \sum_{(n,s)} h(Z_{i_1}, \dots, Z_{i_s}) \quad (2.1)$$

is a U-statistic as introduced by [27] and [28], where the sum is taken over all  $\binom{n}{s}$  subsamples of size  $s$ ; we use the  $(n, s)$  shorthand for this quantity throughout the remainder of this paper. Standard elementary examples of U-statistics include sample mean, sample variance and covariance, and Kendall's  $\tau$ -statistic. When both the kernel  $h$  and rank  $s$  are held fixed, [28] showed that  $U_{n,s}$  tends toward a normal distribution with mean  $\theta$  and variance  $s^2 \zeta_1/n$  where, for any  $1 \leq c \leq s$ ,

$$\zeta_c = \text{cov} (h(Z_1, \dots, Z_c, Z_{c+1}, \dots, Z_s), h(Z_1, \dots, Z_c, Z'_{c+1}, \dots, Z'_s))$$

where  $Z'_{c+1}, \dots, Z'_n$  are i.i.d. from  $F_Z$  and independent of  $Z_1, \dots, Z_n$ .

Throughout the remainder of this paper, we consider a regression framework where the data consist of independent pairs of random variables consisting of covariates and a response  $Z_i = (X_i, Y_i) \in \mathcal{X} \times \mathbb{R}$  ( $i = 1, \dots, n$ ) sampled from a common distribution  $F_Z$ . Unless otherwise stated, we assume  $\mathcal{X} = \mathbb{R}^p$  for analytical convenience.

Given some  $s \leq n$ , let  $Z_{i_1}, \dots, Z_{i_s}$  denote a subsample of size  $s$  and consider a particular location  $x \in \mathbb{R}^p$ . The prediction at  $x$  can be written as  $h_x(Z_{i_1}, \dots, Z_{i_s})$  where the function  $h_x$  takes the subsampled covariates and responses as inputs, forms a regression estimate, and outputs the predicted response at  $x$ . Throughout the remainder of this paper, we drop the subscript  $x$  for notational convenience. Repeating this process on  $N$  subsamples and averaging across predictions gives

$$U_{n,s,N}(x) = \frac{1}{N} \sum_{i=1}^N h(Z_{i_1}, \dots, Z_{i_s})$$

so that our prediction now takes the form of a U-statistic with kernel  $h$ . When all subsamples are used so that  $N = \binom{n}{s}$ , the form is that of a *complete* U-statistic; whenever a smaller number of subsamples are utilized, it is *incomplete*. When the

subsample size  $s$  grows with the sample size  $n$ , these estimators are considered *infinite-order* U-statistics as introduced by [29] and utilized by [25] to establish asymptotic normality of random forests.

In a general supervised learning framework, these kernels can be thought of as base learners in an ensemble. Decision trees are among the most popular choices of base learners and are typically built according to the CART criterion. Here, splits in each cell  $A$  are chosen to maximize

$$L(j, z) = \frac{1}{|A|} \sum_{i=1}^n (Y_i - \bar{Y}_A)^2 \mathbb{1}_{X_i \in A} - \frac{1}{|A|} \sum_{i=1}^n (Y_i - \bar{Y}_{A_L} \mathbb{1}_{X_{j,i} < z} - \bar{Y}_{A_R} \mathbb{1}_{X_{j,i} \geq z})^2 \mathbb{1}_{X_i \in A}$$

across all covariates  $X_j$ ,  $1 \leq j \leq p$ , where  $z \in \mathbb{R}$ ,  $A_L = \{X \in A : X_j < z\}$ ,  $A_R = \{X \in A : X_j \geq z\}$ , and for any set  $S$ ,  $\bar{Y}_S$  denotes the average response value for observations  $X \in S$ . When trees are built with bootstrap samples, the resulting ensembles produce *bagged* estimates as discussed in [30]. The random forest extension of bagging introduced by [1] inserts additional independent randomness into each tree, typically to determine the subset of  $\mathbf{mtry} \leq p$  features eligible for splitting at each node. The subsampled version of this procedure thus produces estimates at  $x$  of the form

$$\tilde{U}_{n,s,N,\omega}(x) = \frac{1}{N} \sum_{i=1}^N h(Z_{i1}, \dots, Z_{is}; \omega). \quad (2.2)$$

Note that for each decision tree we consider an i.i.d. sample of randomness  $\omega_i$  but for notational convenience, we refer to this as simply  $\omega$  for all trees. Furthermore, in a similar fashion as above, define  $\zeta_{c,\omega}$  ( $c = 1, \dots, s-1$ ) and  $\zeta_s$  as

$$\begin{aligned} \zeta_{c,\omega} &= \text{cov} \left( h(\dots, Z_c, Z_{c+1}, \dots, Z_s; \omega), h(\dots, Z_c, Z'_{c+1}, \dots, Z'_s; \omega') \right) \\ \zeta_s &= \text{cov} \left( h(\dots, Z_c, Z_{c+1}, \dots, Z_s; \omega), h(\dots, Z_c, Z_{c+1}, \dots, Z_s; \omega) \right) \end{aligned} \quad (2.3)$$

and note that  $\zeta_s$  is simply the variance of the kernel with randomization parameter  $\omega$ .

[25] provide asymptotic distributional results for  $\tilde{U}_{n,s,N,\omega}$  with respect to their individual means that cover all possible growth rates of  $N$  with respect to  $n$ , though the form of the result provided has several practical limitations. In particular, the authors require that  $\zeta_{1,\omega}$  does not approach 0, but for most practical base learners, the correlation between estimators with only one observation in common should vanish as the subsample size grows. Indeed, Lemma 1 in Appendix A gives that  $\zeta_{1,\omega} \leq \frac{1}{s} \zeta_{s,\omega} \leq \frac{1}{s} \zeta_s$  so that when  $\zeta_s$  is bounded,  $\zeta_{1,\omega} \rightarrow 0$  as  $s \rightarrow \infty$ . In very recent work, [31] showed that the same result could be obtained under a more mild condition. In both results, however, the subsample size is limited to  $s = o(n^{1/2})$  which can be quite restrictive in practice. In Appendix

A, we demonstrate that this limitation is a result of a reliance on Hájek projections and in fact, whenever such an approach is taken, there is strong reason to believe that a subsampling rate of  $s = o(n^{1/2})$  is the largest possible. As later discussed by [26] however, when  $s$  is small, it is possible that the squared bias decays slower than the variance, thereby producing confidence intervals which, when built according to the stated Gaussian limit distribution, may not cover the true value. [26] provide an alternative central limit theorem for averages over trees built according to honesty and regularity conditions. When base learners conform to such conditions and  $N$  is very large, the authors show that the subsampling rate can be improved to  $s = o(n^\beta)$  for  $0.5 < \beta < 1$  while retaining consistent estimates.

Motivated by the form of (2.2) we now formalize the notion of *generalized U-statistics*.

**Definition 1** (generalized U-statistic). *Suppose  $Z_1, \dots, Z_n$  are i.i.d. samples from  $F_Z$  and let  $h$  denote a (possibly randomized) real-valued function utilizing  $s$  of these samples that is permutation symmetric in those  $s$  arguments. A generalized U-statistic with kernel  $h$  of order (rank)  $s$  refers to any estimator of the form*

$$U_{n,s,N,\omega} = \frac{1}{\hat{N}} \sum_{(n,s)} \rho h(Z_{i_1}, \dots, Z_{i_s}; \omega) \quad (2.4)$$

where  $\omega$  denotes i.i.d. randomness, independent of the original data. The  $\rho$  denotes i.i.d. Bernoulli random variables determining which subsamples are selected and  $\Pr(\rho = 1) = N/\binom{n}{s}$ . The actual number of subsamples selected is given by  $\hat{N} = \sum \rho$  where  $\mathbb{E}[\hat{N}] = N$ . When  $N = \binom{n}{s}$ , the estimator in Eq. (2.4) is a generalized complete U-statistic and is denoted as  $U_{n,s,\omega}$ . When  $N < \binom{n}{s}$ , these estimators are generalized incomplete U-statistics.

Though it is not practical to simulate  $\binom{n}{s}$  Bernoulli random variables, fortunately, it is equivalent to first simulate  $\hat{N} \sim \text{Binomial}(\binom{n}{s}, N/\binom{n}{s})$  and then randomly generate  $\hat{N}$  subsamples without replacement. Note also that while the number of subsamples  $\hat{N}$  in Eq. (2.4) is random, it concentrates around  $N$ .

Allowing for the possibility of a randomized kernel is of benefit here as it allows the results that follow to pertain to the kinds of randomized ensembles often considered in practice. The randomization parameter  $\omega$  might, for example, perform some kind of feature subsampling as is commonly associated with random forests – much further discussion along these lines is provided in Section 3. We stress however that the mere inclusion of such a randomization parameter is not where the true innovation in our work lies, nor should it be viewed as the “essential ingredient” in what we refer to as generalized U-statistics. Indeed, in several of the results that follow, the theoretical details needed to establish them follow a near-identical recipe regardless of whether the kernel itself takes on additional randomness.

Rather, the real benefit of considering generalized U-statistics lies in the form of the estimator itself that allows for, in essence, a random weighting to

be applied to the kernel through the use of  $\rho$ . As a bit of a preview of what is to follow, note also that in this generalized form, an incomplete U-statistic can be viewed as merely a complete U-statistic with a different kernel up to  $N/\hat{N}$ . It is these kinds of realizations that provide significant benefits for theoretical analysis by allowing us to view incomplete U-statistics as merely a modified version of its complete form, rather than as an approximation to it that inherits a remainder term that needs to be controlled. Furthermore, in the complete case, it can be shown that the variance of the U-statistic can be decomposed into a sum over  $s$  terms and that the structure of the statistic itself shrinks the higher-order terms in that sum. This careful examination of higher-order terms allows us to not only establish asymptotic normality, but to provide rates of convergence sharper than any in the existing literature, some of which are based on fundamental work dating back several decades.

In the literature on classic U-statistics, many results are derived by applying a technique called the H-decomposition, which allows the statistic to be written as a sum of uncorrelated terms. Appendix B.1 contains a detailed overview of the classic H-decomposition. The idea was first introduced by [32], but has analogues in many parts of statistics, most notably in the analysis of variance in balanced experimental designs; for a more general result, see [33]. To handle *generalized* U-statistics, we begin by extending the concept of the H-decomposition to this more general setting.

**Definition 2** (H-decomposition). *Suppose that  $Z_1, \dots, Z_s$  are i.i.d. samples from  $F_Z$  and  $h(z_1, \dots, z_s; \omega)$  is a (possibly randomized) real valued function that is permutation-symmetric in  $(z_1, \dots, z_s)$ . Let*

$$h_i(z_1, \dots, z_i) = \mathbb{E}[h(z_1, \dots, z_i, Z_{i+1}, \dots, Z_s; \omega)] - \mathbb{E}[h]$$

for  $i = 1, \dots, s$  and let

$$h^{(i)} = h_i(z_1, \dots, z_i) - \sum_{j=1}^i \sum_{(s,j)} h^{(j)}(z_{i1}, \dots, z_{ij}), \quad \text{for } i = 1, \dots, s-1 \text{ and}$$

$$h^{(s)} = h(z_1, \dots, z_s; \omega) - \sum_{j=1}^{s-1} \sum_{(s,j)} h^{(j)}(z_{i1}, \dots, z_{ij}).$$

The H-decomposition of a generalized complete U-statistic is expressed as

$$U_{n,s,\omega} = \sum_{j=1}^s \binom{s}{j} \binom{n}{j}^{-1} \sum_{(n,j)} h^{(j)}(Z_{i1}, \dots, Z_{is}). \quad (2.5)$$

When no extra randomness is injected into  $h$ , the above definition reduces to the classic H-decomposition. Note that the randomness  $\omega$  is only involved in  $h^{(s)}$ ; for  $h^{(1)}, \dots, h^{(s-1)}$ , it is marginalized out. Note that because each subsample is associated with an i.i.d. draw of the randomness  $\omega$ , each of the  $h^{(s)}$  terms in Eq. (2.5) involve this randomness though this notation is suppressed in Eq. (2.5) for readability.

### 3. Asymptotic normality

Before providing the asymptotic distributional results for generalized U-statistics of the form

$$U_{n,s,N,\omega} = \frac{1}{\widehat{N}} \sum_{(n,s)} \rho h(Z_{i_1}, \dots, Z_{i_s}; \omega) \quad (3.1)$$

we pause to emphasize the value in considering this form of estimator and to distinguish this generalization from the more classical counterparts considered in recent studies. [25] produce a central limit theorem for infinite-order U-statistics, but consider randomized kernels only insofar as establishing that when such randomness is well-behaved, the limiting distributions are equivalent. More recently, [26] analyzed random forests constructed with all possible subsamples where the kernel can thus be written in a form where the additional randomness is marginalized out. Such estimators take the form

$$\binom{n}{s}^{-1} \sum_{(n,s)} \mathbb{E}_\omega h(Z_{i_1}, \dots, Z_{i_s}; \omega) \quad (3.2)$$

so that the kernels themselves are non-random and thus the estimator is simply a complete, infinite-order U-statistic with kernel  $g = \mathbb{E}_\omega h(Z_{i_1}, \dots, Z_{i_s}; \omega)$ .

While more convenient for theoretical analysis, random forests of the form conceived in (3.2) are not generally utilized in practice, even in small-data settings since, by construction, such a statistic involves building every possible randomized tree on every possible subsample of the data. In practice, random forests might be loosely seen as a *double* or *nested* Monte Carlo approximation to the estimators in (3.2), where one source of approximation results from using  $N < \binom{n}{s}$  subsamples and the other results from estimating the kernel itself  $\mathbb{E}_\omega h(Z_{i_1}, \dots, Z_{i_s}; \omega)$  on each subsample. Recent work by [34] provides an analysis of these kinds of nested approximations.

In practice, however, random forests are nearly always constructed by selecting subsamples at random and pairing each with an independently selected randomization instance  $\omega$ , which is itself generally assumed to be selected uniformly at random. Generalized U-statistics therefore provide a direct and accurate representation of such estimators. We begin with a theorem establishing asymptotic normality for complete generalized U-statistics.

**Theorem 1.** *Let  $Z_1, \dots, Z_n$  be i.i.d. from  $F_Z$  and  $U_{n,s,\omega}$  be a generalized complete U-statistic with kernel  $h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$ ,  $\zeta_{1,\omega} = \text{var}(\mathbb{E}[h|Z_1])$  and  $\zeta_s = \text{var}(h)$ . If  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$ , then*

$$\frac{U_{n,s,\omega} - \theta}{\sqrt{s^2 \zeta_{1,\omega} / n}} \rightsquigarrow N(0, 1). \quad (3.3)$$

The proof of Theorem 1 is provided in the Appendix B.2. The general strategy is to find a linear statistic to approximate  $U_{n,s,\omega}$ , and show that the difference is negligible by applying the H-decomposition.

**Remark 1.** *The condition in Theorem 1 that  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$  can be replaced by the weaker condition that  $\frac{s}{n} \left( \frac{\zeta_s}{s\zeta_{1,\omega}} - 1 \right) \rightarrow 0$ . In practice, this condition can be satisfied by choosing  $s$  to grow slow relative to the variance ratio  $\zeta_s/s\zeta_{1,\omega}$ . In particular, whenever the ratio is bounded, choosing  $s = o(n)$  is sufficient. Thus, in establishing asymptotic normality, this weaker condition may be of minimal consequence. However, in quantifying the finite sample deviations from normality via the Berry-Esseen Theorems in Section 4, this alternative condition plays an important role in establishing the bounds provided.*

Similar results for non-generalized U-statistics have appeared in the recent works discussed earlier. Theorem 1 in [25] can be modified slightly to provide an analogous result whenever  $\frac{s^2}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$ . A recent result in [31] proceeds along these lines. Both results, however, could be improved by applying the H-decomposition rather than the Hájek projection. Similarly, Theorem 3.1 in [26] establishes asymptotic normality for non-generalized, complete U-statistics whenever the subsample size  $s$  grows like  $n^\beta$  for some  $\beta < 1$ . Here though the authors are concerned only with base learners that take the form of averages over honest and regular trees and in particular, with controlling the asymptotic bias of the resulting estimator. Thus, with minor modifications, Theorem 3.1 in [26] could be seen as something of a corollary to our Theorem 1 above, corresponding to the special case where the within-kernel randomness is held fixed or marginalized out.

The complete forms of these estimators are almost never utilized in practice due to the computational burden involved with calculating  $\binom{n}{s}$  base learners. Thus, armed with the results for the complete case, we now establish an analogous result for *incomplete* generalized U-statistics.

**Theorem 2.** *Let  $Z_1, \dots, Z_n$  be i.i.d. from  $F_Z$  and  $U_{n,s,N,\omega}$  be a generalized incomplete U-statistic with kernel  $h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$ ,  $\zeta_{1,\omega} = \text{var}(\mathbb{E}[h|Z_1])$  and  $\zeta_s = \text{var}(h)$ . Suppose that  $\mathbb{E}[|h - \theta|^{2k}]/\mathbb{E}^2[|h - \theta|^k]$  is uniformly bounded for  $k = 2, 3$  and for all  $s$ . If  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$  and  $N \rightarrow \infty$ , then*

$$\frac{U_{n,s,N,\omega} - \theta}{\sqrt{s^2\zeta_{1,\omega}/n + \zeta_s/N}} \rightsquigarrow N(0, 1). \quad (3.4)$$

**Remark 2.** *Note that the variance in the theorem above takes a different form than in Theorem 1 but the requirement that  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$  remains the same. Indeed, whenever this condition is satisfied, the complete U-statistic analogue is also asymptotically normal and normality of the incomplete version can be established as a by-product. However, this condition and more generally, asymptotic normality of the complete version, is not necessary. In such cases, choosing a very small ensemble size (e.g.  $N = o(n/s)$ ) is sufficient. More details and related results are provided in Appendix B.2 along with the proof of Theorem 2.*

Taken together, Theorems 1 and 2 provide the asymptotic distribution of generalized U-statistics for all possible growth rates on the number of subsamples  $N$  relative to  $n$ . Besides the regularity conditions on the kernel, these results require

only that  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$ . This condition, similar to the notion of  $\nu$ -incrementality discussed in [26], is not overly strong but may appear somewhat arbitrary. In the following subsection we investigate the behavior of this ratio for a variety of base learners.

### 3.1. Variance ratio behavior

For a given kernel  $h$ , let  $\hat{h}$  be the projection of  $h$  onto the linear space. We have that  $\hat{h} = \sum_{i=1}^s h_1(Z_i)$  and thus

$$\frac{\text{var}(h)}{\text{var}(\hat{h})} = \frac{\text{var}(h)}{s\text{var}(\mathbb{E}[h|Z_1])} = \frac{\zeta_s}{s\zeta_{1,\omega}}. \quad (3.5)$$

Since  $\zeta_s$  is the overall variance and  $\zeta_{1,\omega}$  can be written as the variance of the expectation of the kernel conditioning on one argument, we can view the ratio in Eq. (3.5) as a measure of the potential influence of one single observation on the output of the kernel. When  $\zeta_s/s\zeta_{1,\omega} \rightarrow 1$ ,  $h$  itself is asymptotically linear. More generally though, Theorems 1 and 2 require only that  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$  in order for the generalized U-statistic to be asymptotically normal. Thus, if the limiting behavior of the variance ratio  $\zeta_s/s\zeta_{1,\omega}$  is understood, the subsampling rate can be chosen to ensure the entire term approaches 0.

For simple kernels such as the sample mean and sample variance, it is straightforward to show that the limit of this variance ratio is 1, though this can also be shown to hold for more standard regression estimates such as ordinary least squares; see Appendix B.3 for explicit calculations. Here we first focus our attention on nearest-neighbor estimators and linear smoothers as these are more directly relatable to the tree-style base learners often used in practice, and then examine *double sample* trees, a modified CART-style decision tree.

**Proposition 1.** *Let  $Z_1, \dots, Z_s$  denote i.i.d. pairs of random variables  $(X_i, Y_i)$  and suppose  $Y_i = f(X_i) + \epsilon_i$  where  $f$  is continuous,  $\epsilon_i$  has mean 0 and variance  $\sigma^2$ , and  $X_i$  and  $\epsilon_i$  are independent. Let  $\varphi$  denote the standard  $k$ -nearest neighbor ( $k$ NN) estimator. Then*

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s}{s\zeta_1} \leq c(k)$$

where

$$c(k) = 2k / \left[ \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \frac{(i+j)!}{i!j!} \frac{1}{2^{i+j}} \right]$$

so that  $c(k)$  is decreasing in  $k$  and  $1 < c(k) \leq 2$ .

A sketch of  $c(k)$  for  $k = 1, \dots, 50$  is shown in Fig. 1. The proof of Proposition 1 is provided in Appendix B.3. Note that  $k$ NN is a nonadaptive linear smoother, the variance ratio of which is bounded above by a constant. The following result gives an upper bound for the more general class of all linear smoothers.

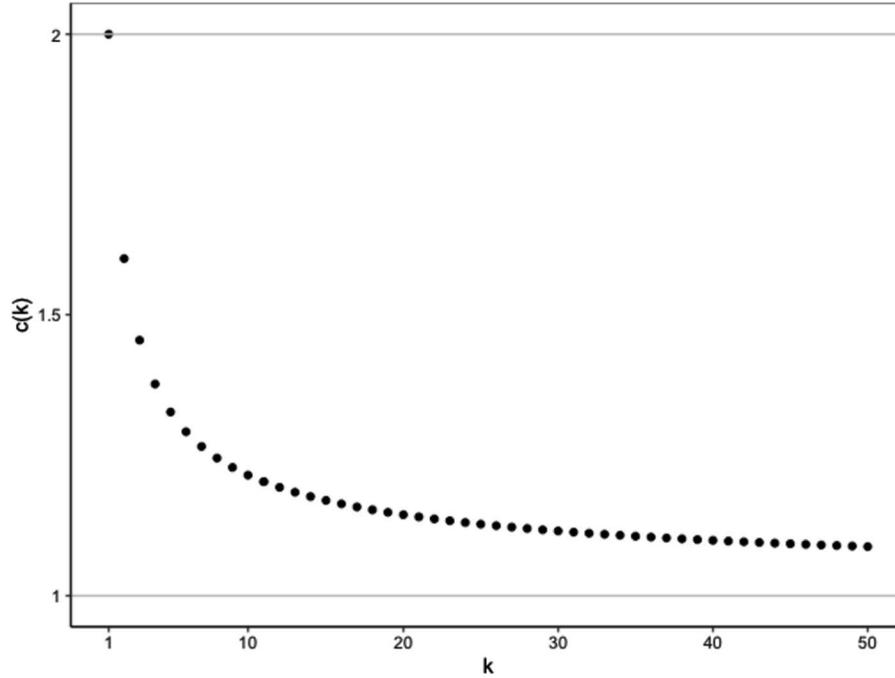


FIG 1. A plot of  $c(k)$  for  $k = 1, \dots, 50$ .

**Proposition 2.** Let  $Z_1, \dots, Z_s$  denote i.i.d. pairs of random variables  $(X_i, Y_i)$  and suppose  $Y_i = f(X_i) + \epsilon_i$  where  $f$  is bounded,  $\epsilon$  has mean 0 and variance  $\sigma^2$ , and  $X_i$  and  $\epsilon_i$  are independent. Let  $\varphi = \sum_{i=1}^s w(i, x, \mathbf{X})Y_i$  such that  $\sum_{i=1}^s w(i, x, \mathbf{X}) = 1$ , where  $\mathbf{X}$  denotes  $\{X_i\}_{i=1}^s$ . Then

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s}{s^2 \zeta_1} < \infty.$$

The results above demonstrate that the behavior of the variance ratio is manageable for  $k$ -nearest neighbor base learners and more generally, linear smoothers. Recent work [35, 36] has sought to draw a connection between these estimators and the CART-style trees utilized in Breiman's original random forests. The purely random forest [12] that determines splits completely at random, for example, is exactly a linear smoother and thus by the above result, has a variance ratio that behaves like  $O(s)$ . In work dating back even further, [10] introduced the concept of *potential* nearest neighbors (PNNs) and showed that random forests can be viewed as an adaptively weighted  $k$ -PNN method.

