

Estimation and Inference of Heterogeneous Treatment Effects using Random Forests*

Stefan Wager
Department of Statistics
Stanford University
swager@stanford.edu

Susan Athey
Graduate School of Business
Stanford University
athey@stanford.edu

November 22, 2015

Abstract

Many scientific and engineering challenges—ranging from personalized medicine to customized marketing recommendations—require an understanding of treatment effect heterogeneity. In this paper, we develop a non-parametric *causal forest* for estimating heterogeneous treatment effects that extends Breiman’s widely used random forest algorithm. Given a potential outcomes framework with unconfoundedness, we show that causal forests are pointwise consistent for the true treatment effect, and have an asymptotically Gaussian and centered sampling distribution. We also discuss a practical method for constructing asymptotic confidence intervals for the true treatment effect that are centered at the causal forest estimates. Our theoretical results rely on a generic Gaussian theory for a large family of random forest algorithms; to our knowledge, this is the first set of results that allows any type of random forest, including classification and regression forests, to be used for provably valid statistical inference. In experiments, we find causal forests to be substantially more powerful than classical methods based on nearest-neighbor matching, especially as the number of covariates increases.

KEYWORDS: Adaptive nearest neighbors matching; asymptotic normality; potential outcomes; unconfoundedness.

1 Introduction

A large body of research across many disciplines seeks to draw inferences from data about the causal effect of a treatment. Examples include medical studies about the effect of a drug on health outcomes, studies of the impact of advertising or marketing offers on consumer purchases, evaluations of the effectiveness of government programs or public policies, and “A/B tests” (large-scale randomized experiments) commonly used by technology firms to select algorithms for ranking search results or making recommendations. Historically, most datasets have been too small to meaningfully explore heterogeneity of treatment effects

*Part of the results developed in this paper were made available as an earlier technical report “*Asymptotic Theory for Random Forests*”, available at <http://arxiv.org/abs/1405.0352>.

beyond dividing the sample into a few subgroups. But recently there has been an explosion of empirical settings where it is potentially feasible to customize estimates for individuals.

An impediment to exploring heterogeneous treatment effects is the fear that researchers will iteratively search for subgroups with high treatment levels, and then report only the results for subgroups with extreme effects, thus highlighting heterogeneity that may be purely spurious [Assmann et al., 2000, Cook et al., 2004]. For this reason, protocols for clinical trials must specify in advance which subgroups will be analyzed, and other disciplines such as economics have instituted protocols for registering pre-analysis plans for randomized experiments or surveys. However, it may be possible to discover unexpected and powerful treatment effect heterogeneity if the researcher has a data-driven method that does not impose artificial restrictions. In this paper, we seek to address this challenge by developing a nonparametric method for heterogeneous treatment effect estimation that yields valid asymptotic confidence intervals for the true underlying treatment effect.

Classical approaches to nonparametric estimation of heterogeneous treatment effects include nearest-neighbor matching, kernel methods, and series estimation; see, e.g., Crump et al. [2008], Lee [2009], and Willke et al. [2012]. These methods perform well in applications with a small number of covariates, but quickly break down as the number of covariates increases. In this paper, we explore the use of ideas from the machine learning literature to improve the performance of these classical methods with many covariates. We focus on the family of random forest algorithms introduced by Breiman [2001a], which allow for flexible modeling of interactions in high dimensions by building a large number of regression trees and averaging the models. Random forests are related to kernels and nearest-neighbor methods in that they make predictions using a weighted average of “nearby” observations; however, random forests differ in that they have a data-driven way to determine which nearby observations receive more weight, something that is especially important in environments with many covariates or complex interactions among covariates.

Despite their widespread success at prediction and classification, there are important hurdles that need to be cleared before random forests are directly useful to causal inference. Ideally, an estimator should be consistent with a well-understood asymptotic sampling distribution, so that a researcher can use it to test hypotheses and establish confidence intervals. For example, when deciding to use a drug for an individual, we may wish to test the hypothesis that the expected benefit from the treatment is less than the treatment cost. Asymptotic normality results are especially important in the causal inference setting, both because many policy applications require confidence intervals for decision-making, and because it can be difficult to directly evaluate the model’s performance using, e.g., cross validation, when estimating causal effects. Yet, the asymptotics of random forests have been largely left open, even in the standard regression or classification contexts.

This paper addresses these limitations, developing a forest-based method for treatment effect estimation that allows for a tractable asymptotic theory and valid statistical inference. Following Athey and Imbens [2015], the trees comprising our proposed forest are *causal trees* that estimate the effect of the treatment at the leaves of the trees; we thus refer to our algorithm as a *causal forest*.

In the interest of generality, we begin our theoretical analysis by developing the desired consistency and asymptotic normality results in the context of regression forests. We prove these results for a particular variant of regression forests that uses subsampling to generate a variety of different trees, while it relies on deeply grown trees that satisfy a condition we call “honesty” to reduce bias. An example of an honest tree is one where the tree is grown using one subsample, while the predictions at the leaves of the tree are estimated using a

different subsample. We also show that the heuristically-motivated infinitesimal jackknife for random forests developed by Efron [2014] and Wager et al. [2014] is consistent for the asymptotic variance of random forests in this setting. Our proof builds on classical ideas due to Efron and Stein [1981], Hájek [1968], and Hoeffding [1948], as well as the adaptive nearest neighbors interpretation of random forests of Lin and Jeon [2006]. Given these general results, we next show that our consistency and asymptotic normality results extend from the regression setting to estimating heterogeneous treatment effects in a potential outcomes framework with unconfoundedness [Holland, 1986, Imbens and Rubin, 2015, Rosenbaum and Rubin, 1983, Rubin, 1974].

Although our main focus in this paper is causal inference, we note that there are a variety of important applications of the asymptotic normality result in a pure prediction context. For example, Kleinberg et al. [2015] seek to improve the allocation of medicare funding for hip or knee replacement surgery by detecting patients who had been prescribed such a surgery, but were in fact likely to die of other causes before the surgery would have been useful to them. Here we need predictions for the probability that a given patient will survive for more than, say, one year that come with rigorous confidence statements; our results are the first that enable the use of random forests for this purpose.

Finally, we compare the performance of the causal forest algorithm against classical k -nearest neighbor matching using simulations, finding that the causal forest dominates in terms of both bias and variance in a variety of settings, and that its advantage increases with the number of covariates. We also examine coverage rates of our confidence intervals for heterogeneous treatment effects.

1.1 Related Work

There has been a longstanding understanding in the machine learning literature that prediction methods such as random forests ought to be validated empirically [Breiman, 2001b]: if the goal is prediction, then we should hold out a test set, and the method will be considered as good as its error rate is on this test set. However, there are fundamental challenges with applying a test set approach in the setting of causal inference. In the widely used potential outcomes framework we use to formalize our results [Holland, 1986, Imbens and Rubin, 2015, Rosenbaum and Rubin, 1983, Rubin, 1974], a treatment effect is understood as a difference between two potential outcomes: e.g., would the patient have died if they received the drug vs. if they didn't receive it. Only one of these potential outcomes can ever be observed in practice, and so direct test-set evaluation is in general impossible.¹ Thus, when evaluating estimators of causal effects, asymptotic theory plays a much more important role than in the standard prediction context.

From a technical point of view, the main contribution of this paper is an asymptotic normality theory enabling us to do statistical inference using random forest predictions. Recent results by Biau [2012], Meinshausen [2006], Scornet et al. [2015], Mentch and Hooker [2014], and others have established asymptotic properties of particular variants and simplifications of the random forest algorithm. To our knowledge, however, we provide the first set of conditions under which predictions made by random forests are both asymptotically unbiased and Gaussian, thus allowing for classical statistical inference; the extension to the causal

¹Athey and Imbens [2015] have proposed indirect approaches to mimic test-set evaluation for causal inference. However, these approaches require an estimate of the true treatment effect for each observation in the test set, which creates a new set of challenges. In the absence of an observable ground truth in a test set, statistical theory plays a more central role in evaluating the noise in estimates of causal effects.

forests proposed in this paper is also new. We review the existing theoretical literature on random forests in more detail in Section 3.1.

A small but growing literature seeks to address the problem of flexibly estimating heterogeneous treatment effects. [Green and Kern \[2012\]](#) and [Hill \[2011\]](#) have proposed the use of forest-based algorithms for estimating heterogeneous treatment effects. These papers use the Bayesian additive regression tree (BART) method of [Chipman et al. \[2010\]](#), and report posterior credible intervals obtained by Markov-chain Monte Carlo (MCMC) sampling based on a convenience prior. Although BART is often successful in practice, there are currently no results guaranteeing posterior concentration around the true conditional mean function, or convergence of the MCMC sampler in polynomial time. Another approach to the problem is to divide the training data by treatment status, and apply supervised learning methods to each group separately. For example, [Foster et al. \[2011\]](#) use random forests to estimate the effect of covariates on outcomes in treated and control groups. They then take the difference in predictions as data and project treatment effects onto units' attributes using regression or classification trees. Our approach differs in that we modify the standard random forest algorithm to focus on directly estimating heterogeneity in causal effects, and further we establish the statistical properties of our estimator.