**Definition 3** ( $k$ -potential nearest neighbors [10]). A sample point  $Z_i = (X_i, Y_i)$  is called a  $k$ -potential nearest neighbor ( $k$ -PNN) of a target point  $x$  if and only if there are fewer than  $k$  sample points other than  $X_i$  in the hyperrectangle defined by  $x$  and  $X_i$ .

Typically, the number of potential nearest neighbors is much larger than the number of nearest neighbors. Existing nearest-neighbor methods, both adaptive and nonadaptive, predict by selecting and averaging over  $k$  points from the set of all  $k$ -PNNs. The classical kNN procedure non-adaptively chooses the  $k$  points as those closest to  $x$  under some metric whereas commonly used tree-based methods may have a terminal size bounded by  $k$  and adaptively select points from the  $k$ -PNNs based on empirical relationships in the data.

Moving closer to this, consider the base learner that forms a prediction at  $x$  by simply choosing  $k$  of the  $s$  observations in the subsample uniformly at random and averaging the corresponding response values. In Appendix B.4, we show that the resulting variance ratio for this naive estimator is given by

$$\frac{\zeta_s}{s\zeta_{1,\omega}} = \frac{s}{k} = \frac{1/k \cdot s^2}{s \cdot 1}.$$

Now reconsider the kNN base learner. We can view such an estimator as “randomly” selecting  $k$  points from the  $k$ -NNs and again predicting by taking the average. In this case, the variance ratio can be written as

$$\frac{\zeta_s}{s\zeta_{1,\omega}} = O(1) = O\left(\frac{1/k \cdot k^2}{s \cdot \mathbb{E}[\text{Pr}^2(X_1 \in \text{kNN} \mid X_1)]}\right).$$

The form of this result may naturally lead one to conjecture that for any base learner that predicts by randomly selecting and averaging over points in some set  $A$ , the resulting ratio may have the form

$$\frac{\zeta_s}{s\zeta_{1,\omega}} = O\left(\frac{1/k \cdot |A|^2}{s \cdot \mathbb{E}[\text{Pr}^2(X_1 \in A \mid X_1)]}\right). \quad (3.6)$$

Consider then a simple tree-style estimator that predicts at  $x$  by sampling  $k$  points uniformly at random from its  $k$ -PNNs and averaging the corresponding response values; we refer to these random potential nearest neighbor estimators as RP trees. The additional difficulty introduced with RP trees is that the size of this set of potential nearest neighbors is itself random, though from [10], we know that the expected number of  $k$ -PNNs is  $O(k(\log s)^{p-1})$  and so extending our conjecture, we arrive at the following proposition. We have

**Proposition 3.** *Let  $Z_1, \dots, Z_s$  denote i.i.d. pairs of random variables  $(X_i, Y_i)$  and suppose  $Y_i = f(X_i) + \epsilon_i$  where  $f$  is bounded,  $\epsilon$  has mean 0 and variance  $\sigma^2$ , and  $X_i$  and  $\epsilon_i$  are independent. Suppose further that the density of  $X$  is bounded away from 0 and infinity in  $[0, 1]^p$ . Then for the RP tree estimator, we have*

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s / s\zeta_{1,\omega}}{(\log s)^{2p-2}} < \infty. \quad (3.7)$$

The proof of Proposition 3 is provided in Appendix B.4. Here, asymptotic normality can be ensured by insisting on the same subsample sizes put forth in [26], namely that  $s = o(n^\beta)$  for some  $0.5 < \beta < 1$ .

To bridge the gap even further between these completely non-adaptive trees that decide splits independent of any data and the completely adaptive CART-style trees often employed in practice that utilize all available data to decide both the tree structure and final predictions, we turn now to the idea of *honesty* and *double sample* trees.

**Definition 4.** We call a tree-style predictor a non-adaptive  $k$ -tree if (a) the splitting does not depend on the response values of the data, but could potentially depend on the features and some other random mechanism that is independent of the data, and (b) the terminal size is bounded between  $k$  and  $2k - 1$ .

Note that part (a) of the above definition is what was referred to as the “honesty” condition in [26]. According to this definition, RP trees are simply a special type of non-adaptive  $k$ -tree. We show these RP trees are a worst case non-adaptive  $k$ -tree in the sense that the  $\zeta_1$  parameter resulting from a non-adaptive  $k$ -tree can be lower bounded by that of an RP tree.

**Proposition 4.** Under the same conditions on the underlying distribution as in Proposition 3, for the non-adaptive  $k$ -tree, we have

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s / s \zeta_{1,\omega}}{(\log s)^{2p-2}} < \infty. \quad (3.8)$$

The proof of Proposition 4 is provided in Appendix B.5. A similar result was obtained in [26] where the denominator was given as  $(\log s)^p$ . Appendix B.5 provides an alternative proof.

As is becoming evident, we now find ourselves caught in a rut familiar to most researchers studying the theoretical properties of statistical and machine learning procedures. On one hand, the non-adaptive nature of the methods described above helps a great deal in providing mathematical tractability – in our case, by making it feasible to bound  $\zeta_{1,\omega}$ . On the other, the strong performance of CART-style decision trees in practice is generally attributed, at least in large part, to their adaptive nature. Double sample trees, introduced in [26], can be thought of as something of a compromise between these two extremes, perhaps even offering the best of both worlds in some settings.

Double sample trees are constructed by first randomly splitting the training samples into two halves -  $\mathcal{D}_1$  and  $\mathcal{D}_2$ . The first half is used to obtain the structure of the tree and the second half is used to form predictions by averaging across the response values in  $\mathcal{D}_2$  from those points that fall in the same terminal node as the target point. Here we also impose that the terminal nodes in the trees contain between  $k$  and  $2k - 1$  observations from  $\mathcal{D}_2$ .

The advantage of double sample trees is somewhat obvious: the space can still be partitioned into small regions based on the CART criterion using data originating from the same underlying process, yet, because that data is independent of the second set of data used to form predictions, we can still guarantee control of  $\zeta_{1,\omega}$ . On the other hand, because only half of the data is being used for each process, one may worry that such models would lose predictive accuracy. Recent

work on regularization in random forests, however, (see, e.g., [37, 38]) suggests that this could actually be potentially advantageous in low-signal settings.

To see how control of  $\zeta_{1,\omega}$  can be established, consider some  $Z_1 \in \mathcal{D}_2$ . If we regard  $\mathcal{D}_1$  as the source of the extra randomness, then the double sample tree can be viewed as a non-adaptive  $k$ -tree with respect to  $\mathcal{D}_2$ . In particular, we can establish good control of  $\zeta_{1,\omega}$  due to Proposition 4, where  $\zeta_{1,\omega} = O(1/s(\log s)^{2p-2})$ . Indeed, letting  $h_1(Z_1) = \mathbb{E}[h(\mathcal{D}_1, \mathcal{D}_2; \omega) \mid Z_1]$ , by the law of total variation, we have

$$\begin{aligned} \zeta_{1,\omega} &= \text{var}(h_1(Z_1)) \\ &\geq \mathbb{E}[\text{var}(h_1(Z_1) \mid Z_1 \in \mathcal{D}_1) \mathbf{1}_{Z_1 \in \mathcal{D}_1}] + \mathbb{E}[\text{var}(h_1(Z_1) \mid Z_1 \in \mathcal{D}_2) \mathbf{1}_{Z_1 \in \mathcal{D}_2}] \\ &\geq \Pr(Z_1 \in \mathcal{D}_2) \text{var}(h_1(Z_1) \mid Z_1 \in \mathcal{D}_2) \\ &= \frac{1}{2} \text{var}(h_1(Z_1) \mid Z_1 \in \mathcal{D}_2) \\ &= O(1/(s(\log s)^{2p-2})). \end{aligned} \tag{3.9}$$

Note that [26] obtained a similar result for double sample trees by imposing further regularity conditions on the splitting process. We do not require any such conditions here.

In general, calculating the variance ratio for adaptive base learners without imposing specific constraints on the base learners and/or data is quite challenging. However, we conclude our discussion here by noting that the previous calculations offer some encouragement. Given the  $k$ -PNNs of some target point  $x$  and considering estimators that predict by averaging over some subset of these, we showed that for non-adaptive estimators like  $k$ NN, the variance ratio is bounded. On the other hand, when the samples are selected uniformly at random from all  $k$ -PNNs, the variance ratio is on the order of  $(\log s)^{2p-2}$ . Further and more generally, when the samples are selected in any non-adaptive fashion from all  $k$ -PNNs – including via the use of an external, independent dataset – the variance ratio remains upper bounded by  $O((\log s)^{2p-2})$ .

Tree-based estimators, by definition, predict by averaging over subsets of potential nearest neighbors, though as discussed above, the particular fashion in which those neighbors are chosen in practice is often heavily data-dependent. If, however, we are in a common regression setting where the response is regressed on covariates that contain some signal, then trees may heavily weight only a subset of the potential nearest neighbors, particularly in directions that can account for some of the variability in the response. In such settings, the variance ratio may be approximately of the form in (3.6) for some smaller set  $S$  and therefore be smaller than that of RP trees. In Appendix B.6, we provide a simple simulation study that suggests that the variance ratios resulting from ensembles are well-behaved in practice regardless of whether the base learners are simple non-adaptive nearest neighbor methods or CART-based decision trees. We also provide several plots demonstrating the asymptotic normality of the resulting predictions.

We close this section by noting that in order for these asymptotic distributions to be used for inferential purposes, one must estimate the corresponding parameters. In practice, estimating  $\zeta_{1,\omega} = \text{var}(h_1(z))$  is often quite challenging. Since  $h_1(Z_i)$  is generally unknown and we don't know the underlying distribution  $F_Z$ , we cannot generate new data and use a Monte Carlo method. Rather, we can only sample only from the original data  $(Z_1, \dots, Z_n)$ , which inevitably introduces dependence among  $\hat{h}_1(Z_1), \dots, \hat{h}_1(Z_n)$ .

Recent work, however, showed that when the conditions for Theorem 1 hold, the proposed ‘‘pseudo’’ infinitesimal jackknife estimator (closely related to that introduced in [23]) is consistent whenever  $N = \binom{n}{s}$  [39]. The insight is that although the dependence introduces bias, that bias will be dominated by  $\zeta_{1,\omega}$  if the U-statistic  $U_{n,s,\omega}$  is dominated by its linear term. Even more importantly, this work also showed that so long as  $N$  is large enough to ensure that  $n/(Ns\zeta_{1,\omega}) \rightarrow 0$ , the finite sample version of pseudo infinitesimal jackknife estimator remains consistent. Note that  $s\zeta_{1,\omega} \leq \zeta_s$  and  $\zeta_s$  could tend towards 0, so it is necessary that  $N \gg n$ . One unfortunate practical takeaway is that the number of trees required for accurately estimating  $\zeta_{1,\omega}$  can potentially be much larger than the number required to stabilize predictions from the random forest. Interested readers are also invited to see recent work by [40] that explores this variance estimation problem through the lens of V-statistics where subsamples are drawn with replacement. Note that if the asymptotic linearity condition in Theorem 1 does not hold, to the best of our knowledge, there are no methods in the existing literature capable of estimating  $\zeta_{1,\omega}$  consistently given  $s = o(n)$ . This can be explained by the fact that since  $\zeta_{1,\omega}$  represents the covariance between two trees sharing only a single common sample, whenever  $s$  is large, we do not have sufficiently many such pairs with such little dependence.

Finally, we note that estimating these variance parameters is not always essential for carrying out formal inference in a random forest context. Recognizing the potential computational strain involved in estimating  $\zeta_{1,\omega}$ , [41] recently developed an asymptotically valid permutation test for assessing variable importance in random forests. Here, rather than explicitly estimate the sampling distribution of predictions, the authors suggest building just one additional forest in which the effects of the variables under investigation are muted. Predictions are made using the trees in each forest and the individual trees are then randomly permuted between forests and their respective accuracies are recalculated. The intuition here is that if the features in question have no effect on predictive accuracy, forests containing trees that are unable to utilize those features should be as accurate as those with no such restrictions.

#### 4. Berry-Esseen bounds

Given i.i.d. random variables  $Z_1, \dots, Z_n$  with mean  $\mu$  and variance  $\sigma^2$ , the Berry-Esseen theorem [42, 43] provides a classical result describing the rate of convergence of  $S_n = \sum_{i=1}^n (Z_i - \mu)/\sigma\sqrt{n}$  to the normal distribution. It states

that provided the third moment  $v_3 = \mathbb{E}|Z - \mu|^3$  is finite,

$$\sup_{z \in \mathbb{R}} |F_n(z) - \Phi(z)| \leq \frac{Cv_3}{\sigma^3 \sqrt{n}}$$

where  $F_n$  is the distribution function of  $S_n$ ,  $\Phi$  is the distribution function of the standard normal, and  $C$  is a constant independent of  $n$  and the  $Z_i$ . Several authors (e.g. [44, 45, 46, 47]) have since contributed various iterations of Berry-Esseen theorems for U-statistics. In the following sections, we derive bounds for *generalized* U-statistics involving  $n, s, N$ , and the moments of the base learner to lend some intuition regarding how these parameters might be chosen in practice. We utilize the H-decomposition along with novel representations of U-statistics in order to provide bounds sharper than previously established in the literature for infinite-order U-statistics as well as first-of-their-kind bounds for *generalized* U-statistics.

#### 4.1. Bounds for generalized U-statistics

We begin with the following result on generalized, complete U-statistics.

**Theorem 3.** *Suppose that  $Z_1, \dots, Z_n$  are i.i.d. from  $F_Z$  and that  $U_{n,s,\omega}$  is a generalized complete U-statistic with kernel  $h = h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$ ,  $\zeta_s = \text{var}(h)$  and  $\zeta_{1,\omega} = \mathbb{E}[g^2(Z_1)]$ , where  $g(z) = \mathbb{E}[h(z, Z_2, \dots, Z_s; \omega)] - \theta$ . Suppose further that  $\zeta_s < \infty$  and  $\zeta_{1,\omega} > 0$ , then*

$$\sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s,\omega} - \theta}{\sqrt{s^2 \zeta_{1,\omega} / n}} \leq z \right\} - \Phi(z) \right| \leq \frac{6.1 \mathbb{E}|g|^3}{n^{1/2} \zeta_{1,\omega}^{3/2}} + (1 + \sqrt{2}) \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s \zeta_{1,\omega}} - 1 \right) \right\}^{1/2}.$$

A number of important points are worth noting here. First, when  $s$  is fixed, the bound has a rate on the order of  $1/\sqrt{n}$  as should be expected since this is the standard rate associated with classic (finite-order), complete U-statistics. Additionally, when the randomness  $\omega$  is held fixed so that the estimator reduces to an infinite-order U-statistic, this bound is sharper than that provided in [47], which, to our knowledge, is the sharpest to date in the existing literature. Specifically, the bound appearing in [47] replaces the term

$$\frac{s}{n} \left( \frac{\zeta_s}{s \zeta_1} - 1 \right)$$

in the bound above, which is on the order of  $s/n$ , with

$$\frac{(s-1)^2 \zeta_s}{s(n-s+1) \zeta_1}$$

which is on the order of  $s^2/n$ . An immediate consequence of this tighter bound is that when kernels are employed such that the resulting terms  $\zeta_s/s\zeta_1$  and  $\mathbb{E}|g|^3/\zeta_1^{3/2}$  are bounded, a subsampling rate of  $s = o(n)$  is sufficient to ensure

the bound converges to 0, whereas a rate of  $s = o(\sqrt{n})$  would be required according to the bound given in [47]. This sharper rate we obtain is ultimately a result of utilizing Lemma 4 together with the H-decomposition. Full details are provided in Appendix C.2.

We note also that recent work by [48] considered a similar high-dimensional statistic and obtain the Berry-Esseen bound

$$O\left(\frac{D_n^2 \log^{q_1}(dn)}{\underline{\sigma}_g^2 n}\right)^{1/6} + O\left(\frac{s^3}{n} \left[\frac{\log^{q_2}(d) D_n^2}{s \underline{\sigma}_g^2}\right]\right)^{1/4} \quad (4.1)$$

where,  $q_1, q_2$  are constants,  $d$  is the dimension of the kernel, and  $D_n$  is a uniform bound on the Orlicz norm of the components of the kernel. Note that the first term in Eq. (4.1) comes from the Berry-Esseen bound for the Hájek projection (which is a high dimensional sample mean; see [49]) and the second term comes from controlling the remainder of that Hájek projection. The two terms in this sum correspond directly to the two terms in Theorem 3.

Comparing the two bounds, we see that the second term is raised to the power of 1/4 in Eq. (4.1) but to the power of 1/2 in our result in Theorem 3. The authors in [48] also argue that the term  $\underline{\sigma}_g^2$  is lower bounded by  $O(s^{-2})$ , ultimately implying that  $s$  is required to be at most  $n^{1/4-\epsilon}$ ,  $\forall \epsilon > 0$ , whereas we require only that  $s = O(n^{1-\epsilon})$ . It seems to be the *multi*-dimensionality alone rather than the *high*-dimensionality that is driving the suboptimal rates in Eq. (4.1), because we do not have a sharp bound on the size of the remainder term nor a sharp concentration inequality like that in Eq. (C.3) to apply in higher dimensions. The authors consider bounding  $\underline{\sigma}_g^2$ , where  $\underline{\sigma}_g^2 = \min \zeta_{1,j}$  for  $j = 1, \dots, d$ , via the Cramér-Rao inequality and end up with  $\zeta_{1,j} = O(s^{-2})$ . This equality may be conservative since the bound considers the worst-case scenario. In most of our examples,  $\zeta_1$  can be on the order of  $s^{-1-\epsilon}$ ,  $\forall \epsilon > 0$ .

Generalized incomplete U-statistics can be viewed as generalized *complete* U-statistics with an alternative kernel up to the scalar  $N/\tilde{N}$ . Recognizing this fact, we can make use of the H-decomposition and Lemma 4 given in the appendix to obtain the following bound for incomplete, generalized U-statistics.

**Theorem 4.** *Suppose that  $Z_1, \dots, Z_n$  are i.i.d. from  $F_z$  and that  $U_{n,s,N,\omega}$  is a generalized incomplete U-statistic with kernel  $h = h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$ ,  $\zeta_s = \text{var}(h)$ , and  $\zeta_{1,\omega} = \mathbb{E}[g^2(Z_1)]$ , where  $g(z) = \mathbb{E}[h(z, Z_2, \dots, Z_s; \omega)] - \theta$  and let  $p = N/\binom{n}{s}$ . Suppose further that  $\zeta_s < \infty$  and  $\zeta_{1,\omega} > 0$ . Then for any  $0 < \eta_0 < 1/2$ ,*

$$\begin{aligned} & \sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s,N,\omega} - \theta}{\sqrt{s^2 \zeta_{1,\omega}/n}} \leq z \right\} - \Phi(z) \right| \\ & \leq C \left\{ \frac{\mathbb{E}|g|^3}{n^{1/2} \zeta_{1,\omega}^{3/2}} + \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s \zeta_{1,\omega}} - 1 \right) \right\}^{1/2} + \left\{ \frac{n}{N} (1-p) \frac{\zeta_s}{s \zeta_{1,\omega}} \right\}^{1/2} + N^{-1/2+\eta_0} \right\}. \end{aligned}$$

The proof of Theorem 4 is provided in Appendix C.2. The preceding theorems indicate that for both infinite-order and generalized U-statistics, when incomplete versions of these estimators are used, these statistics remain asymptotically as efficient as the complete forms so long as  $n = o(N)$ . Comparing Theorems 3 and 4, note that these bounds differ only by the inclusion of two additional final terms in the sum, where  $N^{-1/2+\eta_0}$  is due to the variation of  $\hat{N}$  with respect to  $N$  and  $\left\{\frac{n}{N}(1-p)\frac{\zeta_s}{s\zeta_{1,\omega}}\right\}^{1/2}$  is due to incompleteness, which is close to 0 in such large- $N$  settings.

However, in small- $N$  settings, this term can become quite large, leading to a bound nearing or even exceeding 1, thereby making it of little use. Theorem 5 below provides improved Berry-Esseen bounds for such small- $N$  settings where relatively few base learners are employed. To achieve this, rather than writing the estimators as linear statistics plus a small additional manageable term, we take an alternative approach that views incomplete U-statistics as complete U-statistics plus a remainder. This strategy is similar to that used in [50] who recently derived non-asymptotic Gaussian approximation error bounds for high-dimensional, incomplete U-statistics, but for kernels with fixed (finite) rank. Proofs of the following results are provided in Appendix C.2.

**Theorem 5.** *Suppose that  $Z_1, \dots, Z_n$  are i.i.d. from  $F_Z$  and that  $U_{n,s,\omega,N}$  is a generalized incomplete U-statistic with kernel  $h = h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$ ,  $\zeta_s = \text{var}(h)$ , and  $\zeta_{1,\omega} = \mathbb{E}[g^2(Z_1)]$ , where  $g(z) = \mathbb{E}[h(z, Z_2, \dots, Z_s; \omega)] - \theta$ . Suppose further that  $\zeta_s < \infty$  and  $\zeta_{1,\omega} > 0$ . If  $\mathbb{E}[|h-\theta|^{2k}]/\mathbb{E}^2[|h-\theta|^k]$  is uniformly bounded for  $k = 2, 3$  and for all  $s$ . Then for any  $0 < \eta_0 < 1/2$ ,*

$$\begin{aligned} & \sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s,N,\omega} - \theta}{\sqrt{s^2 \zeta_{1,\omega}/n + \zeta_s/N}} \leq z \right\} - \Phi(z) \right| \\ & \leq C \left\{ \frac{\mathbb{E}|g|^3}{n^{1/2}(\mathbb{E}|g|^2)^{3/2}} + \frac{\mathbb{E}|h-\theta|^3}{N^{1/2}(\mathbb{E}|h-\theta|^2)^{3/2}} + \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s\zeta_{1,\omega}} - 1 \right) \right\}^{1/2} \right. \\ & \quad \left. + N^{-1/2+\eta_0} + \left( \frac{s}{n} \right)^{1/3} \right\} \end{aligned}$$

for some constant  $C > 0$ .

Here we see that when  $s$  is fixed, the Berry-Esseen bound is on the order of  $n^{-1/3}$ . When  $s$  grows with  $n$ , the bound converges to zero as long as  $s/n \rightarrow 0$  and  $N \rightarrow \infty$  with some mild conditions on  $h$ . The fundamental task in producing this result is to show that the convolution of the two independent sequences approaches a normal distribution. A number of approximations are required, though we give a nearly sharp bound on each in order to provide the Berry-Esseen bound shown. As noted earlier, [50] recently investigated a similar setup for higher-dimensional kernels assumed to be of fixed rank. In our case, the use of an infinite-order kernel injected with extra randomness introduces additional technical difficulties, though the restriction to one-dimensional settings allows us to incorporate more useful concentration inequalities. Thus, even for ker-

nels assumed to have a fixed, finite rank, the result above is sharper than that provided in [50] for the one-dimensional setting.

As an additional benefit, we note that the bound above consisting of a five-term sum contains insightful terms not produced in [50]. In particular, the second term corresponds to the bound that would be available for an estimator that takes an average of i.i.d. random variables, while the first term plus the third term gives the bound for the complete infinite-order U-statistic setting. This leads to the very natural intuition that when  $N$  is quite small, the bound produced is approximately what would be expected by averaging over independent base learners whereas when  $N$  is large, the bound is approximately what we would expect for a complete infinite-order U-statistic. Again,  $N^{-1/2+\eta_0}$  is due to the variation of  $\hat{N}$  with respect to  $N$ . To see where the final term  $(s/n)^{1/3}$  comes from, we now delve into the proof details in the following subsection.