A few papers use tree-based methods for estimating heterogeneous treatment effects; for the trees that underly our random forest estimator, we follow most closely the approach proposed in [Athey and Imbens \[2015\]](#), who propose honest, causal trees, and obtain valid confidence intervals for average treatment effects for each of the subpopulations (leaves) identified by the algorithm. This approach requires the number of observations within a leaf to grow to infinity, and does not personalize predictions for each individual. Other related approaches include those of [Su et al. \[2009\]](#) and [Zeileis et al. \[2008\]](#), which build a tree for treatment effects in subgroups and use statistical tests to determine splits; however, these papers do not analyze bias or consistency properties.

Finally, we note a growing literature on estimating heterogeneous treatment effects using different machine learning methods. [Imai and Ratkovic \[2013\]](#), [Signorovitch \[2007\]](#), [Tian et al. \[2014\]](#) and [Weisberg and Pontes \[2015\]](#) develop lasso-like methods for causal inference in a sparse high-dimensional linear setting. [Beygelzimer and Langford \[2009\]](#), [Dudík et al. \[2011\]](#), and others discuss procedures for transforming outcomes that enable off-the-shelf loss minimization methods to be used for optimal treatment policy estimation. In the econometrics literature, [Dehejia \[2005\]](#), [Hirano and Porter \[2009\]](#), [Manski \[2004\]](#), [Bhattacharya and Dupas \[2012\]](#) estimate parametric or semi-parametric models for optimal policies, relying on regularization for covariate selection in the case of [Bhattacharya and Dupas \[2012\]](#). [Taddy et al. \[2014\]](#) use Bayesian nonparametric methods with Dirichlet priors to flexibly estimate the data-generating process, and then project the estimates of heterogeneous treatment effects down onto the feature space using regularization methods or regression trees to get low-dimensional summaries of the heterogeneity; but again, there are no guarantees about asymptotic properties.

2 Causal Forests

2.1 Treatment Estimation with Unconfoundedness

Suppose we have access to n independent and identically distributed training examples labeled $i = 1, \dots, n$, each of which consists of a feature vector $X_i \in [0, 1]^d$, a response $Y_i \in \mathbb{R}$, and a treatment indicator $W_i \in \{0, 1\}$. Following the potential outcomes framework

[Holland, 1986, Imbens and Rubin, 2015, Rosenbaum and Rubin, 1983, Rubin, 1974], we then posit the existence of potential outcomes $Y_i^{(1)}$ and $Y_i^{(0)}$ corresponding respectively to the response the i -th subject would have experienced with and without the treatment, and define the treatment effect at x as

$$\tau(x) = \mathbb{E} \left[Y_i^{(1)} - Y_i^{(0)} \mid X_i = x \right]. \quad (1)$$

Our goal is to estimate this function $\tau(x)$. The main difficulty is that we can only ever observe one of the two potential outcomes $Y_i^{(0)}, Y_i^{(1)}$ for a given training example, and so cannot directly train machine learning method on differences of the form $Y_i^{(1)} - Y_i^{(0)}$.

In general, we cannot estimate $\tau(x)$ simply from the observed data (X_i, Y_i, W_i) without further restrictions on the data generating distribution. A standard way to make progress is to assume unconfoundedness [Rosenbaum and Rubin, 1983], i.e., that the treatment assignment W_i is independent of the potential outcomes for Y_i conditional on X_i :

$$\left\{ Y_i^{(0)}, Y_i^{(1)} \right\} \perp\!\!\!\perp W_i \mid X_i. \quad (2)$$

The motivation behind this unconfoundedness is that, given continuity assumptions, it effectively implies that we can treat nearby observations in x -space as having come from a randomized experiment; thus, nearest-neighbor matching and other local methods will in general be consistent for $\tau(x)$.

An immediate consequence of unconfoundedness is that

$$\mathbb{E} \left[Y_i \left(\frac{W_i}{e(x)} - \frac{1 - W_i}{1 - e(x)} \right) \mid X_i = x \right] = \tau(x), \quad \text{where } e(x) = \mathbb{E} [W_i \mid X_i = x] \quad (3)$$

is the propensity of receiving treatment at x . Thus, if we knew $e(x)$, we would have access to a simple unbiased estimator for $\tau(x)$; this observation lies at the heart of methods based on propensity weighting [e.g., Hirano et al., 2003]. Many early applications of machine learning to causal inference effectively reduce to estimating $e(x)$ using, e.g., boosting, a neural network, or even random forests, and then transforming this into an estimate for $\tau(x)$ using (3) [e.g., McCaffrey et al., 2004, Westreich et al., 2010]. In this paper, we take a more indirect approach: we show that, under regularity assumptions, causal forests can use the unconfoundedness assumption (2) to achieve consistency without needing to explicitly estimate the propensity $e(x)$.

2.2 From Regression Trees to Causal Trees and Forests

At a high level, trees and forests can be thought of as nearest neighbor methods with an adaptive neighborhood metric. Given a test point x , classical methods such as k -nearest neighbors seek the k closest points to x according to some pre-specified distance measure, e.g., Euclidean distance. In contrast, tree-based methods also seek to find training examples that are close to x , but now closeness is defined with respect to a decision tree, and the closest points to x are those that fall in the same leaf as it. The advantage of trees is that their leaves can be narrower along the directions where the signal is changing fast and wider along the other directions, potentially leading a to a substantial increase in power when the dimension of the feature space is even moderately large.

In this section, we seek to build causal trees that resemble their regression analogues as closely as possible. In an effort to elucidate this connection, suppose first that we only

observe independent samples (X_i, Y_i) , and want to build a CART regression tree. We start by recursively splitting the feature space until we have partitioned it into a set of leaves L , each of which only contains a few training samples. Then, given a test point x , we evaluate the prediction $\hat{\mu}(x)$ by identifying the leaf $L(x)$ containing x and setting

$$\hat{\mu}(x) = \frac{1}{|\{i : X_i \in L(x)\}|} \sum_{\{i : X_i \in L(x)\}} Y_i. \quad (4)$$

Heuristically, this strategy is well-motivated if we believe the leaf $L(x)$ to be small enough that the responses Y_i inside the leaf are roughly identically distributed. There are several procedures for how to place the splits in the decision tree; see, e.g., [Hastie et al. \[2009\]](#).

In the context of causal trees, we analogously want to think of the leaves as small enough that the (Y_i, W_i) pairs corresponding to the indices i for which $i \in L(x)$ act as though they had come from a randomized experiment. Then, it is natural to estimate the treatment effect for any $x \in L$ as

$$\hat{\tau}(x) = \frac{1}{|\{i : W_i = 1, X_i \in L\}|} \sum_{\{i : W_i = 1, X_i \in L\}} Y_i - \frac{1}{|\{i : W_i = 0, X_i \in L\}|} \sum_{\{i : W_i = 0, X_i \in L\}} Y_i. \quad (5)$$

In the following sections, we will establish that such trees can be used to grow causal forests that are consistent for $\tau(x)$.²

Finally, given a procedure for generating a single causal tree, a causal forest generates an ensemble of B such trees, each of which outputs an estimate $\hat{\tau}_b(x)$. The forest then aggregates their predictions by averaging them: $\hat{\tau}(x) = B^{-1} \sum_{b=1}^B \hat{\tau}_b(x)$. We always assume that the individual causal trees comprising the forest are built using random subsamples of s training examples, where $s/n \ll 1$; for our theoretical results, we will assume that $s \asymp n^\beta$ for some $\beta < 1$. The advantage of a forest over a single tree is that it is not always clear what the “best” causal tree is. In this case, as shown by [Breiman \[2001a\]](#), it is often better to generate many different decent-looking trees and average their predictions, instead of seeking a single highly-optimized tree. In practice, this aggregation scheme helps reduce variance and smooths sharp decision boundaries.

2.3 Asymptotic Inference with Causal Forests

Our results require some conditions on the forest-growing scheme, outlined in [Section 2.4](#): the trees comprising the forest must be grown on subsamples of the training data, and the splitting rule must not inappropriately incorporate information about the outcomes Y_i . However, given these high level conditions, we obtain a widely applicable consistency result that applies to several different interesting causal forest algorithms.

Our first result is that causal forests are consistent for the true treatment effect $\tau(x)$. To achieve pointwise consistency, we need to assume that the conditional mean functions $\mathbb{E}[Y^{(0)} | X = x]$ and $\mathbb{E}[Y^{(1)} | X = x]$ are both Lipschitz continuous. To our knowledge, all existing results on pointwise consistency of regression forests [e.g., [Meinshausen, 2006](#), [Biau,](#)

²The causal tree algorithm presented above is a simplification of the method of [Athey and Imbens \[2015\]](#). The main difference between our approach and that of [Athey and Imbens \[2015\]](#) is that they seek to build a single well-tuned tree; to this end, they use fairly large leaves and apply a form propensity weighting based on [\(3\)](#) within each leaf to correct for variations in $e(x)$ inside the leaf. In contrast, we follow [Breiman \[2001a\]](#) and build our causal forest using deep trees. Since our leaves are small, we do not need to apply any additional corrections inside them.

2012] require an analogous condition on $\mathbb{E}[Y \mid X = x]$. This is not particularly surprising, as forests generally have smooth response surfaces [Bühlmann and Yu, 2002].

In addition to continuity assumptions, we also need to assume that we have overlap, i.e., for some $\varepsilon > 0$ and all $x \in [0, 1]^d$,

$$\varepsilon < \mathbb{P}[W = 1 \mid X = x] < 1 - \varepsilon. \quad (6)$$

This condition effectively guarantees that, for large enough n , there will be enough treatment and control units near any test point x for local methods to work.