#### 4.2. A tighter bound

In order to obtain the previous bounds, we first condition on  $Z_1, \dots, Z_n$  and obtain a Berry-Esseen bound for the difference between the infinite-order forms of incomplete and complete U-statistics,  $U_{n,s,N} - U_{n,s}$ . The terms involved in this bound are themselves infinite-order U-statistics with kernels that are power functions of the original kernel  $h$ . We make use of Chebyshev's inequality to replace those infinite-order U-statistics by their population mean, the application of which requires no particular assumptions on the tail behavior of the kernel. This approach, however, leads to the non-optimal term of  $(\frac{s}{n})^{1/3}$ . We thus conclude our discussion on Berry-Esseen Theorems by showing in this final subsection that placing additional assumptions on the kernel  $h$  can allow the application of sharper concentration inequalities that can therefore allow the term  $(\frac{s}{n})^{1/3}$  to be replaced by  $(\frac{s}{n})^{1/2}$ .

**Theorem 6.** *Suppose that  $Z_1, \dots, Z_n$  are i.i.d. from  $F_Z$  and that  $U_{n,s,N,\omega}$  is a generalized incomplete U-statistic with kernel  $h = h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$ ,  $\zeta_s = \text{var}(h)$  and  $\zeta_{1,\omega} = \mathbb{E}[g^2(Z_1)]$ , where  $g(z) = \mathbb{E}[h(z, Z_2, \dots, Z_s; \omega)] - \theta$ . Suppose further that  $\zeta_s < \infty$  and  $\zeta_{1,\omega} > 0$ . If  $|h - \theta|^k$  is sub-Gaussian after standardization with variance proxy that is uniformly bounded for  $k = 2, 3$  and all  $s$ , then for any  $0 < \eta_0, \eta < 1/2$ ,*

$$\begin{aligned} & \sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s,N,\omega} - \theta}{\sqrt{s^2 \zeta_{1,\omega}/n + \zeta_s/N}} \leq z \right\} - \Phi(z) \right| \\ & \leq C \left\{ \frac{\mathbb{E}|g|^3}{n^{1/2}(\mathbb{E}|g|^2)^{3/2}} + \frac{\mathbb{E}|h - \theta|^3}{N^{1/2}(\mathbb{E}|h - \theta|^2)^{3/2}} + \left[ \frac{s}{n} \left( \frac{\zeta_s}{s\zeta_{1,\omega}} - 1 \right) \right]^{1/2} \right. \\ & \quad \left. + N^{-1/2+\eta_0} + \left( \frac{s}{n} \right)^{-1/2+\eta} \right\}, \end{aligned}$$

where  $C > 0$  is some constant.

Note that since there is a trade-off between the probability and concentration bound, larger  $\eta$  or  $\eta_0$  requires a larger  $n$  to ensure the above inequality holds. Proof details of Theorem 6 are given in Appendix C.3 for the incomplete infinite-order U-statistic setting; the extension to *generalized* incomplete U-statistics follows in an identical fashion.

## 5. Discussion

The previous sections establish distributional results for random forest estimators, which take the form of *generalized* U-statistics. We showed that under mild regularity conditions, such estimators tend to a normal distribution so long as  $\frac{s}{n}(\frac{\zeta_s}{s\zeta_{1,\omega}} - 1) \rightarrow 0$ . When kernels are well-behaved, this thus implies that subsamples may be taken on the order of  $n$  while retaining the asymptotic normality of the estimator. In practice, we expect that this condition is often most naturally satisfied by subsampling at a slower rate with  $s = o(n)$  and ensuring that the corresponding variance ratio  $\zeta_s/s\zeta_{1,\omega}$  is bounded. In Section 3 we showed that the variance ratio is well-behaved for a number of nearest-neighbor-type base learners. In general though, such behavior is not well-understood, particularly for adaptive learners and such investigations present promising opportunities for future research. In Section 4 we provide Berry-Esseen bounds to quantify the proximity of these estimators to the normal distribution. Theorem 3 provides the sharpest bound to date on this rate for complete, infinite-order U-statistics, while the bounds that follow are each the first of their kind.

Throughout this work, asymptotic results for incomplete U-statistics are often based on the asymptotic result for the complete analogue. Indeed, this is how the condition on the variance ratio arises. While natural, this approach is not strictly necessary for establishing asymptotic normality of the incomplete version, particularly when the number of subsamples is quite small. Further discussion on this somewhat special case is provided in the appendix.

## Appendix A: Proofs in Section 2

Here we provide a fuller discussion of the previously established central limit theorems for randomized, incomplete, infinite-order U-statistics, paying particular attention to the relationship between the projection method utilized and the resulting subsampling rate necessary in order to retain asymptotic normality. As noted in Section 2, [25] provided one such theorem, but with somewhat strict conditions. First, the authors require that for all  $\delta > 0$ ,

$$\frac{1}{\zeta_{1,\omega}} \int_{|h_1(z)| \geq \delta \sqrt{n\zeta_{1,\omega}}} h_1^2(z) dP \rightarrow 0 \quad (n \rightarrow \infty)$$

where  $h_1(z) = \mathbb{E}[h(z, Z_2, \dots, Z_s; \omega)] - \theta$ . Note however that so long as  $\mathbb{E}[h^2(Z_1, \dots, Z_s; \omega)] < \infty$ ,

$$\frac{1}{\zeta_{1,\omega}} \int_{|h_1| \geq \delta \sqrt{n \zeta_{1,\omega}}} h_1^2(Z) dP = \int_{\left| \frac{h_1}{\sqrt{\zeta_{1,\omega}}} \right| \geq \delta \sqrt{n}} \left( \frac{h_1}{\sqrt{\zeta_{1,\omega}}} \right)^2 dP$$

automatically tends to 0 as  $n \rightarrow \infty$  and thus this condition is redundant for kernels assumed to have finite second moment.

In Section 2, we noted that there is strong reason to suspect that a subsampling rate of  $s = o(n^{1/2})$  is the largest possible when the results are established via Hájek projections. We now elaborate on that point here.

Let  $\mathcal{S}$  denote the set of all variables of the form  $\sum_{i=1}^n g_i(Z_i)$  for arbitrary measurable functions  $g_i : \mathbb{R}^d \mapsto \mathbb{R}$  with  $\mathbb{E}[g_i^2(Z_i)] < \infty$  ( $i = 1, \dots, n$ ). The Hájek projection of  $U_{n,s}$  onto  $\mathcal{S}$  is

$$\hat{U}_{n,s} = \theta + \frac{s}{n} \sum_{i=1}^n h_1(Z_i).$$

Now, by the central limit theorem for i.i.d case, we have  $\hat{U}_{n,s}/\sqrt{s^2 \zeta_1/n} \rightsquigarrow N(0, 1)$  and thus by Theorem 11.2 in [51], to obtain the asymptotic normality of U-Statistic, it is sufficient to demonstrate that  $\text{var}(U_{n,s})/\text{var}(\hat{U}_{n,s}) \rightarrow 1$ . This is straightforward when the rank of the kernel is fixed but requires more careful attention whenever  $s$  is allowed to grow with  $n$ . The variance of the U-statistic is

$$\begin{aligned} \text{var}(U_{n,s}) &= \binom{n}{s}^{-1} \sum_{\beta} \sum_{\beta'} \text{cov}(h(Z_{\beta_1}, \dots, Z_{\beta_s}), h(Z_{\beta'_1}, \dots, Z_{\beta'_s})) \\ &= \binom{n}{s}^{-1} \sum_{j=1}^s \binom{s}{j} \binom{n-s}{s-j} \zeta_j \\ &= \sum_{j=1}^s \frac{s!^2}{j!(s-j)!^2} \frac{(n-s) \cdots (n-2s+j+1)}{n(n-1) \cdots (n-s+1)} \zeta_j \end{aligned}$$

where  $\beta$  indexes subsamples of size  $s$ , and the variance of  $\hat{U}_{n,s}$  is

$$\text{var}(\hat{U}_{n,s}) = \frac{s^2}{n} \text{var}(h_1(Z_1)) = \frac{s^2}{n} \zeta_1.$$

The variance ratio is then  $\text{var}(U_{n,s})/\text{var}(\hat{U}_{n,s}) = (a_n + b_n)/c_n$ , where

$$\begin{aligned} a_n &= \frac{s^2}{n} \frac{(n-s) \cdots (n-2s+2)}{(n-1) \cdots (n-s+1)} \zeta_1, \\ b_n &= \binom{n}{s}^{-1} \sum_{j=2}^s \binom{s}{j} \binom{n-s}{s-j} \zeta_j, \\ c_n &= \frac{s^2}{n} \zeta_1. \end{aligned}$$

Thus, in order for the variance ratio to converge to 1, it suffices to show  $a_n/c_n \rightarrow 1$  and  $b_n/c_n \rightarrow 0$ . To transform these two conditions with respect to  $s$  and  $n$ , we introduce the following lemmas.

**Lemma 1** ([52]). For  $1 \leq c \leq d \leq s$ ,  $\zeta_s/c \leq \zeta_d/d$ .

**Lemma 2.** Let  $H(n, s) = \left[ \frac{(n-s) \cdots (n-2s+2)}{(n-1) \cdots (n-s+1)} \right]$ , then  $s/\sqrt{n} \rightarrow 0$  if and only if  $H(n, s) \rightarrow 1$ .

*Proof.* First, when  $s/\sqrt{n} \rightarrow 0$ , we have

$$\begin{aligned} H(n, s) &\geq \left[ \frac{n-2s+2}{n-1} \right]^{s-1} \\ &= \exp \left[ (s-1) \log \left( 1 - \frac{2s-3}{n-1} \right) \right] \\ &\approx \exp \left[ -\frac{2s^2}{n} \right] \rightarrow 1. \end{aligned}$$

Now, to prove the other direction, first note that if there exists a subsequence  $\{s'\}$  such that  $s'/\sqrt{n'} \geq c$  for some constant  $c > 0$ , then

$$\begin{aligned} H(n', s') &\leq \left[ \frac{n' - 3s'/2 + 1}{n' - s'/2} \right]^{s'-1} \\ &= \exp \left[ (s'-1) \log \left( 1 - \frac{s'-1}{n' - s'/2} \right) \right] \\ &\approx \exp \left[ -\frac{s'^2}{n'} \right] < 1. \quad \square \end{aligned}$$

Now, we can transform the conditions on  $a_n, b_n$  and  $c_n$  into conditions on  $n$  and  $s$ . Note that

$$a_n/c_n = H(n, s)$$

and

$$\begin{aligned} b_n/c_n &= \binom{n-1}{s-1}^{-1} \left\{ \sum_{j=1}^{s-1} \frac{1}{j+1} \binom{s-1}{j} \binom{(n-1)-(s-1)}{(s-1)-j} \frac{\zeta_{j+1}}{\zeta_1} \right\} \\ &= \binom{n-1}{s-1}^{-1} \left\{ \sum_{j=1}^{s-1} \binom{s-1}{j} \binom{(n-1)-(s-1)}{(s-1)-j} \frac{\zeta_{j+1}}{(j+1)\zeta_1} \right\} \\ &\geq 1 - \left[ \frac{(n-s) \cdots (n-2s+2)}{(n-1) \cdots (n-s+1)} \right] \\ &= 1 - H(n, s). \end{aligned}$$

Due to Lemma 2,  $s/\sqrt{n} \rightarrow 0$  is the necessary condition for  $b_n/c_n \rightarrow 0$  and  $a_n/c_n \rightarrow 1$ . Thus, if we utilize the Hájek projection and follow the above approach in establishing that the variance ratio converges to 1, there is no apparent way to relax the condition that  $s/\sqrt{n} \rightarrow 0$ . On the other hand, the

H-decomposition we use in Section 3 provides a finer approach and a better method for comparing the variance of  $U_{n,s}$  and  $\hat{U}_{n,s}$  thereby allowing for a faster subsampling rate to be employed.

## Appendix B: Proofs in Section 3

### B.1. H-decomposition

Distributional results for U-statistics are typically established via projection methods whereby some projection  $\hat{U}$  is shown to be asymptotically normal with  $|U - \hat{U}| \rightarrow 0$  in probability. The most popular projections are the Hájek projection and the H-decomposition. We show in Appendix A that the approach of Hájek projection always requires  $s/\sqrt{n} \rightarrow 0$  undesirably. Alternatively, the H-decomposition provides a representation of U-statistics in terms of sums of other uncorrelated U-statistics of rank  $1, \dots, s$ . The form of this decomposition presented here is derived by [32]. We illustrate those techniques in the setting of the original U-statistic  $U_{n,s}$  for simplicity and then extend them to the generalized complete U-statistic  $U_{n,s,\omega}$ . Let

$$h_c(z_1, \dots, z_c) = \mathbb{E}[h(z_1, \dots, z_c, Z_{c+1}, \dots, Z_s)] - \theta,$$

and define kernels  $h^{(1)}, h^{(2)}, \dots, h^{(s)}$  of degree  $1, \dots, s$  recursively as

$$\begin{aligned} h^{(1)} &= h_1(z_1) \\ h^{(2)} &= h_2(z_1, z_2) - h_1(z_1) - h_1(z_2) \\ &\vdots \\ h^{(s)} &= h_s(z_1, \dots, z_s) - \sum_{j=1}^{s-1} \sum_{(s,j)} h^{(j)}(z_{i_1}, \dots, z_{i_j}). \end{aligned} \tag{B.1}$$

These kernel functions have many important and desirable properties, a sample of which are enumerated in the following proposition.

**Proposition 5** ([52]). *For  $h^{(j)}$ ,  $j = 1, \dots, s$  defined as above, we have*

1. For  $c = 1, \dots, j - 1$ ,  $\mathbb{E}[h^{(j)}(z_1, \dots, z_c, Z_{c+1}, \dots, Z_j)] = 0$ .
2.  $\mathbb{E}[h^{(j)}(Z_1, \dots, Z_j)] = 0$ .
3. Let  $j < j'$  and  $S_1$  and  $S_2$  be a  $j$ -subset of  $\{Z_1, \dots, Z_n\}$  and a  $j'$ -subset of  $\{Z_1, \dots, Z_n\}$  respectively, then  $\text{cov}(h^{(j)}(S_1), h^{(j')}(S_2)) = 0$ .
4. Let  $S_1 \neq S_2$  be two distinct  $j$ -subsets of  $\{Z_1, \dots, Z_n\}$ , then  $\text{cov}(h^{(j)}(S_1), h^{(j)}(S_2)) = 0$ .

$h = h(Z_1, \dots, Z_n)$  can be written as  $h = \sum_{j=1}^s \sum_{(s,j)} h(Z_{i_1}, \dots, Z_{i_j})$  and the

expression of  $U_{n,s}$  now follows easily as

$$\begin{aligned}
U_{n,s} - \theta &= \binom{n}{s}^{-1} \sum_{(n,s)} h_s(Z_{i_1}, \dots, Z_{i_s}) \\
&= \binom{n}{s}^{-1} \sum_{(n,s)} \left\{ \sum_{j=1}^s \sum_{(s,j)} h^{(j)}(Z_{i_1}, \dots, Z_{i_j}) \right\} \\
&= \sum_{j=1}^s \binom{n}{s}^{-1} \sum_{(n,s)} \sum_{(s,j)} h^{(j)}(Z_{i_1}, \dots, Z_{i_j}) \\
&= \sum_{j=1}^s \binom{s}{j} H_n^{(j)}
\end{aligned}$$

where  $H_n^{(j)} = \binom{n}{j}^{-1} \sum_{(n,j)} h^{(j)}(Z_{i_1}, \dots, Z_{i_j})$  is itself a U-statistic, the usefulness of which lies in the fact that  $H_n^{(j)}$  ( $j = 1, \dots, n$ ) are uncorrelated and the terms in  $H_n^{(j)}$  are also uncorrelated. Because of the properties above, the variance of the kernel is

$$\text{var}(h) = \text{var} \left\{ \sum_{j=1}^s \sum_{(s,j)} h^{(j)}(Z_{i_1}, \dots, Z_{i_j}) \right\} = \sum_{j=1}^s \binom{s}{j} V_j \quad (\text{B.2})$$

where  $V_j = \text{var}(h^{(j)}(Z_{i_1}, \dots, Z_{i_j}))$ . Similarly, the variance of the U-statistic can be written as

$$\text{var}(U_{n,s}) = \text{var} \left\{ \sum_{j=1}^s \binom{s}{j} H_n^{(j)} \right\} = \sum_{j=1}^s \binom{s}{j}^2 \binom{n}{j}^{-1} V_j. \quad (\text{B.3})$$

Note that the first-order term  $sH_n^{(1)}$  is exactly the same as in the Hájek projection  $\hat{U}_{n,s}$ , but the H-decomposition provides a convenient alternative representation of U-statistics as well as their variance. In Section 3, we exploit this fact to derive a tighter and more general central limit theorem for generalized U-statistics.

## B.2. Proofs of asymptotic normality

**Proof of Theorem 1:** The generalized complete U-statistic and the base learner can be written in terms of the new kernel functions  $h^{(1)}, \dots, h^{(s)}$  defined in relation to the H-decomposition. Let  $V_{i,\omega} = \text{var}(h^{(i)})$  for  $i = 1, \dots, s-1$ ,  $V_s = \text{var}(h^{(s)})$  and define

$$V_{s,\omega} = \text{var} \left\{ h_s(Z_1, \dots, Z_s) - \sum_{j=1}^{s-1} \sum_{(s,j)} h^{(j)}(Z_{i_1}, \dots, Z_{i_j}) \right\}.$$

These new kernels  $h^{(1)}, \dots, h^{(s)}$  still retain the desirable properties in Proposition 5. Thus, similar to Eq. (B.2), Eq. (B.3), we have the following expressions for the variance of the kernel and generalized U-statistic:

$$\begin{aligned}\text{var}(h) &= \zeta_s = \sum_{j=1}^{s-1} \binom{s}{j} V_{j,\omega} + V_s, \\ \text{var}(\hat{U}_{n,s,\omega}) &= \frac{s^2}{n} V_{1,\omega} = \frac{s^2}{n} \zeta_{1,\omega}, \\ \text{var}(U_{n,s,\omega}) &= \sum_{j=1}^{s-1} \binom{s}{j}^2 \binom{n}{j}^{-1} V_{j,\omega} + \binom{n}{s}^{-1} V_s.\end{aligned}\tag{B.4}$$

The sequence  $\hat{U}_{n,s,\omega}/\sqrt{s^2\zeta_{1,\omega}/n}$  converges weakly to  $N(0, 1)$  by the central limit theorem since  $\hat{U}_{n,s,\omega} = \frac{s}{n} \sum_{i=1}^n h_1(Z_i)$  is a sum of i.i.d. random variables for each  $s$ , which satisfies Lindeberg's condition automatically. From Eq. (B.4), we have

$$\begin{aligned}\frac{\text{var}(U_{n,s,\omega})}{\text{var}(\hat{U}_{n,s,\omega})} &= \left(\frac{s^2}{n} V_{1,\omega}\right)^{-1} \left\{ \sum_{j=1}^{s-1} \binom{s}{j}^2 \binom{n}{j}^{-1} V_{j,\omega} + \binom{n}{s}^{-1} V_s \right\} \\ &\leq 1 + \left(\frac{s^2}{n} V_{1,\omega}\right)^{-1} \frac{s^2}{n^2} \left\{ \sum_{j=2}^{s-1} \binom{s}{j} V_{j,\omega} + V_s \right\} \\ &\leq 1 + \frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 1.\end{aligned}\tag{B.5}$$

Thus by Theorem 11.2 in [51], we obtain  $\frac{U_{n,s,\omega}-\theta}{\sqrt{s^2\zeta_{1,\omega}/n}} \rightsquigarrow N(0, 1)$ . ■

The following theorem is the basis for one approach to proving the central limit theorems and will be used extensively in this section.

**Theorem 7** (Lévy-Cramér Continuity Theorem). *Let  $\{F_n\}_{n \geq 1}$  and  $F_\infty$  be cdfs with characteristic functions  $\{\phi_n\}_{n \geq 1}$  and  $\phi_\infty$  respectively.*

(a) *Let  $F_n \rightsquigarrow F_\infty$ . Then, for each  $0 < K < \infty$ ,*

$$\sup_{|t| \leq K} |\phi_n(t) - \phi(t)| \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

*That is,  $\phi_n$  converges to  $\phi$  uniformly on bounded intervals.*

(b) *Let  $\lim_{n \rightarrow \infty} \phi_n(t) = \phi(t)$  exist for all  $t \in \mathbb{R}$  and let  $\phi$  be continuous at 0. Then  $\phi$  is the characteristic function of some distribution  $F$  and  $F_n \rightsquigarrow F$ .*

**Proof of Theorem 2:** Without loss of generality, let  $\theta = 0$  and observe

that

$$\begin{aligned} U_{n,s,N,\omega} &= \frac{1}{\hat{N}} \sum_{(n,s)} \rho h(Z_{i1}, \dots, Z_{is}; \omega) \\ &= \frac{N}{\hat{N}} \left[ U_{n,s,\omega} + \frac{1}{N} \sum_{(n,s)} (\rho - p) h(Z_{i1}, \dots, Z_{is}; \omega) \right] \\ &= \frac{N}{\hat{N}} [A_n + B_n], \end{aligned}$$

where  $A_n, B_n$  are uncorrelated and  $\text{var}(B_n) = (1-p)\zeta_s/N := d_{n,S,N}^2$ . Since  $N/\hat{N} \xrightarrow{p} 1$  as  $N \rightarrow \infty$ , we need only show the distributional result of  $A_n + B_n$  and apply Slutsky's theorem to obtain that of  $U_{n,s,N,\omega}$ .

First, consider the case that  $p = N/\binom{n}{s} \not\rightarrow 0$ . Since  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$ , by Theorem 1 we have  $A_n/\sqrt{s^2\zeta_{1,\omega}/n} \rightsquigarrow N(0, 1)$ . Moreover, we have

$$\frac{\text{var}(B_n)}{\text{var}(A_n)} \rightarrow \frac{N^{-1}(1-p)\zeta_s}{s^2\zeta_{1,\omega}/n} \leq \frac{n^2}{Ns^2} \cdot \frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0.$$

Thus  $U_{n,s,N,\omega}/\sqrt{s^2\zeta_{1,\omega}/n} \rightsquigarrow N(0, 1)$ , implying Eq. (3.4).