Theorem (informal statement). *If we have unconfoundedness (2) and overlap (6), and the above Lipschitz conditions hold, then causal forests are pointwise consistent.*

Beyond consistency, in order to do statistical inference on the basis of the estimated treatment effects $\hat{\tau}(x)$, we need to understand their asymptotic sampling distribution. Using the potential nearest neighbors construction of Lin and Jeon [2006] and classical analysis tools going back to Hoeffding [1948] and Hájek [1968], we show that—provided the sub-sample size s scales appropriately with n —the predictions made by a causal forest are asymptotically Gaussian and unbiased.

Moreover, we show that the asymptotic variance of causal forests can be accurately estimated. To do so, we use the infinitesimal jackknife for random forests developed by Efron [2014] and Wager et al. [2014], based on the original infinitesimal jackknife procedure of Jaeckel [1972]. This method assumes that we have taken the number of trees B to be large enough that the Monte Carlo variability of the forest does not matter; and only measures the randomness in $\hat{\tau}(x)$ due to the training sample.

To define the variance estimates, let $\hat{\tau}_b^*(x)$ be the treatment effect estimate given by the b -th tree, and let $N_{ib}^* \in \{0, 1\}$ indicate whether or not the i -th training example was used for the b -th tree.³ Then, we set

$$\widehat{V}_{IJ}(x) = \frac{n-1}{n} \left(\frac{n}{n-s} \right)^2 \sum_{i=1}^n \text{Cov}_* [\hat{\tau}_b^*(x), N_{ib}^*]^2, \quad (7)$$

where the covariance is taken with respect to the set of all the trees $b = 1, \dots, B$ used in the forest. The term $n(n-1)/(n-s)^2$ is a finite-sample correction for forests grown by subsampling without replacement; see Proposition 10.

Theorem (informal statement). *Given the conditions required for consistency, suppose moreover that the subsample size s scales as $s \asymp n^\beta$ for some $\beta_{\min} < \beta < 1$. Then causal forests are asymptotically normal with*

$$(\hat{\tau}(x) - \tau(x)) / \sqrt{\widehat{V}_{IJ}(x)} \Rightarrow \mathcal{N}(0, 1). \quad (8)$$

2.4 Honest Splitting Rules

In our discussion so far, we have emphasized the flexible nature of our results: for a wide variety of causal forests that can be tailored to the application area, we achieve both consistency and centered asymptotic normality, provided the sub-sample size s scales at an

³For double-sample trees defined in Procedure 1, $N_{ib}^* = 1$ if the i -th example appears in either the \mathcal{I} -sample or the \mathcal{J} -sample.

Procedure 1. DOUBLE-SAMPLE TREES

Double-sample trees split the available training data into two parts: one half for estimating the desired response inside each leaf, and another half for placing splits.

Input: n training examples of the form (X_i, Y_i) for regression trees or (X_i, Y_i, W_i) for causal trees, where X_i are features, Y_i is the response, and W_i is the treatment assignment. A minimum leaf size k .

1. Draw a random subsample of size s from $\{1, \dots, n\}$ without replacement, and then divide it into two disjoint sets of size $|\mathcal{I}| = \lfloor s/2 \rfloor$ and $|\mathcal{J}| = \lceil s/2 \rceil$.
2. Grow a tree via recursive partitioning. The splits are chosen using any data from the \mathcal{J} sample and X - or W -observations from the \mathcal{I} sample, but without using Y -observations from the \mathcal{I} -sample.
3. Estimate leaf-wise responses using only the \mathcal{I} -sample observations.

Double-sample *regression* trees make predictions $\hat{\mu}(x)$ using (4) on the leaf containing x , only using the \mathcal{I} -sample observations. The splitting criteria is the standard for CART regression trees (minimizing mean-squared error of predictions). Splits are restricted so that each leaf of the tree must contain k or more \mathcal{I} -sample observations.

Double-sample *causal* trees are defined similarly, except that for prediction we estimate $\hat{\tau}(x)$ using (5) on the \mathcal{I} sample. Following [Athey and Imbens \[2015\]](#), the splits of the tree are chosen by maximizing the variance of $\hat{\tau}(X_i)$ for $i \in \mathcal{J}$; see [Remark 1](#) for details. In addition, each leaf of the tree must contain k or more \mathcal{I} -sample observations of *each* treatment class.

appropriate rate. Our results do, however, require the splitting rule to satisfy a fairly strong condition, which we call honesty: a splitting rule is honest if, for each training example i , it only uses the response Y_i to estimate the within-leaf treatment effect τ using (5) or to decide where to place the splits, but not both. We discuss two causal forest algorithms that satisfy this condition.

Our first algorithm, which we call a double-sample tree, achieves honesty by dividing its training subsample into two halves \mathcal{I} and \mathcal{J} . Then, it uses the \mathcal{J} -sample to place the splits, while holding out the \mathcal{I} -sample to do within-leaf estimation; see [Procedure 1](#) for details. A similar family of algorithms was discussed in detail by [Denil et al. \[2014\]](#), who showed that such forests could achieve competitive performance relative to standard tree algorithms that do not divide their training samples.

Another way to build honest trees is to ignore the outcome data Y_i when placing splits, and instead first train a classification tree for the treatment assignments W_i ([Procedure 2](#)). Such propensity trees can be particularly useful in observational studies, where we want to minimize bias due to variation in $e(x)$. Seeking estimators that match training examples based on estimated propensity is a longstanding idea in causal inference, going back to [Rosenbaum and Rubin \[1983\]](#).

It appears that our honesty condition is required to achieve valid statistical inference

Procedure 2. PROPENSITY TREES

Propensity trees use only the treatment assignment indicator W_i to place splits, and save the responses Y_i for estimating τ .

Input: n training examples (X_i, Y_i, W_i) , where X_i are features, Y_i is the response, and W_i is the treatment assignment. A minimum leaf size k .

1. Draw a random subsample $\mathcal{I} \in \{1, \dots, n\}$ of size $|\mathcal{I}| = s$ (no replacement).
2. Train a classification tree using sample \mathcal{I} where the outcome is the treatment assignment, i.e., on the (X_i, W_i) pairs with $i \in \mathcal{I}$. Each leaf of the tree must have k or more observations of *each* treatment class.
3. Estimate $\tau(x)$ using (5) on the leaf containing x .

In step 2, the splits are chosen by optimizing, e.g., the Gini criterion used by CART for classification [Breiman et al., 1984].

using causal forests in particular, and also random forests in general. Recent theoretical advances have established that regression forests can be consistent even if this honesty condition does not hold [Scornet et al., 2015, Wager and Walther, 2015]. However, the rates given in both papers allow for the forests to be asymptotically bias- instead of variance-dominated. Thus, at least until we develop techniques for debiasing random forests, honesty appears to be a necessary ingredient in valid inference for τ . We discuss this issue further in Section 3.

Remark 1. For completeness, we briefly outline the motivation for the splitting rule of Athey and Imbens [2015] we use for our double-sample trees. This method is motivated by an algorithm for minimizing the squared-error loss in regression trees. Because regression trees compute predictions $\hat{\mu}$ by averaging training responses over leaves, we can verify that

$$\sum_{i \in \mathcal{J}} (\hat{\mu}(X_i) - Y_i)^2 = \sum_{i \in \mathcal{J}} Y_i^2 - \sum_{i \in \mathcal{J}} \hat{\mu}(X_i)^2. \quad (9)$$

Thus, finding the squared-error minimizing split is equivalent to maximizing the variance of $\hat{\mu}(X_i)$ for $i \in \mathcal{J}$; note that $\sum_{i \in \mathcal{J}} \hat{\mu}(X_i) = \sum_{i \in \mathcal{J}} Y_i$ for all trees, and so maximizing variance is equivalent to maximizing the sum of the $\hat{\mu}(X_i)^2$. In Procedure 1, we emulate this algorithm by picking splits that maximize the variance of $\hat{\tau}(X_i)$ for $i \in \mathcal{J}$.

3 Asymptotic Theory for Random Forests

In order to use random forests to provide formally valid statistical inference, we need an asymptotic normality theory for random forests. In the interest of generality, we first develop such a theory in the context of classical regression forests, as originally introduced by Breiman [2001a]. We then return to our main topic and show how to apply this theory to estimating heterogeneous treatment effects with causal forests in Section 4.

In this section, we assume that we have training examples $Z_i = (X_i, Y_i)$ for $i = 1, \dots, n$, a test point x , and we want to estimate true conditional mean function

$$\mu(x) = \mathbb{E}[Y \mid X = x]. \quad (10)$$

We also have access to a regression tree T which can be used to get estimates of the conditional mean function at x of the form $T(x; \xi, Z_1, \dots, Z_n)$, where $\xi \sim \Xi$ is a source of auxiliary randomness. Our goal is to use this tree-growing scheme to build a random forest that can be used for valid statistical inference about $\mu(x)$.