Next, consider the case that  $p \rightarrow 0$  and define

$$\begin{aligned} \phi_{A_n+B_n}(t) &:= \mathbb{E} \left[ \exp \left( it \left( \frac{s^2\zeta_{1,\omega}}{n} + \frac{\zeta_s}{N} \right)^{-\frac{1}{2}} (A_n + B_n) \right) \right] \\ &= \mathbb{E} \left[ \exp \left( it \left( \frac{s^2\zeta_{1,\omega}}{n} + \frac{\zeta_s}{N} \right)^{-\frac{1}{2}} A_n \right) \times \right. \\ &\quad \left. \mathbb{E} \left[ \exp \left( it \left( \frac{s^2\zeta_{1,\omega}}{n} + \frac{\zeta_s}{N} \right)^{-\frac{1}{2}} B_n \right) \mid Z_1, \dots, Z_n; \omega \right] \right] \\ &= \mathbb{E} \left[ \hat{\phi}_{A_n}(t) \hat{\phi}_{B_n}(t) \right]. \end{aligned}$$

According to the Theorem 7-(b), it is sufficient to show that  $\phi_{A_n+B_n} \rightarrow e^{-\frac{t^2}{2}}$ , of which the key step is to demonstrate that  $\hat{\phi}_{B_n}(t)$  is well behaved with high probability. Note that  $U_2 = \binom{n}{s}^{-1} \sum_{(n,s)} h^2(Z_{i1}, \dots, Z_{is}; \omega)$  is a complete U-statistic with kernel  $h^2$ . For any  $\delta > 0$ , by Chebyshev's inequality, we have

$$\Pr \{ |U_2 - \mathbb{E}[h^2]| \geq \delta \mathbb{E}[h^2] \} \leq \frac{s}{n} \frac{\mathbb{E}[h^4]}{\delta^2 \mathbb{E}^2[h^2]}.$$

A similar inequality holds for  $U_3 = \binom{n}{s}^{-1} \sum_{(n,s)} |h(Z_{i1}, \dots, Z_{is}; \omega)|^3$ . Let  $D = \{U_2/\mathbb{E}[h^2] \in [1-\delta, 1+\delta]\} \cap \{U_3/\mathbb{E}[|h|^3] \in [1-\delta, 1+\delta]\}$ . Since  $\mathbb{E}[|h|^{2k}]/(\mathbb{E}[|h|^k])^2$  is uniformly bounded for  $k = 2, 3$  and all  $s$  by our assumption, for

any  $\epsilon > 0$ ,  $D$  holds with probability at least  $1 - \epsilon$  for  $n$  sufficiently large. Let  $\hat{d}_{n,s,N} = [(1-p)U_2/N]^{1/2}$  and consider  $B_n/\hat{d}_{n,s,N} \mid Z_1, Z_2, \dots, Z_n; \omega$ , which is a sum of independent random variables. Consider the Lyapunov's condition, we have

$$\mathcal{L} = \frac{\binom{n}{s}^{-1} \sum_{(n,s)} |h|^3}{\left( \binom{n}{s}^{-1} \sum_{(n,s)} |h|^2 \right)^{3/2}} \frac{1-2p+2p^2}{\sqrt{N(1-p)}} = \frac{1-2p+2p^2}{\sqrt{N(1-p)}} \frac{U_3}{U_2^{3/2}}.$$

Thus, as  $N \rightarrow \infty$  and  $p \rightarrow 0$ ,

$$\mathcal{L} \leq \frac{1+\delta}{(1-\delta)^{3/2}} \frac{\mathbb{E}[|h|^3]}{\mathbb{E}^{3/2}[|h|^2]} \frac{1-2p+2p^2}{\sqrt{N(1-p)}} \rightarrow 0$$

uniformly with respect to  $Z_1, \dots, Z_n$  and  $\omega$  over  $D$ . Note that since the distance between the characteristic function of  $B_n/\hat{d}_{n,s,N}$  and that of a standard normal distribution can be controlled by  $\mathcal{L}$ , we have

$$\begin{aligned} & \left| \hat{\phi}_{B_n}(t) - \exp\left(\frac{t^2}{2} \left( (1-p) \frac{\zeta_s/N}{s^2\zeta_{1,\omega}/n + \zeta_s/N} \right) \zeta_s^{-1} U_2 \right) \right| \\ &= \left| \mathbb{E} \left[ \exp\left( t \left( (1-p) \frac{\zeta_s/N}{s^2\zeta_{1,\omega}/n + \zeta_s/N} \right)^{\frac{1}{2}} (\zeta_s^{-1} U_2)^{\frac{1}{2}} B_n / \hat{d}_{n,s,N} \right) \mid Z_1, \dots, Z_n; \omega \right] \right. \\ & \quad \left. - \exp\left( -\frac{t^2}{2} \left( (1-p) \frac{\zeta_s/N}{s^2\zeta_{1,\omega}/n + \zeta_s/N} \right) \zeta_s^{-1} U_2 \right) \right| \leq \epsilon \end{aligned}$$

over  $D$  for  $n, N$  sufficiently large. Now, let

$$\phi_B(t) = \exp\left( -\frac{t^2}{2} \left( (1-p) \frac{\zeta_s/N}{s^2\zeta_{1,\omega}/n + \zeta_s/N} \right) \right).$$

Then by the uniform continuity of  $e^x$  over any finite interval, there exists  $\delta' = O(\delta)$  such that

$$\left| \exp\left( -\frac{t^2}{2} \left( (1-p) \frac{\zeta_s/N}{s^2\zeta_{1,\omega}/n + \zeta_s/N} \right) \zeta_s^{-1} U_2 \right) - \phi_B(t) \right| \leq \delta'$$

over  $D$ . Finally, for  $n$  and  $N$  sufficiently large, we have

$$\left| \hat{\phi}_{B_n}(t) - \phi_B(t) \right| 1_D \leq (\epsilon + \delta') 1_D. \quad (\text{B.6})$$

Next, consider  $\hat{\phi}_{A_n} \cdot A_n / \sqrt{s^2\zeta_{1,\omega}/n} \rightsquigarrow N(0, 1)$  by Theorem 1, hence by Theorem 7-(a), we have

$$\left| \mathbb{E} \left[ \exp\left( it \left( \frac{s^2\zeta_{1,\omega}/n}{s^2\zeta_{1,\omega}/n + \zeta_s/N} \right)^{\frac{1}{2}} A_n / \sqrt{s^2\zeta_{1,\omega}/n} \right) \right] - e^{-\frac{t^2}{2} \frac{s^2\zeta_{1,\omega}/n}{s^2\zeta_{1,\omega}/n + \zeta_s/N}} \right| \leq \epsilon$$

for  $n$  sufficiently large. Let  $\phi_A(t) = e^{-\frac{t^2}{2} \frac{s^2 \zeta_{1,\omega}/n}{s^2 \zeta_{1,\omega}/n + \zeta_s/N}}$  and consequently, we have

$$\left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t) \right] - \phi_A(t) \right| \leq \epsilon. \quad (\text{B.7})$$

Combining Eq. (B.6) and Eq. (B.7) gives

$$\begin{aligned} |\phi_{A_n+B_n}(t) - \phi_A(t)\phi_B(t)| &= \left| \mathbb{E} \left[ \hat{\phi}_A(t)\hat{\phi}_{B_n}(t) \right] - \phi_A(t)\phi_B(t) \right| \\ &\leq \left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t)\phi_B(t)1_D \right] - \phi_A(t)\phi_B(t) \right| \\ &\quad + \left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t) \left( \hat{\phi}_{B_n}(t) - \phi_B(t) \right) 1_D \right] \right| + \epsilon \\ &\leq \left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t)\phi_B(t)1_D \right] - \phi_A(t)\phi_B(t) \right| + (\epsilon + \delta') + \epsilon \\ &\leq \left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t)1_D \right] - \phi_A(t) \right| + 2\epsilon + \delta' \quad (\text{B.8}) \\ &\leq \left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t) \right] - \phi_A(t) \right| + \left| \mathbb{E} \left[ \hat{\phi}_{A_n}(t)1_{D^c} \right] \right| + 2\epsilon + \delta' \\ &\leq \epsilon + \epsilon + 2\epsilon + \delta' \\ &= 4\epsilon + \delta'. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \phi_A(t)\phi_B(t) &= \exp \left( -\frac{t^2}{2} \left[ \frac{s^2 \zeta_{1,\omega}/n}{s^2 \zeta_{1,\omega}/n + \zeta_s/N} \right. \right. \\ &\quad \left. \left. + \left( \left( 1 - \frac{N}{\binom{n}{s}} \right) \frac{\zeta_s/N}{s^2 \zeta_{1,\omega}/n + \zeta_s/N} \right) \right] \right) \\ &= e^{-\frac{t^2}{2}} \exp \left( -\frac{t^2}{2} \left[ \frac{N}{\binom{n}{s}} \frac{\zeta_s/N}{s^2 \zeta_{1,\omega}/n + \zeta_s/N} \right] \right) \\ &\rightarrow e^{-\frac{t^2}{2}}, \end{aligned}$$

implying that for  $n, N$  sufficiently large, we have

$$\left| \phi_A(t)\phi_B(t) - e^{-\frac{t^2}{2}} \right| \leq \epsilon. \quad (\text{B.9})$$

Combining Eq. (B.8) and Eq. (B.9) yields that  $\phi_{A_n+B_n}(t) \rightarrow e^{-\frac{t^2}{2}}$ , which implies Eq. (3.4).  $\blacksquare$

Following the statement of Theorem 2 in the main text, we remarked that the form of the result provided included the condition that  $\frac{s}{n} \frac{\zeta_s}{s \zeta_{1,\omega}} \rightarrow 0$ , which thus implies that the complete analogue of the incomplete U-statistic is also asymptotically normal, but that such a condition is not necessary. We elaborate on this point here by providing an alternative result.

**Theorem 8.** Let  $Z_1, \dots, Z_n$  be i.i.d. from  $F_Z$  and  $U_{n,s,N,\omega}$  be a generalized incomplete U-statistic with kernel  $h = h(Z_1, \dots, Z_s; \omega)$ . Let  $\theta = \mathbb{E}[h]$  and  $\zeta_s = \text{var}(h)$ . Suppose that  $\mathbb{E}[|h - \theta|^{2k}] / \mathbb{E}^2[|h - \theta|^k]$  is uniformly bounded for  $k = 2, 3$  and for  $s$ . Then

1.  $U_{n,s,N,\omega} - \theta = \frac{N}{\hat{N}}(A_n + B_n)$ , where  $B_n = N^{-1} \sum_{(n,s)} (\rho - p)(h(Z_{i_1}, \dots, Z_{i_s}; \omega) - \theta)$  and  $A_n = U_{n,s,\omega} - \theta$ . If  $s/n \rightarrow 0$  and  $N \rightarrow \infty$  with  $p = N/\binom{n}{s} \rightarrow 0$ , then

$$\frac{B_n}{\sqrt{\zeta_s/N}} \rightsquigarrow N(0, 1). \quad (\text{B.10})$$

2. In addition to the conditions in 1, If  $\text{var}(A_n)/\text{var}(B_n) \rightarrow 0$ , then

$$\frac{U_{n,s,N,\omega} - \theta}{\sqrt{\zeta_s/N}} \rightsquigarrow N(0, 1). \quad (\text{B.11})$$

*Proof.* For 1, without loss of generality, let  $\theta = 0$ . Observe that

$$\begin{aligned} U_{n,s,N,\omega} &= \frac{1}{\hat{N}} \sum_{(n,s)} \rho h(Z_{i_1}, \dots, Z_{i_s}; \omega) \\ &= \frac{1}{\hat{N}} \left[ \sum_{(n,s)} (\rho - p) h(Z_{i_1}, \dots, Z_{i_s}; \omega) + U_{n,s,\omega} \right] \\ &= \frac{N}{\hat{N}} [B_n + A_n], \end{aligned}$$

where  $A_n, B_n$  are uncorrelated and  $\text{var}(B_n) = (1 - p)\zeta_s/N := d_{n,s,N}^2$ . First,  $U_2 = \binom{n}{s}^{-1} \sum_{(n,s)} h^2(Z_{i_1}, \dots, Z_{i_s}; \omega)$  is a complete U-statistic with kernel  $h^2$ . For any  $\delta > 0$ , by Chebyshev's inequality, we have

$$\Pr \{ |U_2 - \mathbb{E}[h^2]| \geq \delta \mathbb{E}[h^2] \} \leq \frac{s}{n} \frac{\mathbb{E}[h^4]}{\delta^2 \mathbb{E}^2[h^2]} \leq \frac{s}{n} \frac{C}{\delta^2}.$$

Let  $D_2 = \{ \zeta_s^{-1} U_2 \in [1 - \delta, 1 + \delta] \}$ . Thus for any  $\delta, \epsilon > 0$ ,  $\Pr(D_2) \geq 1 - \epsilon$  for  $n$  sufficiently large. Let  $\hat{d}_{n,s,N} = [(1 - p)U_2/N]^{1/2}$ , then

$$\begin{aligned} \phi_{B_n}(t) &:= \mathbb{E} [\exp(itB_n/d_{n,s,N})] \\ &= \mathbb{E} [\mathbb{E} [\exp(itB_n/d_{n,s,N}) \mid Z_1, \dots, Z_n; \omega]] \\ &= \mathbb{E} \left[ \mathbb{E} \left[ \exp \left( it(\zeta_s^{-1} U_2)^{1/2} \cdot B_n / \hat{d}_{n,s,N} \right) \mid Z_1, \dots, Z_n; \omega \right] \right], \end{aligned}$$

where  $B_n / \hat{d}_{n,s,N} \mid Z_1, Z_2, \dots, Z_n; \omega$  is a sum of independent random variables. Consider the Lyapunov's condition, we have

$$\mathcal{L} = \frac{\binom{n}{s}^{-1} \sum_{(n,s)} |h|^3}{\left( \binom{n}{s}^{-1} \sum_{(n,s)} |h|^2 \right)^{3/2}} \frac{1 - 2p + 2p^2}{\sqrt{N(1 - p)}} = \frac{1 - 2p + 2p^2}{\sqrt{N(1 - p)}} \frac{U_3}{U_2^{3/2}}$$

where  $U_3 = \binom{n}{s}^{-1} \sum_{(n,s)} |h(Z_{i_1}, \dots, Z_{i_s}; \omega)|^3$ . Since  $\mathbb{E}[h^6]/\mathbb{E}^2[|h|^3]$  is uniformly bounded, by a similar argument as for  $U_2$ ,  $D_3 = \{U_3/\mathbb{E}[|h|^3] \in [1 - \delta, 1 + \delta]\}$  holds with probability at least  $1 - \epsilon$ . Then as  $N \rightarrow \infty$  and  $p \rightarrow 0$ ,

$$\mathcal{L} \leq \frac{1 + \delta}{(1 - \delta)^{3/2}} \frac{\mathbb{E}[|h|^3]}{\mathbb{E}^{3/2}[|h|^2]} \frac{1 - 2p + 2p^2}{\sqrt{N(1 - p)}} \rightarrow 0$$

uniformly with respect to  $Z_1, \dots, Z_n$  and  $\omega$  over  $D_2 \cap D_3$ . Note that since the distance between the characteristic function of  $B_n/\hat{d}_{n,s,N}$  and that of a standard normal distribution can be controlled by  $\mathcal{L}$ , we have

$$\left| \mathbb{E} \left[ \exp \left( it(\zeta_s^{-1} U_2)^{1/2} \cdot B_n / (\hat{d}_{n,s,N}) \right) \mid Z_1, \dots, Z_n; \omega \right] - e^{-\frac{t^2}{2} \zeta_s^{-1} U_2} \right| \leq \epsilon$$

over  $D_2 \cap D_3$  for  $N$  sufficiently large. Therefore,

$$\begin{aligned} & \left| \phi_{B_n}(t) - e^{-\frac{t^2}{2}} \right| \\ & \leq \mathbb{E} \left[ \left| \mathbb{E} \left[ \exp \left( it(\zeta_s^{-1} U_2)^{\frac{1}{2}} \cdot B_n / \hat{d}_{n,s,N} \right) \mid Z_1, \dots, Z_n; \omega \right] - e^{-t^2/2} \right| \right] \\ & \leq \mathbb{E} \left[ \left( \left| \mathbb{E} \left[ \exp \left( it(\zeta_s^{-1} U_2)^{\frac{1}{2}} \cdot B_n / \hat{d}_{n,s,N} \right) \mid Z_1, \dots, Z_n; \omega \right] - e^{-\frac{t^2}{2} \zeta_s^{-1} U_2} \right| \right. \right. \\ & \quad \left. \left. + \left| e^{-\frac{t^2}{2} \zeta_s^{-1} U_2} - e^{-\frac{t^2}{2}} \right| \right) 1_{D_2 \cap D_3} \right] + 4\epsilon \\ & \leq \mathbb{E} \left[ \left( \epsilon + \left| e^{-\frac{t^2}{2} \zeta_s^{-1} U_2} - e^{-\frac{t^2}{2}} \right| \right) 1_{D_2 \cap D_3} \right] + 4\epsilon \\ & \leq (\epsilon + \delta') + 4\epsilon. \end{aligned}$$

Since  $\epsilon$  and  $\delta$  can be arbitrarily small, we have  $\phi_{B_n}(t) \rightarrow e^{-\frac{t^2}{2}}$  and thus  $B_n/d_{n,s,N} \rightsquigarrow N(0, 1)$ , which implies Eq. (B.10) since  $N/\binom{n}{s} \rightarrow 0$ .

For **2**,  $\text{var}(A_n)/\text{var}(B_n) \rightarrow 0$  implies that  $\text{var}(A_n/\sqrt{\zeta_s/N}) = o(1)$ . Thus, we have

$$\frac{U_{n,s,N,\omega}}{\sqrt{\zeta_s/N}} = \frac{B_n}{\sqrt{\zeta_s/N}} + o_p(1),$$

which implies Eq. (B.11) by applying Slutsky's theorem.  $\square$

Part 1 of Theorem 8 gives that  $B_n$ , the difference between the incomplete and complete generalized U-statistics, is asymptotically normal under quite weak conditions so long as the number of subsamples  $N$  grows slower than  $\binom{n}{s}$ . In particular, no specialized conditions on the resulting variance or variance ratio are required. In Part 2, to establish asymptotic normality of the generalized incomplete U-statistic itself, we do not require the original condition on the variance ratio given in Theorem 2, though we do require that  $\text{var}(A_n)/\text{var}(B_n) \rightarrow 0$ . Such a condition remains difficult to verify for general kernels, but can always be satisfied, for example, by taking  $N = o(n/s)$ . Indeed, note that

$$\frac{s^2}{n} \zeta_{1,\omega} \leq \text{var}(A_n) \leq \frac{s}{n} \zeta_s,$$

and  $\text{var}(B_n) = (1-p)\zeta_s/N$ , thus a sufficient condition for  $(\zeta_s/N)^{-1/2}A_n = o_p(1)$  to hold is letting  $N = o(n/s)$ . Thus, asymptotic normality for incomplete U-statistics can be established without requiring normality of the complete version and in particular, without requiring the specialized condition on the variance ratio discussed at length in Section 3.

### B.3. Simple base learner variance ratios

We now provide explicit calculations for the variance ratios corresponding to the various base learners discussed in Section 3. We begin with simple examples where base learners take the form of a sample mean, sample variance, and least squares estimators.

**Example 1** (sample mean). *Suppose  $Z_1, \dots, Z_s$  are i.i.d. random variables with mean  $\mu$  and let  $h = \bar{Z}$ , then*

$$\frac{\zeta_s}{s\zeta_1} = \frac{\text{var}(\sum_{i=1}^s Z_i/s)}{\text{svar}(Z_1/s + (s-1)\mu/s)} = 1. \quad (\text{B.12})$$

Eq. (B.12) holds for any estimators that can be written as a sum of i.i.d. random variables. In such cases, nothing is lost after projecting since  $\hat{h} = h$ .

**Example 2** (sample variance). *Suppose  $Z_1, \dots, Z_s$  are i.i.d. random variables with variance  $\sigma^2$  and fourth central moment  $\mu_4$  and consider the sample variance  $h = \binom{n}{2}^{-1} \sum_{i < j} (Z_i - Z_j)^2$ . Then as  $s \rightarrow \infty$ ,*

$$\frac{\zeta_s}{s\zeta_1} = 1 + \frac{2}{(s-1)} \cdot \frac{\sigma^4}{\mu_4 - \sigma^4} \rightarrow 1.$$

The kernel  $h$  can be written as  $h = \sum_{i=1}^s \frac{(Z_i - \bar{Z})^2}{s-1} \approx \sum_{i=1}^s \frac{(Z_i^2 - \mu^2)}{s-1}$ . Since  $\bar{Z}$  is much more stable than  $Z_i$ ,  $h$  is close to a sum of i.i.d random variables.

**Example 3** (OLS estimator). *Let  $Z_1, \dots, Z_s$  denote i.i.d. pairs of random variables  $(X_i, Y_i)$  and  $Y_i = X_i^T \beta + \epsilon_i$ . Suppose that  $\epsilon_i$  has mean 0 and variance  $\sigma^2$ , and  $\epsilon_i$  is independent of  $X_i$ . Let  $h = (X^T X)^{-1} X^T Y$  be the ordinary least squares (OLS) estimator of  $\beta$ . Then as  $s \rightarrow \infty$ ,*

$$(s\zeta_1)^{-1} \zeta_s \rightarrow I,$$

where  $I$  is the identity matrix.

*Proof.* Let  $\hat{\beta} = GX^T Y$  be the OLS estimator, where  $G = (X^T X)^{-1}$ , then  $\zeta_s = \text{var}(\hat{\beta}) = \mathbb{E}[G]\sigma^2$ . Since  $X_i$  and  $\epsilon_i$  are independent, we have  $\mathbb{E}[\hat{\beta} | X_1, Y_1] = \beta + \mathbb{E}_1[G]X_1 \cdot \epsilon$ , where  $\mathbb{E}_1$  takes the expectation conditioning on  $X_1$ . Then,

$$\zeta_1 = \text{var}(\mathbb{E}[\hat{\beta} | X_1, Y_1]) = \mathbb{E}[\mathbb{E}_1[G]X_1 X_1^T \mathbb{E}_1[G]]\sigma^2.$$

According to the law of large numbers,  $\frac{1}{s} \sum_{i=1}^s X_i X_i^T \rightarrow \Sigma$  a.s., as  $s \rightarrow \infty$  where  $\mathbb{E}[X_i X_i^T] = \Sigma$ . Letting  $\Sigma^{-1} = \Omega$ , we have

$$s \cdot \mathbb{E}[(X^T X)^{-1}] = \mathbb{E} \left[ \left( \frac{1}{s} \sum_{i=1}^s X_i X_i^T \right)^{-1} \right] \rightarrow \Omega,$$

and hence  $s\zeta_s \rightarrow \Omega$ . Furthermore, we have

$$sG \mid X_1 = \left( \frac{1}{s} \left[ X_1 X_1^T + \sum_{i \neq 1} X_i X_i^T \right] \right)^{-1} \mid X_1 \rightarrow \Omega, \quad a.s.$$

Thus,

$$\begin{aligned} s^2 \cdot \zeta_1 / \sigma^2 &= \mathbb{E} [\mathbb{E}_1[G] X_1 X_1^T \mathbb{E}_1[G]] \\ &= \mathbb{E} [\mathbb{E}_1[sG] \cdot X_1 X_1^T \cdot \mathbb{E}_1[sG]] \\ &\rightarrow \Omega \cdot \Sigma \cdot \Omega \\ &= \Omega, \end{aligned}$$

and hence  $\zeta_1$  is of order  $s^{-2}$ . Therefore,  $(s\zeta_1)^{-1} \zeta_s \rightarrow I$ .  $\square$

Here again, note that  $h = \sum_{i=1}^s (X^T X)^{-1} X_i Y_i$ , which is still close to a sum of i.i.d. random variables. These three examples suggest that perhaps for many common base learners,  $\zeta_s$  is of order  $s^{-1}$  and  $\zeta_1$  of order  $s^{-2}$ ; essentially, each individual observation explains roughly  $s^{-1}$  times the variance of the base learner.

**Proof of Proposition 1:** Denote the kNN estimator at  $x$  as  $\varphi(x)$ . Let  $A_i$  be the event that  $X_1$  is the  $i$ th closest point to the target point  $x$  and  $B = \cup_{i=1}^k A_i$ . First, by the continuity of  $f$  at  $x$ , we have  $\text{var}(\varphi(x)) \rightarrow \sigma^2/k$  as  $s \rightarrow \infty$ . Let  $X_1^*, \dots, X_k^*$  be the  $k$ -NNs of  $x$ . Then

$$\begin{aligned} \mathbb{E}[\varphi(x) \mid X_1, Y_1] &= \frac{1}{k} \mathbb{E} \left[ \mathbb{1}_B \left[ Y_1 + \sum_{i=2}^k Y_i^* \right] + \mathbb{1}_{B^c} \left[ \sum_{i=1}^k Y_i^* \right] \mid X_1, Y_1 \right] \\ &= \frac{\epsilon_1}{k} \mathbb{E}[\mathbb{1}_B \mid X_1] + \frac{1}{k} \left[ \sum_{i=1}^k f(X_i^*) \mid X_1 \right], \end{aligned}$$

and thus  $\text{var}(\mathbb{E}[\varphi(x) \mid X_1, Y_1]) \geq \frac{\sigma^2}{k^2} \mathbb{E}[\text{Pr}^2(B \mid X_1)]$ . Next,

$$\begin{aligned} \mathbb{E}[\text{Pr}^2(B \mid X_1)] &= \mathbb{E} \left[ \left( \sum_{i=1}^k \text{Pr}(A_i \mid X_1) \right)^2 \right] \\ &= \sum_{i=1}^k \sum_{j=1}^k \mathbb{E} [\text{Pr}(A_i \mid X_1) \text{Pr}(A_j \mid X_1)] \\ &= \frac{1}{2s-1} \sum_{i=1}^k \sum_{j=1}^k \frac{\binom{s-1}{i-1} \binom{s-1}{j-1}}{\binom{2s-2}{i+j-2}} \end{aligned}$$

$$= \frac{V(k, s)}{2s - 1}, \quad (\text{B.13})$$

where  $V(k, s) = \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \left[ \frac{\binom{s-1}{i} \binom{s-1}{j}}{\binom{2s-2}{i+j}} \right]$ . We have

$$\begin{aligned} V(k) &= \lim_{s \rightarrow \infty} V(k, s) \\ &= \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \frac{(i+j)!}{i!j!} \frac{1}{2^{i+j}} \\ &= \sum_{c=0}^{2k-2} \sum_{i+j=c, 0 \leq i, j \leq k-1} \frac{(i+j)!}{i!j!} \frac{1}{2^{i+j}}. \end{aligned}$$

We can obtain  $k \leq V(k) < 2k - 1$  by simply observing that

$$\sum_{c=0}^{k-1} \sum_{i+j=c} \frac{c!}{i!j!} \frac{1}{2^c} \leq V(k) \leq \sum_{c=0}^{2k-2} \sum_{i+j=c} \frac{c!}{i!j!} \frac{1}{2^c}.$$

Thus,

$$\begin{aligned} \limsup_{s \rightarrow \infty} \frac{\zeta_s}{s\zeta_1} &= \limsup_{s \rightarrow \infty} \frac{\text{var}(\varphi(x))}{s \text{var}(\mathbb{E}[\varphi(x) \mid X_1, Y_1])} \\ &\leq \limsup_{s \rightarrow \infty} \frac{2s-1}{s} \frac{k}{V(k, s)} \\ &= c(k), \end{aligned} \quad (\text{B.14})$$

and  $1 < c(k) \leq 2$ . ■

Note that Proposition 1 holds without imposing any conditions on the regression function  $f$  or the distribution of  $X$ . To see why, note from the proof that both  $\zeta_s$  and  $\zeta_1$  can each be decomposed into two terms, one of which comes from the variation of the regression function while the other is due to the variation of the noise. Here, since  $k$  is fixed, the term involving the variation in the regression function is small relative to the noise term for large  $s$ . When  $k$  grows with  $s$ , more care must be taken in assessing the contribution of the regression function. [53] (Theorem 15.3) discuss the convergence rate of the variance of kNN estimators. This result could potentially enable results to be established for more general nearest neighbor estimators where the number of neighbors  $k$  is permitted to grow with  $s$ , though we do not explore this further here.

**Proof of Proposition 2:** First let  $\tilde{\varphi} = \sum_{i=1}^s w(i, x, \mathbf{X})f(X_i)$  and note that

$$\begin{aligned} \text{var}(\varphi) &= s\mathbb{E}[w^2(1, x, \mathbf{X})]\sigma^2 + \text{var}(\tilde{\varphi}) \\ &\leq \sigma^2 + \|f\|_\infty^2/4. \end{aligned} \quad (\text{B.15})$$

Next,  $\mathbb{E}[\varphi | X_1, \epsilon_1] = \mathbb{E}[\tilde{\varphi} | X_1] + \epsilon_1 \mathbb{E}[w(1, x, \mathbf{X}) | X_1]$ , and thus

$$\begin{aligned} \text{var}(\mathbb{E}[\varphi | X_1, \epsilon_1]) &= \text{var}(\mathbb{E}[\tilde{\varphi} | X_1]) + \sigma^2 \mathbb{E}[\mathbb{E}^2[w(1, x, \mathbf{X}) | X_1]] \\ &\geq \sigma^2 \mathbb{E}[\mathbb{E}^2[w(1, x, \mathbf{X}) | X_1]] \\ &\geq \sigma^2 \mathbb{E}^2[w(1, x, \mathbf{X})] \\ &= \sigma^2/s^2. \end{aligned} \tag{B.16}$$

Therefore,

$$\begin{aligned} \limsup_{s \rightarrow \infty} \frac{\zeta_s}{s\zeta_1} \frac{1}{s} &= \limsup_{s \rightarrow \infty} \frac{s\sigma^2 \mathbb{E}[w^2(1, x, \mathbf{X})] + \text{var}(\tilde{\varphi})}{s^2 (\sigma^2 \mathbb{E}[\mathbb{E}^2[w(1, x, \mathbf{X}) | X_1]] + \text{var}(\mathbb{E}[\tilde{\varphi} | X_1]))} \tag{B.17} \\ &\leq \frac{\sigma^2 + \|f\|_\infty^2/4}{\sigma^2} \\ &< \infty. \end{aligned} \quad \blacksquare$$

We emphasize that the inequalities in Eq. (B.15) and Eq. (B.16) are generally quite loose in order to cover the worst case scenario. As seen in Proposition 1, the order of  $\zeta_1$  can indeed be  $s^{-1}$  rather than  $s^{-2}$ . Nonetheless, Proposition 2 indicates that  $s/\sqrt{n} \rightarrow 0$  is sufficient to ensure that  $\frac{s}{n} \frac{\zeta_s}{s\zeta_1} \rightarrow 0$ .

#### B.4. Variance ratios for RP trees

Now, we turn to the analysis for RP trees, which form predictions by taking a sample uniformly at random from the potential nearest neighbors and averaging the corresponding response values. We begin with a simpler result for base learners that take a naive random average across  $k$  response values selected uniformly at random from the entire dataset.

**Example 4** (naive random average). *Let  $Z_1, \dots, Z_s$  denote i.i.d. pairs of random variables  $(X_i, Y_i)$ . For any target point  $x$ , let  $\varphi(x)$  denote the estimator that forms a prediction by simply selecting  $k$  sample points uniformly at random (without replacement) and averages the selected response values, so that we can write  $\varphi(x) = \frac{1}{k} \sum_{i=1}^s \xi_i Y_i$ , where*

$$\xi_i = \begin{cases} 1, & i^{\text{th}} \text{ sample is selected} \\ 0, & i^{\text{th}} \text{ sample is not selected.} \end{cases}$$

Then  $\zeta_s = \text{var}(Y_1)/k$  and  $\zeta_{1,\omega} = (\text{var}(Y_1))/s^2$ , so that  $\zeta_s/s\zeta_{1,\omega} = s/k$ .

Note in the above example that when  $k$  is fixed,  $s/\sqrt{n} \rightarrow 0$  is sufficient to ensure that  $\frac{s}{n} \frac{\zeta_s}{s\zeta_{1,\omega}} \rightarrow 0$ . However, when  $k$  is assumed to grow with  $n$ , the subsample size  $s$  can grow more quickly. In the adaptive case, where  $w(i, x, \mathbf{X})$  may depend on  $\{Y_i\}_{i=1}^s$ , tree estimators with small terminal node sizes may look less like a linear statistic and in turn may have a larger variance ratio. However, as discussed, for non-adaptive estimators like kNN, the ratio is bounded by a

constant. In this way, well-behaved tree predictors can be seen as similar to kNN and are still more easily controlled than RP trees.

We turn now to the proving Proposition 3, namely that for base learners that are RP trees,

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s / s \zeta_{1,\omega}}{(\log s)^{2p-2}} < \infty.$$

**Proof of Proposition 3:** Denote the RP tree by  $T$  and the set of  $k$ -PNNs by  $\Xi$ . We have

$$\text{var}(T) = \text{var}(\tilde{T}) + \sigma^2/k \leq \sigma^2/k + \|f\|_\infty^2/4,$$

where  $\tilde{T}$  is the RP tree prediction in the noiseless case. Let  $|\Xi|$  denote the number of  $k$ -PNNs of  $x$ , then

$$\mathbb{E}[T | Z_1] = \epsilon_1 \mathbb{E} \left[ \frac{1}{|\Xi|} \mathbf{1}_{X_1 \in \Xi} | X_1 \right] + \mathbb{E}[\tilde{T} | X_1].$$

Since  $\Xi$  is independent of  $\epsilon$ , we have  $\text{var}(\mathbb{E}[T | Z_1]) \geq \sigma^2 \mathbb{E}[\mathbb{E}^2[S_1 | X_1]]$ , where  $S_1 = \frac{1}{|\Xi|} \mathbf{1}_{X_1 \in \Xi}$ . Note that  $\mathbb{E}[S_1] = \mathbb{E}[\sum_{i=1}^s S_i]/s = 1/s$ , and thus we have  $\text{var}(\mathbb{E}[T | Z_1]) \geq \sigma^2/s^2$ . Furthermore, we have

$$\begin{aligned} \mathbb{E}[S_1 | Z_1] &= \Pr_1(X_1 \in \Xi) \mathbb{E}_1 \left[ \frac{1}{|\Xi|} | X_1 \in \Xi \right] \\ &= \sum_{i=0}^{k-1} \binom{s-1}{i} u^i (1-u)^{s-1-i} \cdot \mathbb{E}_1 \left[ \frac{1}{|\Xi|} | X_1 \in \Xi \right] \\ &= \text{I} \cdot \text{II}, \end{aligned} \tag{B.18}$$

where  $\Pr_1(\cdot) = \Pr(\cdot | X_1)$ ,  $\mathbb{E}_1 = \mathbb{E}(\cdot | X_1)$ ,  $u = \Pr_1(X_i \in R)$  and  $R$  is the hyper-rectangle defined by  $x$  and  $X_1$ . Conditioning on  $X_1 \in \Xi$ , define the conditional probability function of  $(X_2, \dots, X_s) - \Pr_1(\cdot | X_1 \in \Xi)$  as  $\tilde{\Pr}_1$ . For I,

$$\begin{aligned} \text{I}^2 &= \left[ \sum_{i=0}^{k-1} \binom{s-1}{i} u^i (1-u)^{s-1-i} \right]^2 \\ &= \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \binom{s-1}{i} \binom{s-1}{j} u^{i+j} (1-u)^{2s-2-i-j}, \end{aligned}$$

thus

$$\begin{aligned} \mathbb{E}[\text{I}^2] &= \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \binom{s-1}{i} \binom{s-1}{j} \mathbb{E} [u^{i+j} (1-u)^{2s-2-i-j}] \\ &= \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \frac{\binom{s-1}{i} \binom{s-1}{j}}{\binom{2s-2}{i+j}} \mathbb{E} \left[ \frac{u^{i+j} (1-u)^{2s-2-i-j}}{\text{B}(i+j+1, 2s-1-i-j)} \right] \end{aligned}$$

$$= \frac{1}{2s-1} \cdot \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \frac{\binom{s-1}{i} \binom{s-1}{j}}{\binom{2s-2}{i+j}} G(i, j), \quad (\text{B.19})$$

where  $u = \Pr_1(X_i \in R) \in (0, 1)$ . If  $u \sim \text{Uniform}(0, 1)$ , then  $G(i, j) = 1$  and (B.19) reduces to (B.13). Let the probability density function of  $u$  be  $p(u)$ , and the probability density function of beta distribution with shape parameters  $\alpha$  and  $\beta$  be  $g(u, \alpha, \beta)$ , then  $G(i, j) = \int_0^1 g(u, \alpha, \beta) p(u) du$ , where  $\alpha = i + j, \beta = s - 1 - \alpha$ . Since  $i + j \leq 2k - 2$ , when  $s \rightarrow \infty$ ,  $g(u, \alpha, \beta)$  is almost singular at  $u = 0$ . Moreover, we can find that at around  $u = 0$ ,  $p(u) \geq 1$ . Thus, there exists some  $c_1 > 0$  such that  $G(i, j) \geq c_1$ . Therefore, we have

$$\mathbb{E}[I^2] \geq \frac{c_1}{2s-1} V(k, s), \quad V(k, s) = \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} \left[ \frac{\binom{s-1}{i} \binom{s-1}{j}}{\binom{2s-2}{i+j}} \right]. \quad (\text{B.20})$$

For II, by Jensen's inequality, we have  $\text{II} = \tilde{\mathbb{E}}_1[1/|\Xi|] \geq 1/\tilde{\mathbb{E}}_1[|\Xi|]$ , and then

$$\mathbb{E}[\mathbb{E}^2[S_1|X_1]] \geq \mathbb{E} \left[ I^2 \cdot \frac{1}{\tilde{\mathbb{E}}_1^2[|\Xi|]} \right]. \quad (\text{B.21})$$

$\tilde{\mathbb{E}}_1[|\Xi|]$  is just the expected number of  $k$ -PNNs conditioning on  $X_1 \in \Xi$ , or equivalently given that there are at most  $k - 1$  sample points in  $R$ . [10] showed that  $\mathbb{E}[|\Xi|]$  is of order  $k(\log s)^{p-1}$  when the probability density function of the features is bounded away from 0 and  $\infty$  in  $[0, 1]^p$ . Since  $k$ -PNN depends only on the relative distance, it can be shown that there exists some  $c_2 > 0$  such that

$$\tilde{\mathbb{E}}_1[|\Xi|] \leq c_2 \mathbb{E}[|\Xi|], \quad \text{for } X_1 \in [0, 1]^p. \quad (\text{B.22})$$

Note that  $V(k, s) \geq k$ . Combining Eq. (B.20), Eq. (B.21) and Eq. (B.22), we have

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s / s \zeta_{1,\omega}}{(\log s)^{2p-2}} \leq \limsup_{s \rightarrow \infty} \frac{(\sigma^2/k + \|f\|_\infty^2/4)(\log s)^{2-2p}}{s \left( c_1 \frac{V(k,s)}{2s-1} \cdot c_2^{-2} (k(\log s)^{p-1})^{-2} \cdot \sigma^2 \right)} < \infty, \quad (\text{B.23})$$

thus achieving what was claimed in Proposition 3.  $\blacksquare$

### B.5. Variance ratio for non-adaptive $k$ -trees

Finally, we turn to the analysis for non-adaptive  $k$ -trees, which form predictions by taking a sample non-adaptively from the potential nearest neighbors and averaging the corresponding response values.

**Proof of Proposition 4:** Denote a non-adaptive  $k$ -tree by  $T = T(Z_1, \dots, Z_s; \omega)$  and the set of  $2k$ -PNNs by  $\Xi$ . We have

$$\text{var}(T) = \text{var}(\tilde{T}) + \sigma^2/k \leq \sigma^2/k + \|f\|_\infty^2/4,$$

where  $\tilde{T}$  is the tree prediction in the noiseless case. Let  $L(x; \omega)$  be the terminal node that  $x$  belongs to, then  $\mathbb{E}[T | Z_1] = \epsilon_1 \mathbb{E} \left[ \frac{1}{|L(x; \omega)|} \mathbb{1}_{X_1 \in L(x; \omega)} | X_1 \right] + \mathbb{E}[\tilde{T} | X_1]$ . Since  $L(x; \omega)$  is independent of  $\epsilon$ , we have

$$\text{var}(\mathbb{E}[T | Z_1]) \geq \sigma^2 \mathbb{E} \left[ \mathbb{E}^2 \left[ \frac{1}{|L(x; \omega)|} \mathbb{1}_{X_1 \in L(x; \omega)} | X_1 \right] \right] = \sigma^2 \mathbb{E}[S_1^2], \quad (\text{B.24})$$

where  $S_1 = \mathbb{E}[\frac{1}{|L(x; \omega)|} \mathbb{1}_{X_1 \in L(x; \omega)} | X_1]$ . Note that  $\mathbb{E}[S_1] = \mathbb{E}[\sum_{i=1}^s S_i] / s = 1/s$ , and thus we have  $\text{var}(\mathbb{E}[T | Z_1]) \geq \sigma^2/s^2$ . Furthermore, since  $L(x; \omega) \leq 2k - 1$ , then  $L(x; \omega) \subset 2k$ -PNNs and we have

$$\begin{aligned} S_1 &= \mathbb{E} \left[ \frac{1}{|L(x; \omega)|} \mathbb{1}_{X_1 \in L(x; \omega)} | X_1 \right] \\ &\leq \frac{1}{k} \mathbb{E}[\mathbb{1}_{X_1 \in \Xi} | X_1] = \frac{1}{k} \cdot \mathbf{I}, \end{aligned} \quad (\text{B.25})$$

where  $\mathbf{I}$  is the same as in Eq. (B.20). We get rid of the extra randomness  $\omega$  here by upper bounding  $|L(x; \omega)|$  by  $k$ . From the proof of Proposition 3, we know  $\mathbb{E}[\mathbf{I}^2] = O(1/s)$  and according to [10], we have  $\mathbb{E}[\mathbf{I}] = O(2k(\log s)^{p-1}/s)$ . Note that for any random variable  $X$  and any real number  $x$ , we always have

$$\begin{aligned} \mathbb{E}[X^2] &\geq \mathbb{E}[X^2; X \geq x] \\ &= \Pr(X \geq x) \mathbb{E}[X^2 | X \geq x] \\ &\geq \Pr(X \geq x) \mathbb{E}^2[X | X \geq x] \\ &= \frac{\mathbb{E}^2[X; X \geq x]}{\Pr(X \geq x)}. \end{aligned} \quad (\text{B.26})$$

Now, let  $x = s^{-(1+\alpha)}$  for any  $\alpha > 0$ , then  $\Pr(\mathbf{I} \geq s^{-(1+\alpha)}) \leq \frac{\mathbb{E}^2[\mathbf{I}; \mathbf{I} \geq s^{-(1+\alpha)}]}{\mathbb{E}[\mathbf{I}^2]} = O((\log s)^{2p-2}/s)$ . Then, by applying Eq. (B.26) again, we have

$$\mathbb{E}[S_1^2] \geq \frac{\mathbb{E}^2[S_1; S_1 \geq k^{-1}s^{-(1+\alpha)}]}{\Pr(S_1 \geq k^{-1}s^{-(1+\alpha)})}. \quad (\text{B.27})$$

Since  $\mathbb{E}[S_1] = 1/s$ ,  $\mathbb{E}^2[S_1; S_1 \geq k^{-1}s^{-(1+\alpha)}] = O(1/s^2)$ , and hence  $\mathbb{E}[S_1^2] \geq \frac{O(1/s)}{\Pr(\mathbf{I} \geq s^{-(1+\alpha)})} = O(1/s(\log s)^{2p-2})$ . Thus,  $\text{var}(\mathbb{E}[T|Z_1]) \geq \sigma^2 O(1/(s \log s)^{2p-2})$ . Therefore,

$$\limsup_{s \rightarrow \infty} \frac{\zeta_s/s\zeta_1}{(\log s)^{2p-2}} \leq \limsup_{s \rightarrow \infty} \frac{(\sigma^2/k + \|f\|_\infty^2/4)(\log s)^{2-2p}}{s(s^{-1}((\log s)^{p-1})^{-2} \cdot \sigma^2)} < \infty, \quad (\text{B.28})$$

thus achieving what was claimed in Proposition 4. ■

Note that the proof strategy here follows closely to that in [26]; the main difference above is in how we obtain  $\Pr(\mathbf{I} \geq s^{-(1+\alpha)})$ .

### B.6. Simulation

We present here a small simulation study in order to illustrate the behavior of the variance ratio of when the base learners employed are decision trees or nearest neighbor methods, as well to visually demonstrate the asymptotic normality derived in Section 3. We consider the underlying regression function:

$$y = 1.0 \sin(\pi x_1 x_2) + 0.2(x_3 - 0.5)^2 + 0.5x_4 + 0.1x_5 + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 0.3), \quad (\text{B.29})$$

where  $(x_1, \dots, x_5) \sim \mathcal{N}(\mu, \Sigma)$ ,  $\mu = (0.2, 0.3, 0.2, 0.7, 0.4)$ ,  $\Sigma_{(i,i)} = 1$  and  $\Sigma_{(i,j)} = 0.2$  for  $1 \leq i, j \leq 5$ . This function has slightly different forms, but was originally used in the MARS paper [54] and has since been frequently used in more recent work exploring the behavior of random forests [14, 25, 37]. We consider the prediction at  $x_0 = (0.5, 0.5, 0.5, 0.5, .0.5)$  for the base learners.  $\zeta_s$  and  $\zeta_{1,\omega}$  are obtained separately. For  $\zeta_s$ , we generate  $n_{\text{out}}$  independent samples of size  $s$ , build base learners on each, and then use the sample variance of the predictions as our estimate. The parameter  $\zeta_{1,\omega} = \text{var}(\mathbb{E}[h | Z_1])$  is estimated in two steps. We begin by sampling one observation  $Z_1$ , which we refer to as the initial fixed point. We then generate several samples of size  $s - 1$  and combine them with  $Z_1$  to build a base learner and then record the mean of the predictions at  $x_0$ . Let  $n_{\text{in}}$  denote the number of samples drawn so that this average is taken over  $n_{\text{in}}$  predictions. We then repeat the process for  $n_{\text{out}}$  initial sets of fixed points and take our final estimate of  $\zeta_{1,\omega}$  as the variance over the  $n_{\text{out}}$  final averages, yielding the estimator.

It is worth pausing for a moment to note that previous work has suggested that a larger  $n_{\text{in}}$  seems to provide better results even when  $n_{\text{out}}$  is relative small [25]. An explanation for this is as follows. Denote  $\mathbb{E}[h | Z_1 = z] = h_1(z)$  and note that the inner approximation is trying to estimate  $h_1(z)$ . The outer approximation is intended to estimate  $\text{var}(h_1(Z))$ , but it is actually estimating  $\text{var}(\hat{h}_1(Z))$ . By the law of total variance, we have

$$\begin{aligned} \text{var}_n(\hat{h}_1(Z)) &= \text{var}_n(\mathbb{E}[\hat{h}_1(Z) | Z]) + \mathbb{E}_n[\text{var}(\hat{h}_1(Z) | Z)] \\ &= \text{var}_n(h_1(Z)) + \mathbb{E}_n[\text{var}(\hat{h}_1(Z) | Z)], \end{aligned} \quad (\text{B.30})$$

where  $\text{var}_n(\cdot)$  denotes sample variance and  $\mathbb{E}_n(\cdot)$  denotes sample mean. Then

$$\begin{aligned} \frac{\mathbb{E}[\text{var}_n(\hat{h}_1(Z))]}{\text{var}(h_1(Z))} &= 1 + \frac{\mathbb{E}[\text{var}(\hat{h}_1(Z) | Z)]}{\text{var}(h_1(Z))} \\ &= 1 + \frac{\mathbb{E}[\text{var}(\hat{h}_1(Z) | Z)]}{\zeta_{1,\omega}}. \end{aligned} \quad (\text{B.31})$$

On the other hand, note that  $\text{var}(\hat{h}_1(Z) | Z) = \text{var}(h | Z_1 = Z)/n_{\text{in}}$ . When  $s$  is large, knowing one observation will not influence the final prediction by much and so  $\text{var}(h | Z_1 = Z) \approx \zeta_s$ . Therefore, the right hand side of Eq. (B.31) is close to  $1 + \frac{s}{n_{\text{in}}} \frac{\zeta_s}{s\zeta_{1,\omega}}$ . Since  $\frac{\zeta_s}{s\zeta_{1,\omega}} \geq 1$ , to make  $\text{var}_n(\hat{h}_1(Z))$  less biased, we need to make  $n_{\text{in}}$  large enough relative to  $s$ .

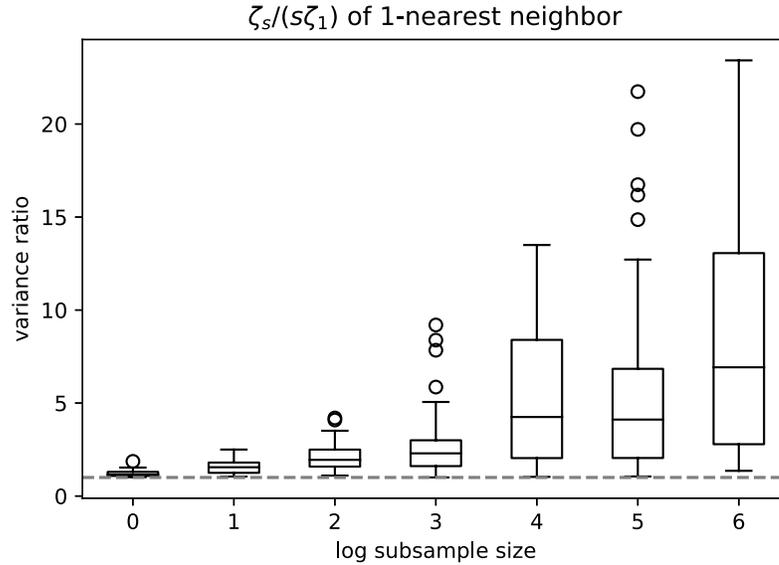


FIG 2. Variance ratio of 1-NN of different subsample size. We estimate  $\zeta_s/s\zeta_{1,\omega}$  30 times to obtain the boxplots. We can see that as the sample size increases, the variance ratio appears to slowly get larger.