We begin by precisely describing how we aggregate individual trees into a forest. For us, a random forest is an average of trees trained over all possible size- s subsamples of the training data, marginalizing over the auxiliary noise ξ . In practice, we compute such a random forest by Monte Carlo averaging, and set

$$\text{RF}(x; Z_1, \dots, Z_n) \approx \frac{1}{B} \sum_{b=1}^B T(x; \xi_b^*, Z_{b1}^*, \dots, Z_{bs}^*), \quad (11)$$

where $\{Z_{b1}^*, \dots, Z_{bs}^*\}$ is drawn without replacement from $\{Z_1, \dots, Z_n\}$, ξ_b^* is a random draw from Ξ , and B is the number of Monte Carlo replicates we can afford to perform. The formulation (12) arises as the $B \rightarrow \infty$ limit of (11); thus, our theory effectively assumes that B is large enough for Monte Carlo effects not to matter. The effects of using a finite B are studied in detail by [Mentch and Hooker \[2014\]](#).

Definition 1. The *random forest* with base learner T and subsample size s is

$$\text{RF}(x; Z_1, \dots, Z_n) = \binom{n}{s}^{-1} \sum_{1 \leq i_1 < i_2 < \dots < i_s \leq n} \mathbb{E}_{\xi \sim \Xi} [T(x; \xi, Z_{i_1}, \dots, Z_{i_s})]. \quad (12)$$

Next, as described in Section 2, we require that the trees T comprising our forest be honest. Double-sample trees, as defined in Procedure 1, can always be used to build honest trees with respect to the \mathcal{I} -sample. In the context of causal trees for observational studies, propensity trees (Procedure 2) provide a simple recipe for building honest trees without sample splitting.

Definition 2. A tree grown on a training sample $(Z_1 = (X_1, Y_1), \dots, Z_s = (X_s, Y_s))$ is *honest* if the tree does not use the responses Y_1, \dots, Y_s in choosing where to place its splits.

In order to guarantee consistency, we also need to enforce that the leaves of the trees become small in *all* dimensions of the feature space as n gets large.⁴ Here, we achieve this effect by enforcing some randomness in the way trees choose the variables they split on: at each step, each variable is selected with probability at least π/d for some $0 < \pi \leq 1$. Formally, the randomness in how to pick the splitting features is contained in the auxiliary random variable ξ .

Definition 3. A tree is a *random-split* tree if at every step of the tree-growing procedure, marginalizing over ξ , the probability that the next split occurs along the j -th feature is bounded below by π/d for some $0 < \pi \leq 1$, for all $j = 1, \dots, d$.

⁴[Biau \[2012\]](#) and [Wager and Walthers \[2015\]](#) consider the estimation of low-dimensional signals embedded in a high-dimensional ambient space using random forests; in this case, the variable selection properties of trees also become important. We leave a study of asymptotic normality of random forests in high dimensions to future work.

The remaining definitions are more technical. We use regularity to control the shape of the tree leaves, while symmetry is used to apply classical tools in establishing asymptotic normality.

Definition 4. A tree predictor grown by recursive partitioning is *regular* if either (a) (*standard case*) each split leaves at least a fraction $\alpha > 0$ of the available training examples on each side of the split and, moreover, the trees are fully grown to depth k for some $k \in \mathbb{N}$, i.e., there are between k and $2k - 1$ observations in each terminal node of the tree; or (b) (*double sample case*) if the predictor is a double-sample tree as in Procedure 1, the tree satisfies part (a) for the \mathcal{I} sample.

Definition 5. A predictor is *symmetric* if the (possibly randomized) output of the predictor does not depend on the order ($i = 1, 2, \dots$) in which the training examples are indexed.

Finally, in the context of classification and regression forests, we estimate the asymptotic variance of random forests using the original infinitesimal jackknife of Wager et al. [2014], i.e.,

$$\widehat{V}_{IJ}(x) = \frac{n-1}{n} \left(\frac{n}{n-s} \right)^2 \sum_{i=1}^n \text{Cov}_* [\hat{\mu}_b^*(x), N_{ib}^*]^2, \quad (13)$$

where $\hat{\mu}_b^*(x)$ is the estimate for $\mu(x)$ given by a single regression tree. We note that the finite-sample correction $n(n-1)/(n-s)^2$ did not appear in Wager et al. [2014], as their paper focused on subsampling with replacement, whereas this correction is only appropriate for subsampling without replacement.

Given these preliminaries, we can state our main result on the asymptotic normality of random forests. As discussed in Section 2.3, we require that the conditional mean function $\mu(x) = \mathbb{E}[Y | X = x]$ be Lipschitz continuous. The asymptotic normality result requires for the subsample size s to scale within the bounds given in (14). If the subsample size grows slower than this, the forest will still be asymptotically normal, but the forest may be asymptotically biased. For clarity, we state the following result with that makes the dependence of $\hat{\mu}_n(x)$ and s_n on n explicit; in most of the paper, however, we drop the subscripts to $\hat{\mu}_n(x)$ and s_n when there is no risk of confusion.

Theorem 1. *Suppose that we have n independent and identically distributed training examples $Z_i = (X_i, Y_i) \in [0, 1]^d \times [-M, M]$ for some $M > 0$. Suppose moreover that the features are independently and uniformly distributed⁵ $X_i \sim U([0, 1]^d)$, that $\mu(x) = \mathbb{E}[Y | X = x]$ and $\mu_2(x) = \mathbb{E}[Y^2 | X = x]$ are Lipschitz-continuous, and $\text{Var}[Y | X = x] > 0$. Given this data-generating process, let T be an honest, regular, and symmetric random-split tree in the sense of Definitions 2, 3, 4, and 5, and let $\hat{\mu}_n(x)$ be the estimate for $\mu(x)$ given by a random forest with base learner T and a subsample size s_n . Finally, suppose that the subsample size s_n scales as*

$$s_n \asymp n^\beta \text{ for some } \beta_{\min} := 1 - \left(1 + \frac{d}{\pi} \frac{\log(\alpha^{-1})}{\log((1-\alpha)^{-1})} \right)^{-1} < \beta < 1. \quad (14)$$

Then, random forest predictions are asymptotically Gaussian:

$$\frac{\hat{\mu}_n(x) - \mu(x)}{\sigma_n(x)} \Rightarrow \mathcal{N}(0, 1) \text{ for a sequence } \sigma_n(x) \rightarrow 0. \quad (15)$$

⁵The result also holds with a density that is bounded away from 0 and infinity; however, we assume uniformity for simpler exposition.

Moreover, the asymptotic variance σ_n can be consistently estimated using the infinitesimal jackknife (7):

$$\widehat{V}_{IJ}(x) / \sigma_n^2(x) \rightarrow_p 1. \quad (16)$$

Remark 2 (binary classification). We note that Theorem 1 also holds for binary classification forests with leaf size $k = 1$, as is default in the R-package `randomForest` [Liaw and Wiener, 2002]. Here, we treat the output $\text{RF}(x)$ of the random forests as an estimate for the probability $\mathbb{P}[Y = 1 \mid X = x]$; Theorem 1 then lets us construct valid confidence intervals for this probability. For classification forests with $k > 1$, the proof of Theorem 1 still holds if the individual classification trees are built by *averaging* observations within a leaf, but not if they are built by *voting*. Extending our results to voting trees is left as further work.

The proof of this result is organized as follows. In Section 3.2, we provide bounds for the bias $\mathbb{E}[\hat{\mu}_n(x) - \mu(x)]$ of random forests, while Section 3.3 studies the sampling distributions of $\hat{\mu}_n(x) - \mathbb{E}[\hat{\mu}_n(x)]$ and establishes Gaussianity. Given a subsampling rate satisfying (14), the bias decays faster than the variance, thus allowing for (15). Before beginning the proof, however, we relate our result to existing results about random forests in Section 3.1.

3.1 Theoretical Background

There has been considerable work in understanding the theoretical properties of random forests. The convergence and consistency properties of trees and random forests have been studied by, among others, Biau [2012], Biau et al. [2008], Breiman [2004], Breiman et al. [1984], Meinshausen [2006], Scornet et al. [2015], and Wager and Walther [2015]. Meanwhile, their sampling variability has been analyzed by Duan [2011], Lin and Jeon [2006], Mentch and Hooker [2014], Sexton and Laake [2009], and Wager et al. [2014]. However, to our knowledge, our Theorem 1 is the first result establishing conditions under which predictions made by random forests are asymptotically unbiased and normal.

Probably the closest existing result is that of Mentch and Hooker [2014], who showed that random forests based on subsampling are asymptotically normal under substantially strong conditions than us: they require that the subsample size s grows slower than \sqrt{n} , i.e., that $s_n/\sqrt{n} \rightarrow 0$. However, under these conditions, random forests will not in general be asymptotically unbiased. As a simple example, suppose that $d = 2$, that $\mu(x) = \|x\|_1$, and that we evaluate an honest random forest at $x = 0$. A quick calculation shows that the bias of the random forest decays as $1/\sqrt{s_n}$, while its variance decays as s_n/n . If $s_n/\sqrt{n} \rightarrow 0$, the squared bias decays slower than the variance, and so the forest cannot be asymptotically unbiased. Thus, although the result of Mentch and Hooker [2014] may appear qualitatively similar to ours, it cannot be used for valid statistical asymptotic inference about $\mu(x)$.