To carry out these simulations with kNN base learners, we use the Python package `sklearn.neighbors.KNeighborsRegressor` with the distance metric is set to as “minkowski”. Fig. 2 uses 1-nearest neighbor base learners and shows how the estimated variance ratio changes with (log) subsample size. Here  $n_{\text{in}} = 1000$  and  $n_{\text{out}} = 50$  and we obtain the boxplots of the ratio  $\zeta_s/s\zeta_1$  for  $s = 2^i$ , for  $i = 1, \dots, 7$ . We see that as  $s$  increases, so does the variance in our estimation. This is because  $\zeta_1$  is getting smaller and becoming harder to estimate. However, if look at the median of the estimated variance ratio  $\zeta_s/s\zeta_1$ , we see that it increases as  $s$  increases at a rate roughly on the order of  $\log(s)$ .

In Fig. 3, we fix the subsample size  $s$  to be 300 and consider  $k$ -nearest neighbor base learners with differing values of  $k$ , where  $k = 2^i$  for  $i = 1, \dots, 7$ . Here  $n_{\text{in}} = 1000$  and  $n_{\text{out}} = 50$ . As expected, when  $k$  increases,  $\zeta_s/s\zeta_1$  decreases and gets closer to 1. Again, as  $k$  increases,  $\zeta_1$  gets larger, so the estimation of  $\zeta_1$  also becomes more accurate, which can be seen from the trend in the variation of the estimates.

To carry the same kinds of simulations with decision trees, we use the Python package `sklearn.tree.DecisionTreeRegressor`. We let the minimum number of samples required to split an internal node be  $2k$  and the minimum number of samples required in a (terminal) leaf node be  $k$  so that the terminal node size is always bounded between  $k$  and  $2k - 1$ . The number of features randomly made eligible for splitting at each internal node – often referred to as the `mtry` parameter – is set to the square root of total number of features.

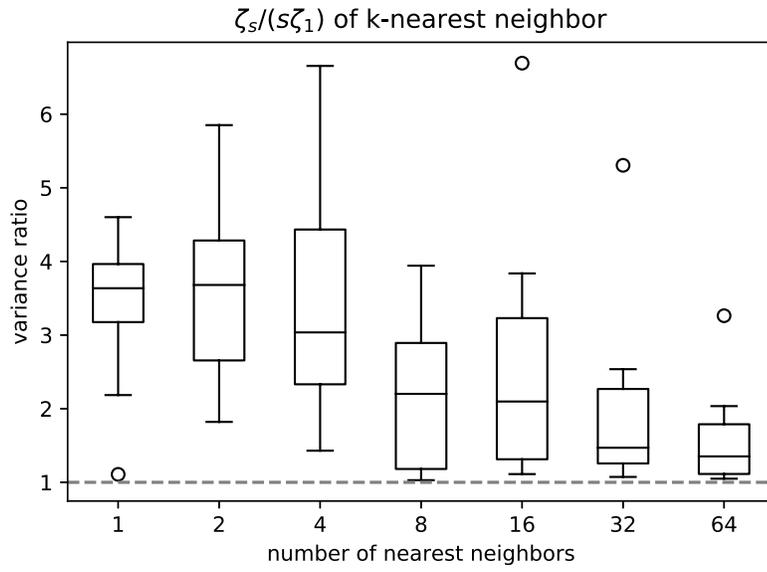


FIG 3. Variance ratio of  $k$ NN with differing numbers of nearest neighbors. We see that the variance ratio appears to get smaller as the number of nearest neighbors increases.

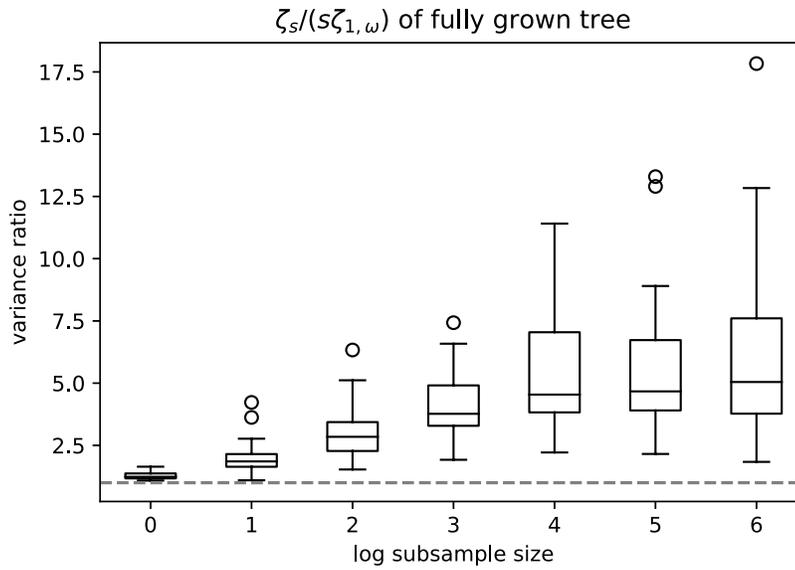


FIG 4. Variance ratio of fully grown trees built with subsamples of varying size. We can see that as the sample size increases, the variance ratio appears to slowly get larger.

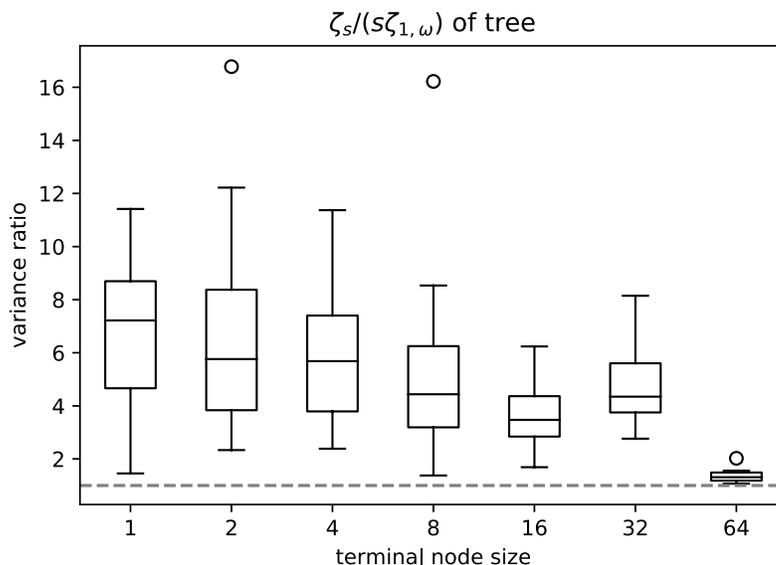


FIG 5. Variance ratio of trees with differing minimal terminal sizes. We see that the variance ratio get smaller as the minimal terminal size increases.

Fig. 4, utilizes fully grown trees with  $n_{\text{in}} = 1000$  and  $n_{\text{out}} = 50$  and again obtain boxplots of the estimated variance ratio  $\zeta_s/s\zeta_{1,\omega}$  for  $s = 2^i$ , for  $i = 1, \dots, 7$ . Once again we see that as  $s$  increases, the variance of our estimates becomes larger and the estimation less accurate. We also again see that the median of the estimated variance ratio  $\zeta_s/s\zeta_{1,\omega}$  increases as  $s$  increases at a rate roughly on the order of  $\log(s)$ .

In Fig. 5, we fix the subsample size  $s$  to be 100 and consider trees with the terminal sizes ranging between  $k$  and  $2k - 1$ , where  $k = 2^i$  for  $i = 1, \dots, 7$ . Compared to ensembles of  $k$ -NN base learners, the trees here have extra randomness through  $\omega$  and so  $\zeta_{1,\omega}$  will be relatively quite small. Thus, we let  $s$  be 100 here instead of 300 in an effort to make  $\zeta_{1,\omega}$  large enough for reasonably accurate estimation. Here we let  $n_{\text{in}} = 2000$  and  $n_{\text{out}} = 50$ . As  $k$  increases, we can see that  $\zeta_s/s\zeta_{1,\omega}$  decreases and gets closer to 1, as expected.

Finally, we turn to illustrating the asymptotic normality of generalized incomplete U-statistics, where the base learners are either nearest neighbor methods (kNN) or decision trees. Here we utilize the same prediction location and MARS function given in Eq. (B.29), letting  $n = 1000$ ,  $N = 2000$  and repeat the process 250 times to get an empirical distribution of the resulting predictions. The size of each subsample is denoted by  $s$  and  $k$  denotes the number of nearest neighbors for kNN or the minimum terminal node size for decision trees. In keeping with our definition of generalized U-statistics, to generate an ensemble, we first simulate  $\hat{N}$  and then randomly generate  $\hat{N}$  subsamples (without replacement) and build a base learner on each. Note that we want  $\hat{N} \sim \text{Binomial}(\binom{n}{s}, N/\binom{n}{s})$ , but

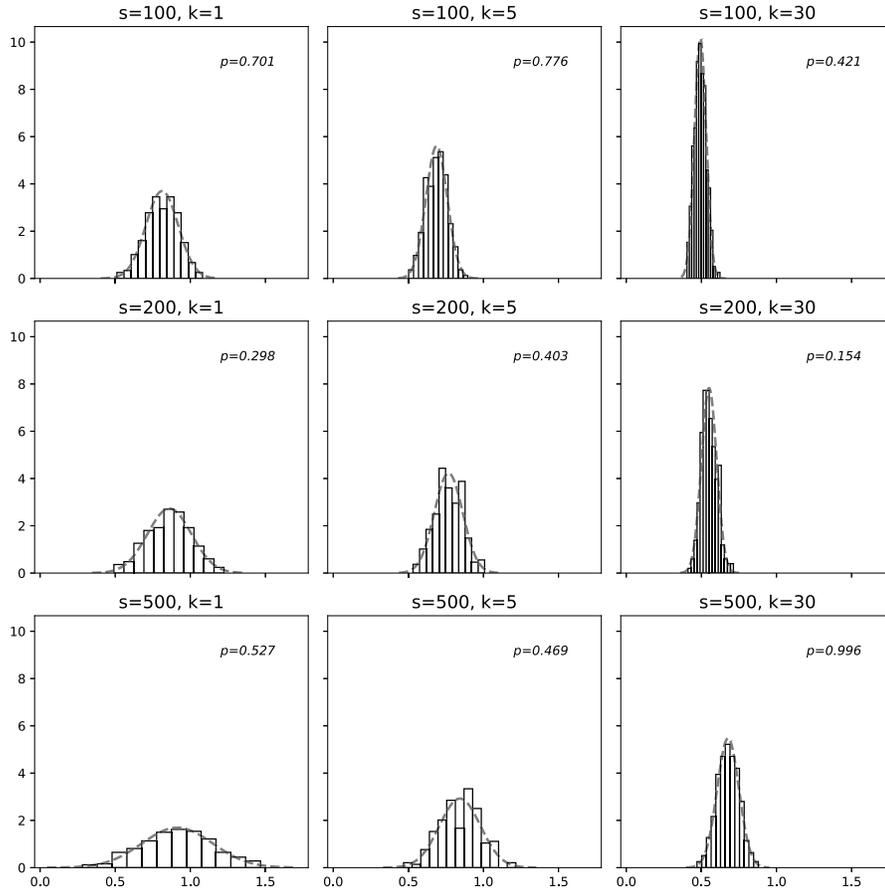


FIG 6. Histograms of predictions from  $k$ NN ensembles. Overlaying each histogram is the normal density estimate formed by the data.

this may be difficult to directly simulate since  $\binom{n}{s}$  could be very large. However, by the Berry-Esseen theorem, we have

$$\sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{\hat{N} - N}{\sqrt{\binom{n}{s} p(1-p)}} \leq z \right\} - \Phi(z) \right| \leq C \frac{1-2p}{\sqrt{\binom{n}{s} p(1-p)}}, \quad (\text{B.32})$$

where  $p = N/\binom{n}{s}$  and  $C < 0.5$ . In our setting  $p \ll 10^{-139}$ , so the right hand side of Eq. (B.32) is  $\leq C/\sqrt{N} \approx 0.011$  and  $\sqrt{\binom{n}{s} p(1-p)} \approx 1/\sqrt{N}$ . Thus, we can simulate  $\hat{N}$  by letting  $\hat{N} \sim \lfloor N + \sqrt{N} \cdot \mathcal{N}(0, 1) \rfloor$ .

Figures 6 and 7 show the distribution of predictions for  $k$ NN and tree ensembles, respectively, at different subsample sizes  $s$  and neighborhood (or terminal node) sizes  $k$ . Note that for both ensemble types, the distribution appears nor-

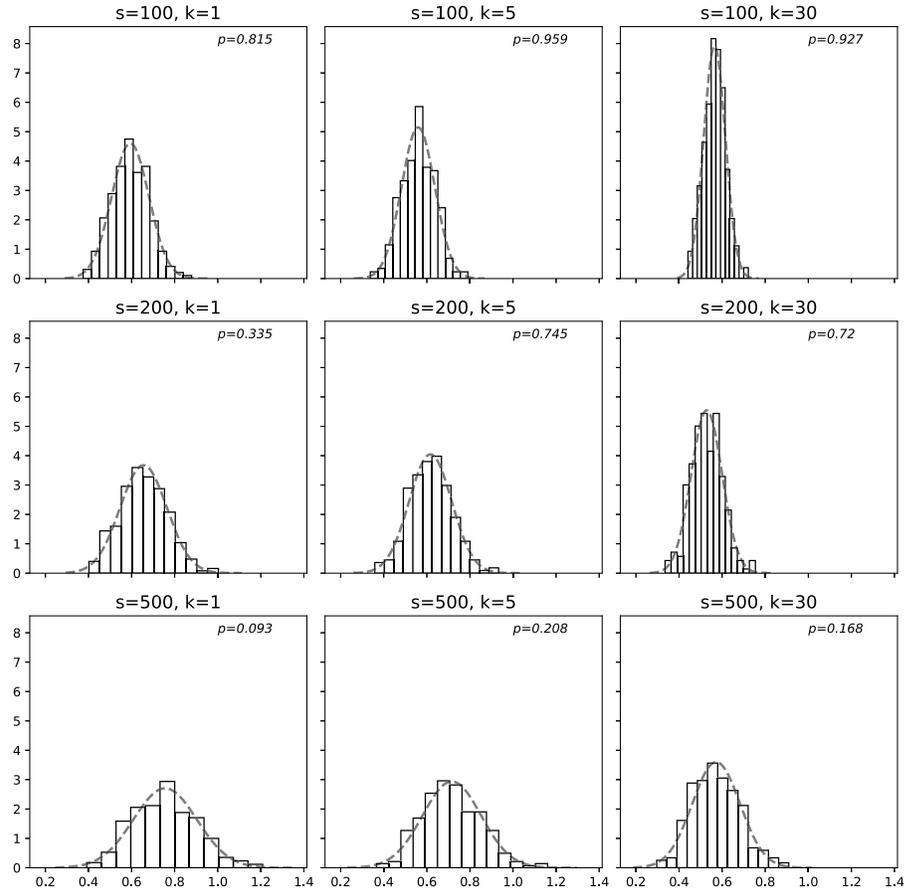


FIG 7. Histograms of predictions from tree-based ensembles. Overlaying each histogram is the normal density estimate formed by the data.

mal at every combination of  $s$  and  $k$ . This was verified with a Shapiro–Wilk test on each; p-values are shown in the upper-right of each plot. For both ensemble types, we see that the resulting variance of the (approximately normal) distribution seems to increase with  $s$  and, at each fixed  $s$ , decrease as  $k$  increases. Note that for large subsamples, the resulting distributions still remain approximately normal, though they do appear to be getting less normal with larger  $s$ , which is to be expected.

## Appendix C: Proofs in Section 4

### C.1. Introduction to Lemma 4

In many cases of interest, a statistic  $T$  can be written as a linear statistic plus a manageable term. [47] used the K-function approach derived from Stein’s

method [55] to build a random concentration inequality for linear statistics. This inequality is an extension of the usual concentration inequalities but the bounds can be random. The authors then apply this randomized concentration inequality to provide a Berry-Esseen bound for  $T$  as in Lemma 4. For completeness, we begin with a brief discussion of this inequality and its derivatives.

Let  $Z_1, \dots, Z_n$  be independent random variables and  $T$  be a statistic of the form

$$T = T(Z_1, \dots, Z_n) = W + \Delta$$

where

$$W = \sum_{i=1}^n g_{n,i}(Z_i), \quad \text{and} \quad \Delta = \Delta(Z_1, \dots, Z_n)$$

for some functions  $g_{n,i}$  and  $\Delta$ . Note here that  $W$  is linear and thus  $T$  takes the form of a linear statistic plus a remainder. Let  $\xi_i = g_{n,i}(Z_i)$  and assume that

$$\mathbb{E}[\xi_i] = 0 \quad (i = 1, \dots, n) \quad \text{and} \quad \sum_{i=1}^n \text{var}(\xi_i) = 1. \quad (\text{C.1})$$

The following randomized concentration inequality can be used to establish uniform Berry-Esseen bounds on  $T$  with optimal asymptotic rates.

**Lemma 3** ([47]). *Let  $\delta > 0$  satisfying*

$$\sum_{i=1}^n \mathbb{E} [|\xi_i| \min(\delta, |\xi_i|)] \geq 1/2. \quad (\text{C.2})$$

*Then for any real-valued random variables  $\Delta_1$  and  $\Delta_2$ ,*

$$\begin{aligned} \Pr(\Delta_1 \leq W \leq \Delta_2) &\leq 4\delta + \mathbb{E}|W(\Delta_2 - \Delta_1)| \\ &+ \sum_{i=1}^n [\mathbb{E}|\xi_i(\Delta_1 - \Delta_{1,i})| + \mathbb{E}|\xi_i(\Delta_2 - \Delta_{2,i})|] \end{aligned} \quad (\text{C.3})$$

*whenever  $\xi_i$  is independent of  $(W - \xi_i, \Delta_{1,i}, \Delta_{2,i})$ .*

For completeness, we replicate the proof of Eq. (C.3) originally given in [47]. The spirit of the proof is to replace bounding the probability by bounding the expectation of some functions.

*Proof.* Let

$$f_{a,b}(w) = \begin{cases} -\frac{1}{2}(b-a) - \delta, & w < a - \delta \\ w - \frac{1}{2}(a+b), & a - \delta \leq w \leq b + \delta \\ \frac{1}{2}(b-a) + \delta, & w > b + \delta \end{cases}$$

and let

$$\hat{K}_i(t) = \xi_i \{1_{-\xi_i \leq t \leq 0} - 1_{0 < t \leq -\xi_i}\}, \quad \hat{K}(t) = \sum_{i=1}^n \hat{K}_i(t).$$

Since  $\xi_i$  and  $f_{\Delta_1, \Delta_2, i}(W - \xi_i)$  are independent for  $1 \leq i \leq n$ , we have

$$\begin{aligned} \mathbb{E}[W f_{\Delta_1, \Delta_2}(W)] &= \sum_{i=1}^n \mathbb{E}[\xi_i (f_{\Delta_1, \Delta_2}(W) - f_{\Delta_1, \Delta_2}(W - \xi_i))] \\ &\quad + \sum_{i=1}^n \mathbb{E}[\xi_i (f_{\Delta_1, \Delta_2}(W - \xi_i) - f_{\Delta_1, \Delta_2, i}(W - \xi_i))] \\ &= H_1 + H_2 \end{aligned}$$

where

$$\begin{aligned} H_1 &= \sum_{i=1}^n \mathbb{E} \left[ \xi_i \int_{-\xi_i}^0 f'_{\Delta_1, \Delta_2}(W + t) dt \right] \\ &= \sum_{i=1}^n \mathbb{E} \left[ \int_{-\infty}^{\infty} f'_{\Delta_1, \Delta_2}(W + t) \hat{K}_i(t) dt \right] \\ &\geq \mathbb{E} \left[ \int_{|t| \leq \delta} f'_{\Delta_1, \Delta_2}(W + t) \hat{K}(t) dt \right] \\ &\geq \mathbb{E} \left[ 1_{\Delta_1 \leq W \leq \Delta_2} \int_{|t| \leq \delta} \hat{K}(t) dt \right] \\ &= \mathbb{E} \left[ 1_{\Delta_1 \leq W \leq \Delta_2} \sum_{i=1}^n |\xi_i| \min(\delta, \xi_i) \right] \\ &\geq H_{1,1} - H_{1,2} \end{aligned}$$

and where

$$H_{1,1} = \Pr(\Delta_1 \leq W \leq \Delta_2) \sum_{i=1}^n \mathbb{E}[|\xi_i| \min(\delta, \xi_i)] \geq \frac{1}{2} \Pr(\Delta_1 \leq W \leq \Delta_2)$$

and

$$\begin{aligned} H_{1,2} &= \mathbb{E} \left| \sum_{i=1}^n [|\xi_i| \min(\delta, \xi_i) - \mathbb{E}[|\xi_i| \min(\delta, \xi_i)]] \right| \\ &\leq \text{var} \left( \sum_{i=1}^n |\xi_i| \min(\delta, \xi_i) \right)^{1/2} \leq \delta. \end{aligned}$$

For  $H_2$ , we have

$$|f_{\Delta_1, \Delta_2}(w) - f_{\Delta_1, \Delta_2, i}(w)| \leq \frac{1}{2} |\Delta_1 - \Delta_{1,i}| + \frac{1}{2} |\Delta_2 - \Delta_{2,i}|$$

which then yields

$$|H_2| \leq \frac{1}{2} (\mathbb{E}|\xi_i(\Delta_1 - \Delta_{1,i})| + \mathbb{E}|\xi_i(\Delta_2 - \Delta_{2,i})|).$$

It follows from the definition of  $f_{a,b}$  that

$$|f_{\Delta_1, \Delta_2}(w)| \leq \frac{1}{2}(\Delta_2 - \Delta_1) + \delta.$$

Hence,

$$\begin{aligned} & \Pr(\Delta_1 \leq W \leq \Delta_2) \\ & \leq 2\mathbb{E}[Wf_{\Delta_1, \Delta_2}(W)] + 2\delta + \sum_{i=1}^n [\mathbb{E}|\xi_i(\Delta_1 - \Delta_{1,i})| + \mathbb{E}|\xi_i(\Delta_2 - \Delta_{2,i})|] \\ & \leq \mathbb{E}|W(\Delta_2 - \Delta_1)| + 2\delta\mathbb{E}|W| \\ & \quad + 2\delta + \sum_{i=1}^n [\mathbb{E}|\xi_i(\Delta_1 - \Delta_{1,i})| + \mathbb{E}|\xi_i(\Delta_2 - \Delta_{2,i})|] \\ & \leq \mathbb{E}|W(\Delta_2 - \Delta_1)| + 4\delta + \sum_{i=1}^n [\mathbb{E}|\xi_i(\Delta_1 - \Delta_{1,i})| + \mathbb{E}|\xi_i(\Delta_2 - \Delta_{2,i})|]. \quad \square \end{aligned}$$

Now, for any estimator of the form  $T = W + \Delta$ , we can write

$$-\Pr(z - |\Delta| \leq W \leq z) \leq \Pr(T \leq z) - \Pr(W \leq z) \leq \Pr(z \leq W \leq z + |\Delta|).$$

Applying Eq. (C.3) to these bounds, we arrive at the following lemma.