The variance estimator \widehat{V}_{IJ} was studied in the context of random forests by Wager et al. [2014], who showed empirically that the method worked well for many problems of interest. Wager et al. [2014] also emphasized that, when using \widehat{V}_{IJ} in practice, it is important to account for Monte Carlo bias. Our analysis provides theoretical backing to these results, by showing that \widehat{V}_{IJ} is in fact a consistent estimate for the variance $\sigma_n^2(x)$ of random forest predictions. The earlier work on this topic [Efron, 2014, Wager et al., 2014] had only motivated the estimator \widehat{V}_{IJ} by highlighting connections to classical statistical ideas, but did not establish any formal justification for it.

Instead of using subsampling, Breiman originally described random forests in terms of

bootstrap sampling, or bagging [Breiman, 1996],

$$\widetilde{\text{RF}}(x; Z_1, \dots, Z_n) = \frac{1}{B} \sum_{b=1}^B T(x; \xi_b^*, Z_{b1}^*, \dots, Z_{bn}^*), \quad (17)$$

where the $Z_{bi}^* \stackrel{\text{iid}}{\sim} \{Z_1, \dots, Z_n\}$ form a bootstrap sample of the training data. Random forests of the form (17), however, have proven to be remarkably resistant to classical statistical analysis. As observed by Buja and Stuetzle [2006], Chen and Hall [2003], Friedman and Hall [2007] and others, estimators of this form can exhibit surprising properties even in simple situations; meanwhile, using subsampling rather than bootstrap sampling has been found to avoid several pitfalls [e.g., Politis et al., 1999]. Although they are less common in the literature, random forests based on subsampling have also been occasionally studied and found to have good practical and theoretical properties [e.g., Bühlmann and Yu, 2002, Mentch and Hooker, 2014, Scornet et al., 2015, Strobl et al., 2007].

Finally, an interesting question for further theoretical study is to understand the optimal scaling of the subsample size s_n for minimizing the mean-squared error of random forests. For subsampled nearest-neighbors estimation, the optimal rate for s_n is $s_n \asymp n^{1-(1+d/4)^{-1}}$ [Biau et al., 2010, Samworth, 2012]. Here, our specific value for β_{\min} depends on the upper bounds for bias developed in the following section. Now, as shown by Biau [2012], under some sparsity assumptions on $\mu(x)$, it is possible to get substantially stronger bounds for this bias; thus, it is plausible that under similar conditions we could push back the lower bound β_{\min} on the growth rate of the subsample size.

3.2 Bias and Honesty

We start by bounding the bias of regression trees. Our approach relies on showing that as the sample size s available to the tree gets large, its leaves get small; Lipschitz-continuity of the conditional mean function and honesty then let us bound the bias. In order to state a formal result, define the *diameter* $\text{diam}(L(x))$ of a leaf $L(x)$ as the length of the longest segment contained inside $L(x)$, and similarly let $\text{diam}_j(L(x))$ denote the length of the longest such segment that is parallel to the j -th axis. The following lemma is a refinement of a result of Meinshausen [2006], who showed that $\text{diam}(L(x)) \rightarrow_p 0$ for regular trees.

Lemma 2. *Let T be a regular, random-split tree and let $L(x)$ denote its leaf containing x . Suppose that $X_1, \dots, X_s \sim U([0, 1]^d)$ independently. Then, for any $0 < \eta < 1$, and for large enough s ,*

$$\mathbb{P} \left[\text{diam}_j(L(x)) \geq \left(\frac{s}{2k-1} \right)^{-\frac{0.99(1-\eta) \log((1-\alpha)^{-1})}{\log(\alpha^{-1})} \frac{\pi}{d}} \right] \leq \left(\frac{s}{2k-1} \right)^{-\frac{\eta^2}{2} \frac{1}{\log(\alpha^{-1})} \frac{\pi}{d}}.$$

This lemma then directly translates into a bound on the bias of a single regression tree. Since a forest is an average of independently-generated trees, the bias of the forest is the same as the bias of a single tree.

Theorem 3. *Under the conditions of Lemma 2, suppose moreover that $\mu(x)$ is Lipschitz continuous, that $|Y| \leq M$, and that the trees T comprising the random forest are honest.*

Then, provided that $\alpha \leq 0.2$, the bias of the random forest at x is bounded by

$$|\mathbb{E} [\hat{\mu}(x)] - \mu(x)| = \mathcal{O} \left(s^{-\frac{1}{2} \frac{\log((1-\alpha)^{-1})}{\log(\alpha^{-1})} \frac{\pi}{d}} \right);$$

the constant in the \mathcal{O} -bound is given in the proof.

Perhaps the largest deviation between our assumptions so far and the commonly-used CART trees of Breiman et al. [1984] is our honesty assumption. CART trees were designed to use the available data as efficiently as possible, so honesty may have appeared prohibitively expensive: if we hold out half the data to get within-leaf estimates $\hat{\mu}_n(x)$, then we only have half as much data that can be used to pick good splits and learn deep trees. In our context, however, since we already train our trees on subsamples of size $s \ll n$, honesty may be less restrictive in practice: if we double the size of s , then we can train trees that are just as deep as before while retaining honesty of predictions.

At a more fundamental level, it is natural to ask whether honesty is required in order to get forests whose predictions are asymptotically unbiased. There has been some recent work showing that some adaptive (i.e., non-honest) trees can be consistent. Scornet et al. [2015] found that, under conditions on the data-generating function, the risk of CART random forests converges to the Bayes risk,⁶ i.e., the test set error of a random forest converges to the Bayes error as $n \rightarrow \infty$. Meanwhile, Wager and Walther [2015] show that adaptive trees with growing leaves are in general consistent under fairly weak conditions, and provide bounds on the decay rate of their bias. However, bounds on the bias established in these papers decay slower than generous upper bounds on the variance of the random forest; thus, we cannot use them to obtain asymptotic unbiasedness in the sense of (15). Given the technology currently available for working with random forests, it appears to us that honesty is a necessary ingredient in obtaining asymptotically normal and unbiased limiting distributions for random forest predictions.

3.3 Asymptotic Normality of Random Forests

Our analysis of the asymptotic normality of random forests builds on ideas developed by Hoeffding [1948] and Hájek [1968] for understanding classical statistical estimators such as U -statistics. We begin by briefly reviewing their results to give some context to our proof. Given a predictor T and independent training examples Z_1, \dots, Z_n , the Hájek projection of T is defined as

$$\overset{\circ}{T} = \mathbb{E} [T] + \sum_{i=1}^n (\mathbb{E} [T | Z_i] - \mathbb{E} [T]). \quad (18)$$

In other words, the Hájek projection of T captures the first-order effects in T . This projection has the properties we would expect: in particular $\text{Var} [\overset{\circ}{T}] \leq \text{Var} [T]$, and if

$$\lim_{n \rightarrow \infty} \text{Var} [\overset{\circ}{T}] / \text{Var} [T] = 1, \text{ then } \lim_{n \rightarrow \infty} \mathbb{E} \left[\left\| \overset{\circ}{T} - T \right\|_2^2 \right] / \text{Var} [T] = 0. \quad (19)$$

Since the Hájek projection $\overset{\circ}{T}$ is a sum of independent random variables, we should expect it to be asymptotically normal under all but pathological conditions. Thus whenever the

⁶We note that this form of average-case consistency does not imply pointwise consistency of the kind sought in this paper. In our technical report [Wager, 2014], we describe a simple example where $\mu(x)$ is constant yet adaptive CART-based random forests appear to be inconsistent.

ratio of the variance of $\overset{\circ}{T}$ to that of T tends to 1, the theory of Hájek projections almost automatically guarantees that T will be asymptotically normal.⁷

If T is a regression tree, however, the condition from (19) does not apply, and we cannot use the classical theory of Hájek projections directly. Our analysis is centered around a weaker form of this condition, which we call ν -incrementality. With our definition, predictors T to which we can apply the argument (19) directly are 1-incremental.

Definition 6. The predictor T is $\nu(s)$ -incremental at x if

$$\text{Var} \left[\overset{\circ}{T}(x; Z_1, \dots, Z_s) \right] / \text{Var} [T(x; Z_1, \dots, Z_s)] \gtrsim \nu(s),$$

where $\overset{\circ}{T}$ is the Hájek projection of T (18). In our notation,

$$f(s) \gtrsim g(s) \text{ means that } \liminf_{s \rightarrow \infty} f(s)/g(s) \geq 1.$$

Our argument proceeds in two steps. First we establish lower bounds for the incrementality of regression trees in Section 3.3.1. Then, in Section 3.3.2 we show how we can turn weakly incremental predictors T into 1-incremental ensembles by subsampling (Lemma 7), thus bringing us back into the realm of classical theory. We also establish the consistency of the infinitesimal jackknife for random forests. Our analysis of regression trees is motivated by the “potential nearest neighbors” model for random forests introduced by Lin and Jeon [2006]; the key technical device used in Section 3.3.2 is the ANOVA decomposition of Efron and Stein [1981]. The discussion of the infinitesimal jackknife for random forest builds on results of Efron [2014] and Wager et al. [2014].

3.3.1 Regression Trees and Incremental Predictors

Analyzing specific greedy tree models such as CART trees can be challenging. We thus follow the lead of Lin and Jeon [2006], and analyze a more general class of predictors—potential nearest neighbors predictors—that operate by doing a nearest-neighbor search over rectangles.

Definition 7. Consider a set of points $X_1, \dots, X_s \in \mathbb{R}^d$ and a fixed $x \in \mathbb{R}^d$. A point X_i is a *potential nearest neighbor* (PNN) of x if the smallest axis-aligned hyperrectangle with vertices x and X_i contains no other points X_j . Extending this notion, a *PNN k -set* of x is a set of points $\Lambda \subseteq \{X_1, \dots, X_s\}$ of size $k \leq |\Lambda| < 2k - 1$ such that there exists an axis aligned hyperrectangle L containing x , Λ , and no other training points. A training example X_i is called a *k -PNN* of x if there exists a PNN k -set of x containing X_i . Finally, a predictor T is a *k -PNN predictor* over $\{Z\}$ if, given a training set

$$\{Z\} = \{(X_1, Y_1), \dots, (X_s, Y_s)\} \in \{\mathbb{R}^d \times \mathcal{Y}\}^s$$

and a test point $x \in \mathbb{R}^d$, T always outputs the average of the responses Y_i over a k -PNN set of x .

This formalism allows us to describe a wide variety of tree predictors. For example, as shown by Lin and Jeon [2006], any decision tree T that makes axis-aligned splits and has

⁷The moments defined in (18) depend on the data-generating process for the Z_i , and so cannot be observed in practice. Thus, the Hájek projection is mostly useful as an abstract theoretical tool. For a review of classical projection arguments, see Chapter 11 of Van der Vaart [2000].

leaves of size between k and $2k - 1$ is a k -PNN predictor. In particular, the base learners originally used by Breiman [2001a], namely CART trees grown up to a leaf size k [Breiman et al., 1984], are k -PNN predictors. Predictions made by k -PNN predictors can always be written as

$$T(x; \xi, Z_1, \dots, Z_s) = \sum_{i=1}^s S_i Y_i, \quad (20)$$

where S_i is a selection variable that takes the value $1/|\{i : X_i \in L(x)\}|$ for indices i in the selected leaf-set $L(x)$ and 0 for all other indices. If the tree is honest, we know in addition that the S_i are independent of the Y_i conditional on the X_i .

An important property of k -PNN predictors is that we can often get a good idea about whether S_i is non-zero even if we only get to see Z_i ; more formally, as we show below, the quantity $s \text{Var} [S_1 | Z_1]$ cannot get too small. Establishing this fact is a key step in showing that k -PNNs are incremental. In the following result, T can be an arbitrary symmetric k -PNN predictor.

Lemma 4. *Suppose that the features X are independently and identically distributed on $[0, 1]^d$ with a density f that is bounded away from infinity, and let T be any symmetric k -PNN predictor. Then, there is a constant $C_{f,d}$ depending only on f and d such that, as s gets large,*

$$s \text{Var} [\mathbb{E} [S_1 | Z_1]] \gtrsim \frac{1}{k} C_{f,d} / \log(s)^d, \quad (21)$$

where S_i is the indicator for whether the observation is selected in the subsample. When f is uniform over $[0, 1]^d$, the bound holds with $C_{f,d} = 2^{-(d+1)} (d-1)!$.

Note that $s \text{Var} [S_1] \sim k^{-1}$, so (21) can be interpreted as a lower bound on how much information Z_1 contains about the selection event S_1 . Thanks to this result, we are now ready to show that all honest and regular random-split trees are incremental. Notice that any symmetric k -regular tree following Definition 4 is also a symmetric k -PNN predictor.

Theorem 5. *Suppose that the conditions of Lemma 4 hold and that T is an honest k -regular symmetric tree in the sense of Definitions 2, 4, and 5. Suppose moreover that the conditional moments $\mu(x)$ and $\mu_2(x)$ are both Lipschitz continuous at x , and that $\mu_2(x)$ is uniformly bounded for all $x \in [0, 1]^d$. Finally, suppose that $\text{Var} [Y | X = x] > 0$. Then T is $\nu(s)$ -incremental at x with*

$$\nu(s) = C_{f,d} / \log(s)^d, \quad (22)$$

where $C_{f,d}$ is the constant from Lemma 4.

Finally, the result of Theorem 5 also holds for double-sample trees of the form described in Procedure 1. To prove the following result, we note that a double-sample tree is an honest, symmetric k -PNN predictor with respect to the \mathcal{I} -sample, while all the data in the \mathcal{J} -sample can be folded into the auxiliary noise term ξ . An immediate application of Theorem 5 then shows us that

$$\text{Var} [\hat{T}] \geq \sum_{i \in \mathcal{I}} \text{Var} [\mathbb{E} [T | Z_i]] \gtrsim \frac{C_{f,d}}{\log(|\mathcal{I}|/2)^d} \text{Var} [T], \quad (23)$$

where T is a double-sample tree built using both the \mathcal{I} and \mathcal{J} samples.

Corollary 6. *Under the conditions of Theorem 5, suppose that T is instead a double-sample tree (Procedure 1) that is honest, k -regular, and symmetric with respect to the \mathcal{I} -sample. Then, T is ν -incremental, with ν as defined in (22).*

3.3.2 Subsampling Incremental Base Learners

In the previous section, we showed that decision trees are ν -incremental, in that the Hájek projection \hat{T} of T preserves at least some of the variation of T . In this section, we show that randomly subsampling ν -incremental predictors makes them 1-incremental; this then lets us proceed with a classical statistical analysis. The following lemma, which flows directly from the ANOVA decomposition of Efron and Stein [1981], provides a first motivating result for our analysis.

Lemma 7. *Let $\hat{\mu}(x)$ be the estimate for $\mu(x)$ generated by a random forest with base learner T as defined in (12), and let $\mathring{\mu}$ be the Hájek projection of $\hat{\mu}$ (18). Then*

$$\mathbb{E} \left[\left(\hat{\mu}(x) - \mathring{\mu}(x) \right)^2 \right] \leq \left(\frac{s}{n} \right)^2 \text{Var} [T(x; \xi, Z_1, \dots, Z_s)]$$

whenever the variance $\text{Var} [T]$ of the base learner is finite.

This technical result paired with Theorem 5 or Corollary 6 leads to an asymptotic Gaussianity result; from a technical point of view, it suffices to check Lindeberg-style conditions for the central limit theorem.

Theorem 8. *Let $\hat{\mu}(x)$ be a random forest estimator trained according the conditions of Theorem 5 or Corollary 6, with Y restricted to a bounded interval $Y \in [-M, M]$. Suppose, moreover, that the subsample size s_n satisfies*

$$\lim_{n \rightarrow \infty} s_n = \infty \text{ and } \lim_{n \rightarrow \infty} s_n \log(n)^d / n = 0.$$

Then, there exists a sequence $\sigma_n(x) \rightarrow 0$ such that

$$\frac{\hat{\mu}_n(x) - \mathbb{E} [\hat{\mu}_n(x)]}{\sigma_n(x)} \Rightarrow \mathcal{N}(0, 1), \quad (24)$$

where $\mathcal{N}(0, 1)$ is the standard normal distribution.

Moreover, as we show below, it is possible to accurately estimate the variance of a random forest using the infinitesimal jackknife for random forests [Efron, 2014, Wager et al., 2014].

Theorem 9. *Let $\widehat{V}_{IJ}(x; Z_1, \dots, Z_n)$ be the infinitesimal jackknife for random forests as defined in (7). Then, under the conditions of Theorem 8,*

$$\widehat{V}_{IJ}(x; Z_1, \dots, Z_n) / \sigma_n^2(x) \rightarrow_p 1. \quad (25)$$

Finally, we end this section by motivating the finite sample correction $n(n-1)/(n-s)^2$ appearing in (13) by considering the simple case where we have trivial trees that do not make any splits: $T(x; \xi, Z_{i_1}, \dots, Z_{i_s}) = s^{-1} \sum_{j=1}^s Y_{i_j}$. In this case, we can verify that the full random forest is nothing but $\hat{\mu} = n^{-1} \sum_{i=1}^n Y_i$, and the standard variance estimator

$$\widehat{V}_{simple} = \frac{1}{n(n-1)} \sum_{i=1}^n (Y_i - \bar{Y})^2, \quad \bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i$$

is well-known to be unbiased for $\text{Var} [\hat{\mu}]$. We show below that, for trivial trees $\widehat{V}_{IJ} = \widehat{V}_{simple}$, implying that our correction makes \widehat{V}_{IJ} exactly unbiased in finite samples for trivial trees. Of course, $n(n-1)/(n-s)^2 \rightarrow 1$, and so Theorem 9 would hold even without this finite-sample correction; however, we find it to substantially improve the performance of our method in practice.

Proposition 10. For trivial trees $T(x; \xi, Z_{i_1}, \dots, Z_{i_s}) = s^{-1} \sum_{j=1}^s Y_{i_j}$, the variance estimate \widehat{V}_{IJ} (13) is equivalent to the standard variance estimator \widehat{V}_{simple} , and $\mathbb{E}[\widehat{V}_{IJ}] = \text{Var}[\hat{\mu}]$.