**Lemma 4** ([47]). *Let  $\xi_1, \dots, \xi_n$  be independent random variables satisfying Eq. (C.1),  $W = \sum_{i=1}^n \xi_i$  and  $T = W + \Delta$ . Let  $\Delta_i$  be a random variable such that  $\xi_i$  and  $(W - \xi_i, \Delta_i)$  are independent. Then for any  $\delta$  satisfying Eq. (C.2), we have*

$$\sup_{z \in \mathbb{R}} |\Pr(T \leq z) - \Pr(W \leq z)| \leq 4\delta + \mathbb{E}|W\Delta| + \sum_{i=1}^n \mathbb{E}|\xi_i(\Delta - \Delta_i)|.$$

In particular,

$$\sup_{z \in \mathbb{R}} |\Pr(T \leq z) - \Pr(W \leq z)| \leq 2(\beta_2 + \beta_3) + \mathbb{E}|W\Delta| + \sum_{i=1}^n \mathbb{E}|\xi_i(\Delta - \Delta_i)| \quad (\text{C.4})$$

and

$$\sup_{z \in \mathbb{R}} |\Pr(T \leq z) - \Phi(z)| \leq 6.1(\beta_2 + \beta_3) + \mathbb{E}|W\Delta| + \sum_{i=1}^n \mathbb{E}|\xi_i(\Delta - \Delta_i)| \quad (\text{C.5})$$

where

$$\beta_2 = \sum_{i=1}^n \mathbb{E}[|\xi_i^2| \mathbb{1}_{|\xi_i| > 1}] \quad \text{and} \quad \beta_3 = \sum_{i=1}^n \mathbb{E}[|\xi_i^3| \mathbb{1}_{|\xi_i| \leq 1}].$$

Note that since  $\sum_{i=1}^n \mathbb{E} [\xi_i^2] = 1$ , if  $\delta > 0$  satisfies

$$\sum_{i=1}^n \mathbb{E} [\xi_i^2 1_{|\xi_i| \geq \delta}] < \frac{1}{2}$$

then Eq. (C.2) holds. In particular, when the  $\xi_i$  are standardized i.i.d. random variables, then  $\delta$  must be on the order of  $1/\sqrt{n}$ . Furthermore, note that when  $\beta_2 + \beta_3 \leq 1$  and  $4\delta \leq 2(\beta_2 + \beta_3)$ , then Eq. (C.2) is automatically satisfied and thus Eq. (C.4) is immediate. Eq. (C.5) is obtained by combining Eq. (C.4) and the sharp Berry-Esseen bound of the sum of independent random variables in [56].

**C.2. Berry-Esseen bounds for generalized U-statistics**

**Proof of Theorem 3:** We provide the proof for  $U_{n,s}$ , the extension to  $U_{n,s,\omega}$  follows in the same fashion with the only difference being in the H-decomposition. Without loss of generality, let  $\theta = 0$ . Observe that

$$U_{n,s} = \sum_{j=1}^s \binom{s}{j} H_n^{(j)} = \frac{s}{n} \sum_{i=1}^n g(Z_i) + \sum_{j=2}^s \binom{s}{j} H_n^{(j)},$$

where  $g(z) = \mathbb{E}[h(z, Z_2, \dots, Z_n)]$  and  $H_n^{(j)} = \binom{n}{j}^{-1} \sum_{(n,j)} h^{(j)}(Z_{i1}, \dots, Z_{ij})$ . Let

$$\Delta = \sqrt{\frac{n}{s^2 \zeta_1}} \sum_{j=2}^s \binom{s}{j} H_n^{(j)}$$

and for  $l \in \{1, \dots, n\}$ , let

$$\Delta_l = \Delta - \sqrt{\frac{n}{s^2 \zeta_1}} \binom{n}{j}^{-1} \sum_{S_j^{(l)}} h^{(j)}(Z_{i1}, \dots, Z_{ij}) \tag{C.6}$$

where  $S_j^{(l)}$  denotes the collection of all subsets of variables of size  $j$  that include the  $l^{th}$  observation. The choice of  $\Delta_l$  plays key role in deciding Berry-Esseen bound, since the closer  $\Delta_l$  is to  $\Delta$ , the tighter the bound in Eq. (C.5). We have

$$\sqrt{\frac{n}{s^2 \zeta_1}} U_{n,s} = W + \Delta, \tag{C.7}$$

where  $W = \sum_{i=1}^n \xi_i$  with  $\xi_i = g(Z_i)/\sqrt{n\zeta_1}$ . For each  $i \in \{1, \dots, n\}$ , the random variable  $W - \xi_i$  and  $\Delta_i$  are functions of  $Z_j, j \neq i$ . Therefore  $\xi_i$  is independent of  $(W - \xi_i, \Delta_i)$ . By Cauchy-Schwarz inequality, we have

$$\mathbb{E}[|W\Delta|] \leq \sqrt{\mathbb{E}|W|^2} \cdot \sqrt{\mathbb{E}|\Delta|^2} = \sqrt{\mathbb{E}|\Delta|^2}$$

and

$$\begin{aligned} \sum_{i=1}^n \mathbb{E}[|\xi_i(\Delta - \Delta_i)|] &\leq \left( \sum_{i=1}^n \mathbb{E}[\xi_i^2] \right)^{\frac{1}{2}} \cdot \left( \sum_{i=1}^n \mathbb{E}|\Delta - \Delta_i|^2 \right)^{\frac{1}{2}} \\ &\leq \sqrt{n} \max(\sqrt{\mathbb{E}|\Delta - \Delta_i|^2}). \end{aligned}$$

Observing the terms on the right, we have

$$\begin{aligned} \frac{s^2 \zeta_1}{n} \mathbb{E}|\Delta|^2 &= \text{var} \left\{ \sum_{j=2}^s \binom{s}{j} H_n^{(j)} \right\} \\ &= \sum_{j=2}^s \binom{s}{j}^2 \binom{n}{j}^{-1} V_j \\ &\leq \frac{s^2}{n^2} (\zeta_s - s\zeta_1). \end{aligned}$$

Similarly, we have

$$\begin{aligned} \frac{s^2 \zeta_1}{n} \mathbb{E}|\Delta - \Delta_i|^2 &= \text{var} \left\{ \sum_{j=2}^s \binom{s}{j} \binom{n}{j}^{-1} \sum_{S_j^{(i)}} h^{(j)}(Z_{i1}, \dots, Z_{ij}) \right\} \\ &= \sum_{j=2}^s \binom{s}{j}^2 \binom{n}{j}^{-2} \binom{n-1}{j-1} V_j \\ &= \sum_{j=2}^s \binom{s}{j}^2 \binom{n}{j}^{-1} \frac{j}{n} V_j \\ &\leq \frac{2s^2}{n^3} \sum_{j=2}^s \binom{s}{j} V_j \\ &\leq \frac{2s^2}{n^3} (\zeta_s - s\zeta_1). \end{aligned}$$

Note that

$$\begin{aligned} \beta_2 + \beta_3 &= \sum_{i=1}^n \mathbb{E} \left[ \left| \frac{g(Z_i)}{\sqrt{n\zeta_1}} \right|^2 \mathbf{1}_{|g(Z_i)| \geq \sqrt{n\zeta_1}} \right] + \sum_{i=1}^n \mathbb{E} \left[ \left| \frac{g(Z_i)}{\sqrt{n\zeta_1}} \right|^3 \mathbf{1}_{|g(Z_i)| \leq \sqrt{n\zeta_1}} \right] \\ &\leq \frac{1}{n^{1/2}} \frac{\mathbb{E}|g|^3}{\zeta_1^{3/2}}. \end{aligned}$$

Finally, by applying Lemma 4, we obtain

$$\sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s}}{\sqrt{s^2 \zeta_1 / n}} \leq z \right\} - \Phi(z) \right| \leq \frac{6.1 \mathbb{E}|g|^3}{n^{1/2} \zeta_1^{3/2}} + (1 + \sqrt{2}) \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s\zeta_1} - 1 \right) \right\}^{1/2}. \quad \blacksquare$$

**Proof of Theorem 4:** We provide a bound for incomplete, infinite-order U-statistics. An analogous result for *generalized* incomplete U-statistics  $U_{n,s,N,\omega}$  can be established by applying the extended form of the H-decomposition. As eluded to earlier, an incomplete U-statistic can be written as

$$U_{n,s,N} = \frac{1}{\hat{N}} \sum_{(n,s)} \rho h(Z_{i_1}, \dots, Z_{i_s}) \quad (\text{C.8})$$

where  $\rho \sim \text{Bernoulli}(p)$  and  $p = N/\binom{n}{s}$ . Note however that Eq. (C.8) can also be written as

$$U_{n,s,N} = \frac{N}{\hat{N}} \left\{ \binom{n}{s}^{-1} \sum_{(n,s)} \frac{\rho}{p} h(Z_{i_1}, \dots, Z_{i_s}) \right\} = \frac{N}{\hat{N}} U_{n,s}^*$$

so that the incomplete U-statistic now takes the form of a scaled, *generalized* complete U-statistic. We thus now consider  $U_{n,s}^*$  and can then extend the results to  $U_{n,s,N}$ . First, note that the variance terms  $\zeta_c^*$  for  $c = 1, \dots, s$  of  $U_{n,s}^*$  are different from those of  $U_{n,s}$  in Eq. (2.1). For  $c = 1, \dots, s-1$ , we have

$$\zeta_c^* = \text{cov}\left(\frac{\rho}{p} h(Z_1, \dots, Z_c, Z_{c+1}, \dots, Z_s), \frac{\rho'}{p} h(Z_1, \dots, Z_c, Z'_{c+1}, \dots, Z'_n)\right) = \zeta_c$$

and

$$\zeta_s^* = \text{cov}\left(\frac{\rho}{p} h(Z_1, \dots, Z_s), \frac{\rho}{p} h(Z_1, \dots, Z_s)\right) = \frac{1}{p} \zeta_s.$$

The H-decomposition will also be different. Here, we have

$$\begin{aligned} h^{(1)*} &= \mathbb{E} \left[ \frac{\rho}{p} h \mid Z_1 \right] = h^{(1)} \\ h^{(2)*} &= \mathbb{E} \left[ \frac{\rho}{p} h \mid Z_1, Z_2 \right] - \mathbb{E} \left[ \frac{\rho}{p} h \mid Z_1 \right] - \mathbb{E} \left[ \frac{\rho}{p} h \mid Z_2 \right] = h^{(2)} \\ &\vdots \\ h^{(s)*} &= \frac{\rho}{p} h - \sum_{j=1}^{s-1} \sum_{(s,j)} h^{(j)}(Z_{i_1}, \dots, Z_{i_j}) \end{aligned}$$

where the  $h$  appearing in the earlier form is replaced here by  $\frac{\rho}{p} h$ . These kernels still retain the desirable properties laid out in Proposition 5. Furthermore, we have

$$V_j^* = V_j \quad \text{for } j = 1, \dots, s-1$$

and

$$V_s^* = \frac{1}{p} \sum_{j=1}^s \binom{s}{j} V_j - \sum_{j=1}^{s-1} \binom{s}{j} V_j = V_s + \frac{(1-p)}{p} \zeta_s.$$

Thus,

$$\begin{aligned} U_{n,s}^* &= \sum_{j=1}^{s-1} \binom{s}{j} H_n^{(j)*} + H_n^{(s)*} \\ &= \sum_{j=1}^{s-1} \binom{s}{j} H_n^{(j)} + \binom{n}{s}^{-1} \sum_{(n,s)} h^{(s)*}(Z_{i1}, \dots, Z_{is}) \\ &= sH_n^{(1)} + \Delta. \end{aligned}$$

Now, because we have rewritten  $U_{n,s}^*$  as a linear term plus a remainder, we can follow the same general strategy as in the complete case above in applying Lemma 4. In particular, let

$$\Delta - \Delta_i = \sum_{j=2}^{s-1} \binom{s}{j} \binom{n}{j}^{-1} \sum_{S_j^{(i)}} h^{(j)}(Z_{i1}, \dots, Z_{ij}) + \binom{n}{s}^{-1} \sum_{S_s^{(i)}} h^{(s)*}(Z_{i1}, \dots, Z_{is}),$$

where  $S_j^{(i)}$  denotes the collection of all subsets of size  $j$  that include the  $i^{\text{th}}$  observation. Then

$$\begin{aligned} \mathbb{E}|\Delta|^2 &= \sum_{j=2}^{s-1} \binom{s}{j}^2 \binom{n}{j}^{-1} V_j + \binom{n}{s}^{-1} V_s^* \\ &= \sum_{j=2}^s \binom{s}{j}^2 \binom{n}{j}^{-1} V_j + \frac{1}{N}(1-p)\zeta_s, \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}|\Delta - \Delta_i|^2 &= \sum_{j=2}^s \binom{s}{j}^2 \binom{n}{j}^{-1} \frac{j}{n} V_j + \binom{n}{s}^{-2} \binom{n-1}{s-1} V_s^* \\ &= \sum_{j=2}^s \binom{s}{j}^2 \binom{n}{j}^{-1} \frac{j}{n} V_j + \frac{s}{n} \frac{1}{N}(1-p)\zeta_s. \end{aligned}$$

Thus

$$\frac{n}{s^2 \zeta_1} \mathbb{E}|\Delta|^2 \leq \frac{s}{n} \left[ \frac{\zeta_s}{s \zeta_1} - 1 \right] + \frac{n}{Ns} (1-p) \frac{\zeta_s}{s \zeta_1}$$

and

$$\sum_{i=1}^n \frac{n}{s^2 \zeta_1} \mathbb{E}|\Delta_i|^2 \leq \frac{2s}{n} \left[ \frac{\zeta_s}{s \zeta_1} - 1 \right] + \frac{n}{N} (1-p) \frac{\zeta_s}{s \zeta_1}.$$

By applying Lemma 4, we have

$$\begin{aligned}
& \sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s}^* - \theta}{\sqrt{s^2 \zeta_1/n}} \leq z \right\} - \Phi(z) \right| \\
& \leq \frac{6.1 \mathbb{E}|g|^3}{n^{1/2} \zeta_1^{3/2}} + (1 + \sqrt{2}) \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s \zeta_1} - 1 \right) \right\}^{1/2} + \\
& \quad \left( 1 + \sqrt{\frac{1}{s}} \right) \left\{ \frac{n}{N} (1-p) \frac{\zeta_s}{s \zeta_1} \right\}^{1/2} \\
& := \epsilon_0.
\end{aligned}$$

Next, we give the Berry-Esseen bound for  $U_{n,s,N}$ .

Denote  $\frac{U_{n,s,N} - \theta}{\sqrt{s^2 \zeta_1/n}} = T = W + \Delta$ , where

$W = \frac{U_{n,s}^* - \theta}{\sqrt{s^2 \zeta_1/n}}$  and  $\Delta = \left( \frac{N}{n} - 1 \right) W$ . Note that we always have

$$-\Pr(z - |\Delta| \leq W \leq z) \leq \Pr(T \leq z) - \Pr(W \leq z) \leq \Pr(z \leq W \leq z + |\Delta|).$$

Since the Berry-Esseen bound for  $W$  is already obtained, it remains only to bound  $\Pr(z \leq W \leq z + |\Delta|)$  and  $\Pr(z - |\Delta| \leq W \leq z)$ . Without loss of generality, we consider bounding  $\Pr(z \leq W \leq z + |\Delta|)$ . By Bernstein's inequality, for any  $\epsilon > 0$ ,

$$\Pr \left( \left| \frac{\hat{N}}{N} - 1 \right| \geq \epsilon \right) \leq 2 \exp \left( - \frac{-\epsilon^2 N}{(1-p) + \epsilon/3} \right). \quad (\text{C.9})$$

By letting  $\epsilon = 2N^{-\beta} < 1/2$  and noting that  $\Pr(|Z| \geq N^\alpha) \leq 2 \exp(-N^{2\alpha}/2)$ , we have

$$\begin{aligned}
\Pr(|\Delta| \geq N^{-2\beta+\alpha}) & \leq \Pr \left( \left| \frac{N}{\hat{N}} - 1 \right| \geq N^{-2\beta} \right) + \Pr(|W| \geq N^\alpha) \\
& \leq \Pr \left( \left| \frac{\hat{N}}{N} - 1 \right| \geq 2N^{-2\beta} \right) + \Pr(|W| \geq N^\alpha) \\
& \leq \Pr \left( \left| \frac{N}{\hat{N}} - 1 \right| \geq 2N^{-2\beta} \right) + \Pr(|Z| \geq N^\alpha) + 2\epsilon_0 \\
& \leq 2 \exp(-N^{2\alpha}/2) + 2 \exp \left( - \frac{4N^{1-2\beta}}{(1-p)(1+2N^{-\beta}/3)} \right) + 2\epsilon_0 \\
& := \epsilon_1 + 2\epsilon_0.
\end{aligned}$$

As a result, we have

$$\begin{aligned}
& \Pr(z \leq W \leq z + |\Delta|) \\
& \leq \Pr(z \leq W \leq z + |\Delta|, |\Delta| \leq N^{-2\beta+\alpha}) + \Pr(|\Delta| \geq N^{-2\beta+\alpha}) \\
& \leq \Pr(z \leq W \leq z + N^{-2\beta+\alpha}) + \epsilon_1 + 2\epsilon_0 \\
& \leq 2\epsilon_0 + \Pr(z \leq Z \leq z + N^{-2\beta+\alpha}) + \epsilon_1 + 2\epsilon_0 \\
& \leq 4\epsilon_0 + \epsilon_1 + \frac{1}{\sqrt{2\pi}} N^{-2\beta+\alpha} \\
& := 4\epsilon_0 + \epsilon_1 + \epsilon_2.
\end{aligned}$$

Let  $\beta = 0.5 - 0.5\eta_0$  and  $\alpha = 0.5\eta_0$ , where  $0 < \eta_0 < 0.5$ . We know that  $\epsilon_1 \ll \epsilon_2$  when  $N$  is large, and therefore

$$\begin{aligned}
& \sup_{z \in \mathbb{R}} \left| \Pr \left\{ \frac{U_{n,s,N} - \theta}{\sqrt{s^2 \zeta_1 / n}} \leq z \right\} - \Phi(z) \right| \\
& = \sup_{z \in \mathbb{R}} |\Pr(T \leq z) - \Phi(z)| \\
& \leq \sup_{z \in \mathbb{R}} |\Pr(W \leq z) - \Phi(z)| + \sup_{z \in \mathbb{R}} |\Pr(T \leq z) - \Pr(W \leq z)| \\
& \leq 5\epsilon_0 + \epsilon_1 + \epsilon_2 \\
& = O \left( \frac{\mathbb{E}[g^2]}{n^{1/2} \zeta_1^{3/2}} + \left[ \frac{s}{n} \left( \frac{\zeta_s}{s \zeta_1} - 1 \right) \right]^{1/2} + \left\{ \frac{n}{N} (1-p) \frac{\zeta_s}{s \zeta_1} \right\}^{1/2} + N^{-1/2+\eta_0} \right).
\end{aligned}$$

Here,  $\eta_0$  can be any constant such that  $0 < \eta_0 < 0.5$ . Note that the factor of  $N/\tilde{N}$  only introduces the extra term  $-N^{-1/2+\eta_0}$  in the ultimate Berry-Esseen bound of  $U_{n,s,N}$  compared to that of  $U_{n,s}^*$ . ■

We now turn to the proof of Theorem 5. Before that, we give the statement and proof of the following simple lemma, which will be used in our proof.

**Lemma 5.**

$$\lim_{a \rightarrow 1^+} \sup_{z \in \mathbb{R}} \left| \frac{\Phi(az) - \Phi(z)}{a-1} \right| < \infty. \quad (\text{C.10})$$

*Proof.* Let  $f(z) = \Phi(az) - \Phi(z)$ , then

$$f'(z) = \frac{1}{\sqrt{2\pi}} \left( a e^{-a^2 z^2 / 2} - e^{-z^2 / 2} \right).$$

Solving  $f'(z) = 0$ , we get  $z^2(a) = \frac{\log(a)}{(a^2-1)/2}$ . Then

$$\begin{aligned}
\lim_{a \rightarrow 1^+} \sup_{z \in \mathbb{R}} \left| \frac{\Phi(az) - \Phi(z)}{a-1} \right| &= \lim_{a \rightarrow 1^+} \left| \frac{\Phi(az(a)) - \Phi(z(a))}{a-1} \right| \\
&= \lim_{a \rightarrow 1^+} |\Phi'(az(a))(az(a))' - \Phi'(z(a))z'(a)|.
\end{aligned} \quad (\text{C.11})$$

According to the Taylor expansion of  $\log(a)$  at  $a = 1$ , we have

$$z^2(a) = \frac{2}{a+1} \left( 1 - \frac{a-1}{2} + \dots \right) = 1 + o(a-1).$$

Hence  $\lim_{a \rightarrow 1^+} z^2(a) = 1$  and  $\lim_{a \rightarrow 1^+} z'(a) = -\frac{1}{2}$ , and therefore we have

$$\lim_{a \rightarrow 1^+} \sup_{z \in \mathbb{R}} \left| \frac{\Phi(az) - \Phi(z)}{a-1} \right| = \frac{e^{-1/2}}{\sqrt{2\pi}} < \infty \quad (\text{C.12})$$

as desired.  $\square$

**Proof of Theorem 5:** We begin with a bound for incomplete, infinite-order U-statistics. The extension of this result to the generalized setting and can be derived in the same fashion. First rewrite Eq. (C.8) as

$$U_{n,s,N} = \frac{1}{\hat{N}} \sum_i \rho_i h(\mathbf{Z}_i) \quad (\text{C.13})$$

where  $\rho_i \sim \text{Bernoulli}(N/\binom{n}{s})$  are i.i.d. and  $\mathbf{Z}_i = (Z_{i1}, \dots, Z_{is})$  denotes a subsample with index  $i$  and the sum is taken over all subsamples. We can rewrite  $U_{n,s,N}$  in Eq. (C.13) as a complete U-statistic  $U_{n,s}$  plus some manageable term up to the scalar  $N/\hat{N}$ . We have

$$\begin{aligned} U_{n,s,N} &= \frac{\hat{N}}{N} \left[ \binom{n}{s}^{-1} \sum_i h(\mathbf{Z}_i) + \frac{1}{N} \sum_i (\rho_i - p) h(\mathbf{Z}_i) \right] \\ &= \frac{\hat{N}}{N} \left[ U_{n,s} + \left( \sqrt{1-p} \right) \frac{1}{N} \sum_i \frac{\rho_i - p}{\sqrt{1-p}} h(\mathbf{Z}_i) \right] \\ &= \frac{\hat{N}}{N} \left[ A_n + \left( \sqrt{1-p} \right) B_n \right] \\ &= \frac{\hat{N}}{N} W_n. \end{aligned} \quad (\text{C.14})$$

We first consider the Berry-Esseen bound of  $W_n$ . Since we already know the limiting behavior of  $A_n$ , it remains only to control  $B_n$ . Note that

$$\Pr(\sqrt{n}W_n \leq z) = \Pr \left\{ \sqrt{N}B_n \leq \frac{z}{\sqrt{\alpha_n(1-p)}} - \sqrt{\frac{N}{1-p}}A_n \right\}$$

where  $\alpha_n = n/N$ . Conditioning on  $Z_1, \dots, Z_n$ ,  $A_n$  can be treated as a constant and we have

$$\sqrt{N}B_n \mid Z_1, \dots, Z_n = \binom{n}{s}^{-\frac{1}{2}} \sum_i \left[ \frac{(\rho_i - p)}{\sqrt{p(1-p)}} \right] h(\mathbf{Z}_i) \mid Z_1, \dots, Z_n. \quad (\text{C.15})$$

Now,  $\sqrt{N}B_n \mid Z_1, \dots, Z_n$  is a sum of independent random variables with variance  $U_2$ , where

$$U_2 = \binom{n}{s}^{-1} \sum_i h^2(\mathbf{Z}_i). \tag{C.16}$$

Let

$$\xi_i = \frac{(p(1-p))^{-1/2}(\rho_i - p)h(\mathbf{Z}_i)}{\sqrt{\sum_i h^2(\mathbf{Z}_i)}}, \quad a_i = \frac{h(\mathbf{Z}_i)}{\sqrt{\sum_i h^2(\mathbf{Z}_i)}}$$

then

$$\sum_i a_i^2 = 1 \quad \text{and} \quad \xi_i = a_i \left[ \frac{(\rho_i - p)}{\sqrt{p(1-p)}} \right].$$