4 Inferring Heterogeneous Treatment Effects

We now return to our main topic, namely estimating heterogeneous treatment effects using random forests in the potential outcomes framework with unconfoundedness, and adapt our asymptotic theory for regression forests to the setting of causal inference. Here, we again work with training data consisting of tuples $Z_i = (X_i, Y_i, W_i)$ for $i = 1, \dots, n$, where X_i is a feature vector, Y_i is the response, and W_i is the treatment assignment. Our goal is to estimate the conditional average treatment effect $\tau(x) = \mathbb{E}[Y^{(1)} - Y^{(0)} | X = x]$ at a pre-specified test point x . By analogy to Definition 1, we build our causal forest CF by averaging estimates for τ obtained by training causal trees Γ over subsamples:

$$\text{CF}(x; Z_1, \dots, Z_n) = \binom{n}{s}^{-1} \sum_{1 \leq i_1 < i_2 < \dots < i_s \leq n} \mathbb{E}_{\xi \sim \Xi} [\Gamma(x; \xi, Z_{i_1}, \dots, Z_{i_s})]. \quad (26)$$

We seek an analogue to Theorem 1 for such causal forests.

Most of the definitions used to state Theorem 1 apply directly to the context; however, the notions of honesty and regularity need to be adapted slightly. Specifically, an honest causal tree is not allowed to look at the responses Y_i when making splits but can look at the treatment assignments W_i . Meanwhile, a regular causal tree must have at least k examples from both treatment classes in each leaf; in other words, regular causal trees seek to act as fully grown trees for the rare treatment assignment, while allowing for more instances of the common treatment assignment.

Definition 2b. A causal tree grown on a training sample $(Z_1 = (X_1, Y_1, W_1), \dots, Z_s = (X_s, Y_s, W_s))$ is *honest* if the tree does not use the responses Y_1, \dots, Y_s in choosing where to place its splits.

Definition 4b. A causal tree grown by recursive partitioning is *regular* at x if either: (a) (*standard case*) (1) Each split leaves at least a fraction $\alpha > 0$ of the available training examples on each side of the split, (2) The leaf containing x has at least k observations from each treatment group ($W_i \in \{0, 1\}$) for some $k \in \mathbb{N}$, and (3) The leaf containing x has either less than $2k - 1$ observations with $W_i = 0$ or $2k - 1$ observations with $W_i = 1$; or (b) (*double-sample case*) for a double-sample tree as defined in Procedure 1, (a) holds for the \mathcal{I} sample.

Given these assumptions, we show a close analogue to Theorem 1, given below. The main difference relative to our first result about regression forests is that we now rely on unconfoundedness and overlap to achieve consistent estimation of $\tau(x)$. To see how these assumptions enter the proof, recall that an honest causal tree uses the features X_i and the treatment assignments W_i in choosing where to place its splits, but not the responses Y_i . Writing $\mathcal{I}^{(1)}(x)$ and $\mathcal{I}^{(0)}(x)$ for the indices of the treatment and control units in the leaf

around x , we then find that after the splitting stage

$$\begin{aligned}
& \mathbb{E} [\Gamma(x) \mid X, W] && (27) \\
&= \frac{\sum_{\{i \in \mathcal{I}^{(1)}(x)\}} \mathbb{E} [Y^{(1)} \mid X = X_i, W = 1]}{|\mathcal{I}^{(1)}(x)|} - \frac{\sum_{\{i \in \mathcal{I}^{(0)}(x)\}} \mathbb{E} [Y^{(0)} \mid X = X_i, W = 0]}{|\mathcal{I}^{(0)}(x)|} \\
&= \frac{\sum_{\{i \in \mathcal{I}^{(1)}(x)\}} \mathbb{E} [Y^{(1)} \mid X = X_i]}{|\mathcal{I}^{(1)}(x)|} - \frac{\sum_{\{i \in \mathcal{I}^{(0)}(x)\}} \mathbb{E} [Y^{(0)} \mid X = X_i]}{|\mathcal{I}^{(0)}(x)|},
\end{aligned}$$

where the second equality follows by unconfoundedness (2). Thus, it suffices to show that the two above terms are consistent for estimating $\mathbb{E} [Y^{(0)} \mid X = x]$ and $\mathbb{E} [Y^{(1)} \mid X = x]$. To do so, we can essentially emulate the argument leading to Theorem 1, provided we can establish an analogue to Lemma 2 and give a fast enough decaying upper bound to the diameter of $L(x)$; this is where we need the overlap assumption. A formal proof of Theorem 11 is given in the appendix.

Theorem 11. *Suppose that we have n independent and identically distributed training examples $Z_i = (X_i, Y_i, W_i) \in [0, 1]^d \times [-M, M] \times \{0, 1\}$ for some $M > 0$. Suppose, moreover, that the treatment assignment is unconfounded (2) and has overlap (6). Finally, suppose that both potential outcome distributions $(X_i, Y_i^{(0)})$ and $(X_i, Y_i^{(1)})$ satisfy the same regularity assumptions as the pair (X_i, Y_i) did in the statement of Theorem 1. Given this data-generating process, let Γ be an honest, regular, and symmetric random-split causal forest in the sense of Definitions 2b, 3, 4b, and 5, and let $\hat{\tau}(x)$ be the estimate for $\tau(x)$ given by a causal forest with base learner Γ and a subsample size s_n scaling as in (14). Then, the predictions $\hat{\tau}(x)$ are consistent and asymptotically both Gaussian and centered, and the variance of the causal forest can be consistently estimated using the infinitesimal jackknife for random forests, i.e., (8) holds.*

Remark 3. (testing at many points) We note that it is not in general possible to construct causal trees that are regular in the sense of Definition 4b for all x simultaneously. As a simple example, consider the situation where $d = 1$, and $W_i = 1 (\{X_i \geq 0\})$; then, the tree can have at most 1 leaf for which it is regular. In the proof of Theorem 11, we avoided this issue by only considering a single test point x , as it is always possible to build a tree that is regular at a single given point x . In practice, if we want to build a causal tree that can be used to predict at many test points, we may need to assign different trees to be valid for different test points. Then, when predicting at a specific x , we treat the set of trees that were assigned to be valid at that x as the relevant forest and apply Theorem 11 to it.

5 Simulation Experiments

In observational studies, accurate estimation of heterogeneous treatment effects requires overcoming two potential sources of bias. First, we need to identify neighborhoods over which the actual treatment effect $\tau(x)$ is reasonably stable and, second, we need to make sure that we are not biased by varying sampling propensities $e(x)$. The simulations here aim to test the ability of causal forests to respond to both of these factors.

Since causal forests are adaptive nearest neighbor estimators, it is natural to use a non-adaptive nearest neighborhood method as our baseline. We compare our method to the

standard k nearest neighbors (k -NN) matching procedure, which estimates the treatment effect as

$$\hat{\tau}_{KNN}(x) = \frac{1}{k} \sum_{i \in \mathcal{S}_1(x)} Y_i - \frac{1}{k} \sum_{i \in \mathcal{S}_0(x)} Y_i, \quad (28)$$

where \mathcal{S}_1 and \mathcal{S}_0 are the k nearest neighbors to x in the treatment ($W = 1$) and control ($W = 0$) samples respectively. We generate confidence intervals for the k -NN method by modeling $\hat{\tau}_{KNN}(x)$ as Gaussian with mean $\tau(x)$ and variance $(\hat{V}(\mathcal{S}_0) + \hat{V}(\mathcal{S}_1))/(k(k-1))$, where $\hat{V}(\mathcal{S}_{0/1})$ is the sample variance for $\mathcal{S}_{0/1}$.

5.1 Experimental Setup

We describe our experiments in terms of the sample size n , the ambient dimension d , as well as the following functions:

$$\begin{aligned} \text{main effect: } m(x) &= 2^{-1} \mathbb{E} [Y^{(0)} + Y^{(1)} \mid X = x], \\ \text{treatment effect: } \tau(x) &= \mathbb{E} [Y^{(1)} - Y^{(0)} \mid X = x], \\ \text{treatment propensity: } e(x) &= \mathbb{P} [W = 1 \mid X = x]. \end{aligned}$$

In all our examples, we respect unconfoundedness (2), use $X \sim U([0, 1]^d)$, and have homoscedastic noise $Y^{(0/1)} \sim \mathcal{N}(\mathbb{E}[Y^{(0/1)} \mid X], 1)$. We evaluate performance in terms of expected mean-squared error for estimating $\tau(X)$ at a random test example X , as well as expected coverage of $\tau(X)$ with a target coverage rate of 0.95.

In our first experiment, we held the treatment effect fixed at $\tau(x) = 0$, and tested the ability of our method to resist bias due to an interaction between $e(x)$ and $m(x)$. This experiment is intended to emulate the problem that in observational studies, a treatment assignment is often correlated with potential outcomes, creating bias unless the statistical method accurately adjusts for covariates. k -NN matching is a popular approach for performing this adjustment in practice. Here, we set

$$e(X) = \frac{1}{4} (1 + \beta_{2,4}(X_1)), \quad m(X) = 2X_1 - 1, \quad (29)$$

where $\beta_{a,b}$ is the β -density with shape parameters a and b . We used $n = 500$ samples and varied d between 2 and 30. Since our goal is accurate propensity matching, we use propensity trees (Procedure 2) as our base learner; we grew $B = 1000$ trees with $s = 50$.

For our second experiment, we evaluated the ability of causal forests to adapt to heterogeneity in $\tau(x)$, while holding $m(x) = 0$ and $e(x) = 0.5$ fixed. Thanks to unconfoundedness, the fact that $e(x)$ is constant means that we are in a randomized experiment. We set τ to be a smooth function supported on the first two features:

$$\tau(X) = \varsigma(X_1) \varsigma(X_2), \quad \varsigma(x) = 1 + \frac{1}{1 + e^{-20(x-1/3)}}. \quad (30)$$

We took $n = 5000$ samples, while varying the ambient dimension d from 2 to 8. For causal forests, we used double-sample trees with the splitting rule of [Athey and Imbens \[2015\]](#) as our base learner (Procedure 1). We used $s = 2500$ (i.e., $|I| = 1250$) and grew $B = 2000$ trees.