Applying the Berry-Esseen bound in [57] for independent random variables, we have

$$\sup_{z \in \mathbb{R}} \left| \Pr \left( \sqrt{N}B_n \leq z \mid Z_1, \dots, Z_n \right) - \Phi \left( z/\sqrt{U_2} \right) \right| \leq 4.1(\beta_2 + \beta_3), \tag{C.17}$$

where  $\beta_2 = \sum_i \mathbb{E} [|\xi_i|^2 1_{|\xi_i| \geq 1}]$  and  $\beta_3 = \sum_i \mathbb{E} [|\xi_i|^3 1_{|\xi_i| \leq 1}]$ . Next, we show that  $(\beta_2 + \beta_3)$  can be uniformly bounded by a small number with high probability and in the rare case when  $(\beta_2 + \beta_3)$  is large, trivially, we have

$$\left| \Pr \left( \sqrt{N}B_n \leq z \mid Z_1, \dots, Z_n \right) - \Phi \left( z/\sqrt{U_2} \right) \right| \leq 2.$$

Indeed,

$$\begin{aligned} \beta_2 + \beta_3 &\leq \sum_i \mathbb{E} |\xi_i|^3 \\ &= \left\{ \frac{\binom{n}{s}^{-1} \sum_i |h(\mathbf{Z}_i)|^3}{\left( \binom{n}{s}^{-1} \sum_i |h(\mathbf{Z}_i)|^2 \right)^{3/2}} \right\} \cdot \binom{n}{s}^{-\frac{1}{2}} \left[ \frac{2p^2 - 2p + 1}{(p(1-p))^{1/2}} \right] \\ &= \frac{U_3}{U_2^{3/2}} \binom{n}{s}^{-1} \left[ \frac{1 - 2p}{(p(1-p))^{1/2}} \right], \end{aligned} \tag{C.18}$$

where  $U_3 = \binom{n}{s}^{-1} \sum_i |h(\mathbf{Z}_i)|^3$ . The terms of  $U_2$  and  $U_3$  are both complete U-statistics and as such, should be concentrated around their expectations. Let

$$\kappa_1 = \frac{\mathbb{E}|h|^4}{(\mathbb{E}|h|^2)^2}, \quad \kappa_2 = \frac{\mathbb{E}|h|^6}{(\mathbb{E}|h|^3)^2}$$

and recall that  $\kappa_1, \kappa_2$  are uniformly bounded by our assumption. Let  $\nu_2 = \mathbb{E}|h|^2 (= \zeta_s)$ ,  $\delta_2 = \left(\frac{s}{n}\right)^\eta \nu_2$ , where  $\eta > 0$ . Then by Chebyshev's inequality, we have

$$\Pr (|U_2 - \nu_2| \geq \delta_2) \leq \frac{\frac{s}{n} \text{var} (|h|^2)}{\delta_2^2} = \left(\frac{s}{n}\right)^{1-2\eta} (\kappa_1 - 1).$$

A similar inequality holds for  $|U_3 - \nu_3|$  and therefore with probability of at least  $1 - \epsilon_0$ , where  $\epsilon_0 = c_0(\frac{s}{n})^{1-2\eta}$  for some constant  $c_0 > 0$ , we have

$$\left| \frac{U_3}{U_2} \right| = \left| \frac{\frac{U_3}{\nu_3}}{\frac{U_2^{3/2}}{\nu_2^{3/2}}} \cdot \frac{\nu_3}{\nu_2} \right| \leq \left\{ \frac{\frac{\nu_3 + \delta_3}{\nu_3}}{\frac{(\nu_2 - \delta_2)^{3/2}}{\nu_2^{3/2}}} \right\} \frac{\nu_3}{\nu_2} \leq c_1 \frac{\nu_3}{\nu_2^{3/2}},$$

where  $c_1 = \left\{ \frac{1 + (\frac{s}{n})^\eta}{(1 - (\frac{s}{n})^\eta)^{3/2}} \right\}$ . Hence, combining this with Eq. (C.18), with probability of at least  $1 - \epsilon_0$ , we have

$$\begin{aligned} \beta_2 + \beta_3 &\leq c_1 \frac{\nu_3}{\nu_2^{3/2}} \binom{n}{s}^{-\frac{1}{2}} \left\{ \frac{1 - 2p + 2p^2}{(p(1-p))^{1/2}} \right\} \\ &\leq c_1 \frac{\nu_3}{\nu_2^{3/2}} N^{-\frac{1}{2}} \left\{ \frac{1 - 2p + 2p^2}{(1-p)^{1/2}} \right\} \\ &\leq c_1 c_2 \frac{\nu_3}{\nu_2^{3/2}} N^{-\frac{1}{2}}, \end{aligned}$$

where  $c_2 = \frac{1 - 2p + 2p^2}{(1-p)^{1/2}}$ . The next step is to substitute  $U_2$  by  $\zeta_s$  by applying Lemma 5 stated above. We obtain

$$\begin{aligned} \sup_{z \in \mathbb{R}} \left| \Phi\left(z/\sqrt{U_2}\right) - \Phi\left(z/\sqrt{\zeta_s}\right) \right| &\leq c_3 |\sqrt{\zeta_s} \wedge \sqrt{U_2}|^{-1} |\sqrt{U_2} - \sqrt{\zeta_s}| \\ &\leq c_3 |\zeta_s \wedge U_2|^{-1} |U_2 - \zeta_s|. \end{aligned}$$

Under the event that  $|U_2 - \zeta_s| \leq \delta_2$ , we have

$$\sup_{z \in \mathbb{R}} \left| \Phi\left(z/\sqrt{U_2}\right) - \Phi\left(z/\sqrt{\zeta_s}\right) \right| \leq c_3 \frac{\delta_2}{\zeta_s - \delta_2} = c_3 \frac{(\frac{s}{n})^\eta}{1 - (\frac{s}{n})^\eta}.$$

Next, since  $A_n$  is a complete U-statistic, by Theorem 3, we have

$$\sup_{z \in \mathbb{R}} |\Pr(\sqrt{n}A_n \leq z) - \Pr(Y_A \leq z)| \leq \epsilon_2 \tag{C.19}$$

where  $\epsilon_2 = \frac{6.1\mathbb{E}|g|^3}{n^{1/2}\zeta_1^{3/2}} + (1 + \sqrt{2}) \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s\zeta_1} - 1 \right) \right\}^{1/2}$  and  $Y_A \sim N(0, s^2\zeta_1)$ . Lastly,

$$\begin{aligned} \Pr(\sqrt{n}W_n \leq z) &= \mathbb{E} \left[ \Pr \left\{ \sqrt{N}B_n \leq \frac{z}{\sqrt{\alpha_n(1-p)}} - \sqrt{\frac{N}{1-p}}A_n \mid Z_1, \dots, Z_n \right\} \right] \\ &\leq \Pr \left\{ Y_B \leq \frac{z}{\sqrt{\alpha_n(1-p)}} - \sqrt{\frac{N}{1-p}}A_n \right\} + \epsilon_1 \\ &= \Pr \left\{ \sqrt{n}A_n \leq z - \sqrt{\alpha_n(1-p)}Y_B \right\} + \epsilon_1 \end{aligned}$$

where  $Y_B \sim N(0, \zeta_s)$  is independent of  $Z_1, \dots, Z_n$  and

$$\epsilon_1 = 4.1 \left\{ c_1 c_2 \frac{\nu_3}{\nu_2^{3/2}} N^{-1/2} \right\} + c_3 \frac{\left(\frac{s}{n}\right)^\eta}{1 - \left(\frac{s}{n}\right)^\eta} + 2\epsilon_0. \text{ Now, conditioning on } Y_B, \text{ we have}$$

$$\begin{aligned} & \Pr \left\{ \sqrt{n} A_n \leq z - \sqrt{\alpha_n(1-p)} Y_B \mid Y_B \right\} \\ & \leq \Pr \left\{ Y_A \leq z - \sqrt{\alpha_n(1-p)} Y_B \mid Y_B \right\} + \epsilon_2. \end{aligned}$$

Combining Eq. (C.17) and Eq. (C.19), we conclude that

$$\begin{aligned} \Pr \left\{ \sqrt{n} W_n \leq z \right\} & \leq \Pr \left\{ Y_A \leq z - \sqrt{\alpha_n(1-p)} Y_B \right\} + \epsilon_1 + \epsilon_2 \\ & = \Pr \left\{ Y_A + \sqrt{\alpha_n(1-p)} Y_B \leq z \right\} + \epsilon_1 + \epsilon_2 \\ & \leq \Pr \left\{ Y_A + \alpha_n^{1/2} Y_B \leq z \right\} + \epsilon_1 + \epsilon_2 + \epsilon_3. \end{aligned}$$

By Lemma 5, we have

$$\begin{aligned} \epsilon_3 & \leq c_3 \left( s^2 \zeta_1 + \alpha_n(1-p)\zeta_s \right)^{-1} \alpha_n p \zeta_s \\ & = c_3 \left( s^2 \zeta_1 + \alpha_n(1-p)\zeta_s \right)^{-1} \binom{n}{s}^{-1} n \zeta_s \\ & \leq c_3 \min \left\{ p(1-p)^{-1}, \frac{n/s \zeta_s}{\binom{n}{s} s \zeta_1} \right\}. \end{aligned}$$

Thus, in summary,

$$\sup_{z \in \mathbb{R}} \left| \Pr \left\{ \sqrt{N} W_n \leq z \right\} - \Pr \left\{ Y_W \leq z \right\} \right| \leq \epsilon_1 + \epsilon_2 + \epsilon_3$$

where

$$\begin{aligned} \epsilon_1 & = 2c_0 \left(\frac{s}{n}\right)^{1-2\eta} + 4.1 \left\{ c_1 c_2 \frac{\nu_3}{\nu_2^{3/2}} N^{-1/2} \right\} + c_3 \frac{\left(\frac{s}{n}\right)^\eta}{1 - \left(\frac{s}{n}\right)^\eta} \\ \epsilon_2 & = \frac{6.1 \mathbb{E}|g|^3}{n^{1/2} \zeta_1^{3/2}} + (1 + \sqrt{2}) \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s \zeta_1} - 1 \right) \right\}^{1/2} \\ \epsilon_3 & = c_3 \cdot \min \left\{ p(1-p)^{-1}, \frac{n/s \zeta_s}{\binom{n}{s} s \zeta_1} \right\} \end{aligned}$$

and  $Y_W \sim N(0, s^2 \zeta_1/n + \zeta_s/N)$ . Note that  $\epsilon_1$  and  $\epsilon_2$  dominate because of the  $\binom{n}{s}$  in the denominator of  $\epsilon_3$ . Therefore, by choosing  $\eta = 1/3$ , the above bound can be simplified as

$$\begin{aligned} \epsilon_1 + \epsilon_2 + \epsilon_3 & \leq C \left\{ \frac{\mathbb{E}|g|^3}{n^{1/2} (\mathbb{E}|g|^2)^{3/2}} + \frac{\mathbb{E}|h|^3}{N^{1/2} (\mathbb{E}|h|^2)^{3/2}} \right. \\ & \quad \left. + \left\{ \frac{s}{n} \left( \frac{\zeta_s}{s \zeta_1} - 1 \right) \right\}^{1/2} + \left(\frac{s}{n}\right)^{1/3} \right\}. \end{aligned}$$

As in the Proof of Theorem 4, the scalar  $\hat{N}/N$  introduces the extra term  $N^{-1/2+\eta_0}$  in the ultimate Berry-Essen bound of  $U_{n,s,N}$ , where  $0 < \eta_0 < 0.5$ . ■

### C.3. Discussion on a tighter bound

Here we provide a sketch of the proof of Theorem 6. Let  $m = \lfloor n/s \rfloor$  and define

$$V(Z_1, Z_2, \dots, Z_n) = \frac{1}{m} \sum_{j=0}^{m-1} h(Z_{j \cdot s+1}, Z_{j \cdot s+2}, \dots, Z_{j \cdot s+s}).$$

The general form of a complete U-statistic in Eq. (2.1) can be rewritten as

$$U_{n,s} = \frac{1}{n!} \sum_{\beta \in S_n} V(Z_{\beta_1, \beta_2, \dots, \beta_n})$$

where  $S_n$  consists of all permutations of  $(1, 2, \dots, n)$ . Now, suppose that  $(h - \theta)/\sigma$  is sub-Gaussian with variance proxy  $v^2$ , where  $\sigma^2 = \text{var}(h)$ , then by definition, we have

$$\mathbb{E}[\exp(\lambda(h - \theta))] \leq \exp\left(\frac{\lambda^2 \sigma^2 v^2}{2}\right), \quad \lambda \in \mathbb{R} \quad (\text{C.20})$$

and hence we have

$$\begin{aligned} \Pr(U_{n,s} - \theta > t) &\leq \exp(-\lambda t) \mathbb{E}[\exp(\lambda(U_{n,s} - \theta))] \\ &\leq \exp(-\lambda t) \sum_{\beta \in S_n} \frac{1}{n!} \mathbb{E}[\exp(\lambda(V(Z_{\beta_1}, Z_{\beta_2}, \dots, Z_{\beta_n}) - \theta))] \\ &= \exp(-\lambda t) \mathbb{E}[\exp(\lambda(V - \theta))] \\ &\leq \exp\left(-\frac{mt^2}{2\sigma^2 v^2}\right), \quad t > 0. \end{aligned} \quad (\text{C.21})$$

The second inequality in Eq. (C.21) is due to Jensen's inequality and the last inequality is due to Hoeffding inequality. Observe that  $\Pr(U_{n,s} - \theta < t)$  follows in the same manner (recall that Eq. (C.20) holds for all  $\lambda \in \mathbb{R}$ ), and we get

$$\Pr(|U_{n,s} - \theta| \geq t) \leq 2 \exp\left\{-\frac{mt^2}{2\sigma^2 v^2}\right\}. \quad (\text{C.22})$$

Now let  $t = m^{-0.5+\eta}\sigma$  where  $0 < \eta < 0.5$ . Then with probability at least  $1 - 2 \exp(-\frac{1}{2v^2}(\lfloor n/s \rfloor)^{2\eta})$ ,  $|U_{n,s} - \theta| \leq (\lfloor n/s \rfloor)^{1/2-\eta}\sigma$ . Therefore if  $|h - \theta|^2$  and  $|h - \theta|^3$  are sub-Gaussian after being standardized, we can then apply Eq. (C.22) in the proof of Theorem 5 to obtain the improved result.

### Acknowledgments

We would like to thank Larry Wasserman for helpful conversations and feedback, and anonymous reviewers for helpful suggestions.

## References

- [1] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001. [MR3874153](#)
- [2] Mariana Belgiu and Lucian Drăguț. Random forest in remote sensing: A review of applications and future directions. *ISPRS Journal of Photogrammetry and Remote Sensing*, 114:24–31, 2016.
- [3] Yanjun Qi. Random forest for bioinformatics. In *Ensemble machine learning*, pages 307–323. Springer, 2012.
- [4] Luckyson Khaidem, Snehanshu Saha, and Sudeepa Roy Dey. Predicting the direction of stock market prices using random forest. *arXiv preprint arXiv:1605.00003*, 2016.
- [5] Tim Coleman, Lucas Mentch, Daniel Fink, Frank A La Sorte, David W Winkler, Giles Hooker, and Wesley M Hochachka. Statistical inference on tree swallow migrations with random forests. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 69(4):973–989, 2020. [MR4133155](#)
- [6] Manuel Fernández-Delgado, Eva Cernadas, Senén Barro, and Dinani Amorim. Do we need hundreds of classifiers to solve real world classification problems? *The Journal of Machine Learning Research*, 15(1):3133–3181, 2014. [MR3277155](#)
- [7] Leo Breiman, Jerome Friedman, Charles J. Stone, and R.A. Olshen. *Classification and Regression Trees*. Wadsworth, Belmont, CA, 1st edition, 1984. [MR0726392](#)
- [8] Moulinath Banerjee and Ian W McKeague. Confidence sets for split points in decision trees. *The Annals of Statistics*, 35(2):543–574, 2007. [MR2336859](#)
- [9] Hemant Ishwaran. The effect of splitting on random forests. *Machine learning*, 99(1):75–118, 2015. [MR3324236](#)
- [10] Yi Lin and Yongho Jeon. Random forests and adaptive nearest neighbors. *Journal of the American Statistical Association*, 101(474):578–590, 2006. [MR2256176](#)
- [11] Miles E Lopes. Estimating a sharp convergence bound for randomized ensembles. *Journal of Statistical Planning and Inference*, 204:35–44, 2020. [MR3961927](#)
- [12] Gérard Biau, Luc Devroye, and Gábor Lugosi. Consistency of random forests and other averaging classifiers. *Journal of Machine Learning Research*, 9(9), 2008. [MR2447310](#)
- [13] Gérard Biau and Luc Devroye. On the layered nearest neighbour estimate, the bagged nearest neighbour estimate and the random forest method in regression and classification. *Journal of Multivariate Analysis*, 101(10):2499–2518, 2010. [MR2719877](#)
- [14] Gérard Biau. Analysis of a random forests model. *The Journal of Machine Learning Research*, 13:1063–1095, 2012. [MR2930634](#)
- [15] Misha Denil, David Matheson, and Nando Freitas. Consistency of online random forests. In *International Conference on Machine Learning*, pages 1256–1264. PMLR, 2013.
- [16] Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, Michael S

- Lauer, et al. Random survival forests. *The Annals of Applied Statistics*, 2(3):841–860, 2008. [MR2516796](#)
- [17] Hemant Ishwaran and Udaya B Kogalur. Consistency of random survival forests. *Statistics & Probability Letters*, 80(13-14):1056–1064, 2010. [MR2651045](#)
- [18] Yifan Cui, Ruoqing Zhu, Mai Zhou, and Michael Kosorok. Consistency of survival tree and forest models: splitting bias and correction. *arXiv preprint arXiv:1707.09631*, 2019.
- [19] Nicolai Meinshausen and Greg Ridgeway. Quantile regression forests. *Journal of Machine Learning Research*, 7(6), 2006. [MR2274394](#)
- [20] Ruoqing Zhu, Donglin Zeng, and Michael R Kosorok. Reinforcement learning trees. *Journal of the American Statistical Association*, 110(512):1770–1784, 2015. [MR3449072](#)
- [21] Gérard Biau and Erwan Scornet. A random forest guided tour. *Test*, 25(2):197–227, 2016. [MR3493512](#)
- [22] Stefan Wager, Trevor Hastie, and Bradley Efron. Confidence intervals for random forests: The jackknife and the infinitesimal jackknife. *The Journal of Machine Learning Research*, 15(1):1625–1651, 2014. [MR3225243](#)
- [23] Bradley Efron. Estimation and accuracy after model selection. *Journal of the American Statistical Association*, 109(507):991–1007, 2014. [MR3265671](#)
- [24] Erwan Scornet, Gérard Biau, and Jean-Philippe Vert. Consistency of random forests. *The Annals of Statistics*, 43(4):1716–1741, 2015. [MR3357876](#)
- [25] Lucas Mentch and Giles Hooker. Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. *The Journal of Machine Learning Research*, 17(1):841–881, 2016. [MR3491120](#)
- [26] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. [MR3862353](#)
- [27] Paul R Halmos. The theory of unbiased estimation. *The Annals of Mathematical Statistics*, 17(1):34–43, 1946. [MR0015746](#)
- [28] Wassily Hoeffding. A class of statistics with asymptotically normal distribution. In *Breakthroughs in statistics*, pages 308–334. Springer, 1992.
- [29] Edward W Frees. Infinite order u-statistics. *Scandinavian Journal of Statistics*, pages 29–45, 1989. [MR1003967](#)
- [30] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- [31] Joseph P Romano and Cyrus DiCiccio. Multiple data splitting for testing. Technical report, Technical report, 2019.
- [32] Wassily Hoeffding. On sequences of sums of independent random vectors. In *The Collected Works of Wassily Hoeffding*, pages 395–408. Springer, 1994. [MR1307621](#)
- [33] Bradley Efron and Charles Stein. The jackknife estimate of variance. *The Annals of Statistics*, pages 586–596, 1981. [MR0615434](#)
- [34] Tom Rainforth, Rob Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. On nesting monte carlo estimators. In *International Conference on Machine Learning*, pages 4267–4276. PMLR, 2018.
- [35] Erwan Scornet. Random forests and kernel methods. *IEEE Transactions*

- on *Information Theory*, 62(3):1485–1500, 2016. [MR3472261](#)
- [36] Matthew A Olson and Abraham J Wyner. Making sense of random forest probabilities: a kernel perspective. *arXiv preprint arXiv:1812.05792*, 2018.
- [37] Lucas Mentch and Siyu Zhou. Randomization as regularization: A degrees of freedom explanation for random forest success. *Journal of Machine Learning Research*, 21(171):1–36, 2020. [MR4209457](#)
- [38] Lucas Mentch and Siyu Zhou. Getting better from worse: Augmented bagging and a cautionary tale of variable importance. *arXiv preprint arXiv:2003.03629*, 2020.
- [39] Wei Peng, Lucas Mentch, and Leonard Stefanski. Bias, consistency, and alternative perspectives of the infinitesimal jackknife. *arXiv preprint arXiv:2106.05918*, 2021.
- [40] Zhengze Zhou, Lucas Mentch, and Giles Hooker.  $v$ -statistics and variance estimation. *arXiv preprint arXiv:1912.01089*, 2019.
- [41] Tim Coleman, Wei Peng, and Lucas Mentch. Scalable and efficient hypothesis testing with random forests. *arXiv preprint arXiv:1904.07830*, 2019.
- [42] Andrew C Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):122–136, 1941. [MR0003498](#)
- [43] Carl-Gustaf Esseen. On the Liapunov limit error in the theory of probability. *Ark. Mat. Astr. Fys.*, 28:1–19, 1942. [MR0011909](#)
- [44] Herman Callaert, Paul Janssen, et al. The Berry-Esseen theorem for  $u$ -statistics. *The Annals of Statistics*, 6(2):417–421, 1978. [MR0464359](#)
- [45] Y-K Chan and John Wierman. On the Berry-Esseen theorem for  $u$ -statistics. *The Annals of Probability*, pages 136–139, 1977. [MR0433551](#)
- [46] William F Grams, RJ Serfling, et al. Convergence rates for  $u$ -statistics and related statistics. *The Annals of Statistics*, 1(1):153–160, 1973. [MR0336788](#)
- [47] Louis HY Chen, Larry Goldstein, and Qi-Man Shao. *Normal approximation by Stein's method*, volume 2. Springer, 2011. [MR2732624](#)
- [48] Yanglei Song, Xiaohui Chen, and Kengo Kato. Approximating high-dimensional infinite-order  $u$ -statistics: Statistical and computational guarantees. *Electronic Journal of Statistics*, 13(2):4794–4848, 2019. [MR4038726](#)
- [49] Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Central limit theorems and bootstrap in high dimensions. *The Annals of Probability*, 45(4):2309–2352, 2017. [MR3693963](#)
- [50] Xiaohui Chen and Kengo Kato. Randomized incomplete  $u$ -statistics in high dimensions. *The Annals of Statistics*, 47(6):3127–3156, 2019. [MR4025737](#)
- [51] Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000. [MR1652247](#)
- [52] A J Lee. *U-statistics: Theory and Practice*. Routledge, 2019.
- [53] Gérard Biau and Luc Devroye. *Lectures on the nearest neighbor method*, volume 246. Springer, 2015. [MR3445317](#)
- [54] Jerome H Friedman. Multivariate adaptive regression splines. *The Annals of Statistics*, pages 1–67, 1991. [MR1091842](#)
- [55] Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the*

- sixth Berkeley symposium on mathematical statistics and probability, volume 2: Probability theory*, volume 6, pages 583–603. University of California Press, 1972. [MR0402873](#)
- [56] Louis HY Chen, Qi-Man Shao, et al. Normal approximation under local dependence. *The Annals of Probability*, 32(3):1985–2028, 2004. [MR2073183](#)
- [57] Louis HY Chen and Qi-Man Shao. A non-uniform Berry–Esseen bound via stein’s method. *Probability theory and related fields*, 120(2):236–254, 2001. [MR1841329](#)