One weakness of nearest neighbor approaches in general, and random forests in particular, is that they can fill the valleys and flatten the peaks of the true $\tau(x)$ function, especially

d	mean-squared error			coverage		
	CF	10-NN	100-NN	CF	10-NN	100-NN
2	0.02 (0)	0.21 (0)	0.09 (0)	0.95 (0)	0.93 (0)	0.62 (1)
5	0.02 (0)	0.24 (0)	0.12 (0)	0.94 (1)	0.92 (0)	0.52 (1)
10	0.02 (0)	0.28 (0)	0.12 (0)	0.94 (1)	0.91 (0)	0.51 (1)
15	0.02 (0)	0.31 (0)	0.13 (0)	0.91 (1)	0.90 (0)	0.48 (1)
20	0.02 (0)	0.32 (0)	0.13 (0)	0.88 (1)	0.89 (0)	0.49 (1)
30	0.02 (0)	0.33 (0)	0.13 (0)	0.85 (1)	0.89 (0)	0.48 (1)

Table 1: Comparison of the performance of a causal forests (CF) with that of the k -nearest neighbors (k -NN) estimator with $k = 10, 100$, on the setup (29). The numbers in parentheses indicate the (rounded) standard sampling error for the last printed digit, obtained by aggregating performance over 500 simulation replications.

near the edge of feature space. We demonstrate this effect using an example similar to the one studied above, except now $\tau(x)$ has a sharper spike in the $x_1, x_2 \approx 1$ region:

$$\tau(X) = \varsigma(X_1) \varsigma(X_2), \quad \varsigma(x) = \frac{2}{1 + e^{-12(x-1/2)}}. \quad (31)$$

We used the same training method as with (30), except with $n = 10000$, $s = 2000$, and $B = 10000$.

We implemented our simulations in R, using the package `randomForestCI` [Wager et al., 2014] for computing \hat{V}_{IJ} and the package `FNN` [Beygelzimer et al., 2013] for k -NN regression. An R package `causalForest` built on top of `rpart` [Therneau et al., 2014] is under development. All our trees had a minimum leaf size of $k = 1$.

5.2 Results

In our first setup (29), causal forests present a striking improvement over k -NN matching; see Table 1. Causal forests succeed in maintaining a mean-squared error of 0.02 as d grows from 2 to 30, while 10-NN and 100-NN do an order of magnitude worse. We note that the noise of k -NN due to variance in Y after conditioning on X and W is already $2/k$, implying that k -NN with $k \leq 100$ cannot hope to match the performance of causal forests. Here, however, 100-NN is overwhelmed by bias, even with $d = 2$. Meanwhile, in terms of uncertainty quantification, our method achieves nominal coverage up to $d = 10$, after which the performance of the confidence intervals starts to decay. The 10-NN method also achieves decent coverage, however, its confidence intervals are much wider than ours as evidenced by the mean-squared error.

In our second setup (30), causal forests present a similar improvement over k -NN matching when $d > 2$, as seen in Table 2.⁸ Unexpectedly, we find that the performance of causal forests *improves* with d , at least when d is small. To understand this phenomenon, we note that the variance of a forest depends on the product of the variance of individual trees times the correlation between different trees [Breiman, 2001a, Hastie et al., 2009]. Apparently, when d is larger, the individual trees have more flexibility in how to place their splits, thus reducing their correlation and decreasing the variance of the full ensemble.

⁸When $d = 2$, we do not expect causal forests to have a particular advantage over k -NN since the true τ also has 2-dimensional support; our results mirror this, as causal forests appear to have comparable performance to 50-NN.

d	mean-squared error			coverage		
	CF	7-NN	50-NN	CF	7-NN	50-NN
2	0.04 (0)	0.29 (0)	0.04 (0)	0.97 (0)	0.93 (0)	0.94 (0)
3	0.03 (0)	0.29 (0)	0.05 (0)	0.96 (0)	0.93 (0)	0.92 (0)
4	0.03 (0)	0.30 (0)	0.08 (0)	0.94 (0)	0.93 (0)	0.86 (1)
5	0.03 (0)	0.31 (0)	0.11 (0)	0.93 (1)	0.92 (0)	0.77 (1)
6	0.02 (0)	0.34 (0)	0.15 (0)	0.93 (1)	0.91 (0)	0.68 (1)
8	0.03 (0)	0.38 (0)	0.21 (0)	0.90 (1)	0.90 (0)	0.57 (1)

Table 2: Comparison of the performance of a causal forests (CF) with that of the k -nearest neighbors (k -NN) estimator with $k = 7, 50$, on the setup (30). The numbers in parentheses indicate the (rounded) standard sampling error for the last printed digit, obtained by aggregating performance over 25 simulation replications.

d	mean-squared error			coverage		
	CF	10-NN	100-NN	CF	10-NN	100-NN
2	0.02 (0)	0.20 (0)	0.02 (0)	0.94 (0)	0.93 (0)	0.94 (0)
3	0.02 (0)	0.20 (0)	0.03 (0)	0.90 (0)	0.93 (0)	0.90 (0)
4	0.02 (0)	0.21 (0)	0.06 (0)	0.84 (1)	0.93 (0)	0.78 (1)
5	0.02 (0)	0.22 (0)	0.09 (0)	0.81 (1)	0.93 (0)	0.67 (0)
6	0.02 (0)	0.24 (0)	0.15 (0)	0.79 (1)	0.92 (0)	0.58 (0)
8	0.03 (0)	0.29 (0)	0.26 (0)	0.73 (1)	0.90 (0)	0.45 (0)

Table 3: Comparison of the performance of a causal forests (CF) with that of the k -nearest neighbors (k -NN) estimator with $k = 10, 100$, on the setup (31). The numbers in parentheses indicate the (rounded) standard sampling error for the last printed digit, obtained by aggregating performance over 40 simulation replications.

Finally, in the setting (31), Table 3, causal forests still achieve an order of magnitude improvement over k -NN in terms of mean-squared error when $d > 2$, but struggle more in terms of coverage. This appears to largely be a bias effect: especially as d gets larger, the random forest is dominated by bias instead of variance and so the confidence intervals are not centered. Figure 1 illustrates this phenomenon: although the causal forest faithfully captures the qualitative aspects of the true τ -surface, it does not exactly match its shape, especially in the upper-right corner where τ is largest. Our theoretical results guarantee that this effect will go away as $n \rightarrow \infty$. Figure 1 also helps us understand why k -NN performs so poorly in terms of mean-squared error: its predictive surface is both badly biased and noticeably “grainy,” and it suffers from bias not only at the boundary where the treatment effect is largest, but also where the slope of the treatment effect is high in the interior.

These results highlight the promise of causal forests for accurate estimation of heterogeneous treatment effects, all while emphasizing avenues for further work. An immediate challenge is to control the bias of causal forests to achieve better coverage. Using more powerful splitting rules is a good way to reduce bias by enabling the trees to focus more closely on the coordinates with the greatest signal. The study of splitting rules for trees designed to estimate causal effects still is in its infancy and improvements may be possible.

There remains work to be done in understanding how to tune random forests. In our experiments, we manually chose whether to use a double-sample tree or a propensity tree;




Figure 1: The true treatment effect $\tau(X_i)$ at 10,000 random test examples X_i , along with estimates $\hat{\tau}(X_i)$ produced by a causal forest and optimally-tuned k -NN ($k^* = 33$), on data drawn according to (31) with $d = 6$. The test points are plotted according to their first two coordinates; the treatment effect is denoted by color, from dark (low) to light (high). On this simulation instance, causal forests had a mean-squared error of 0.029, in comparison to 0.199 for optimally-tuned k -NN.

an important challenge for future work is to design splitting rules that can automatically choose which characteristic of the training data to split on. A principled and automatic rule for choosing s would also be valuable; in Appendix A we present extensive simulations in the setting of Table 2 while varying both s and n .

6 Discussion

This paper proposed a class of non-parametric methods for heterogeneous treatment effect estimation that allow for data-driven feature selection all while maintaining the benefits of classical methods, i.e., asymptotically normal and unbiased point estimates with valid confidence intervals. Our causal forest estimator can be thought of as an adaptive nearest neighbor method, where the data determines which dimensions are most important to consider in selecting nearest neighbors. Such adaptivity seems essential for modern large-scale applications with many features.

In general, the challenge in using adaptive methods as the basis for valid statistical inference is that selection bias can be difficult to quantify; see Berk et al. [2013], Chernozhukov et al. [2015], Taylor and Tibshirani [2015], and references therein for recent advances. In this paper, pairing “honest” trees with the subsampling mechanism of random forests enabled us to accomplish this goal in a simple yet principled way. In our simulation experiments, our method provides dramatically better mean-squared error than classical methods while achieving nominal coverage rates in moderate sample sizes.

A number of important extensions and refinements are left open. Our current results only provide pointwise confidence intervals for $\tau(x)$; extending our theory to the setting of global functional estimation seems like a promising avenue for further work. Another challenge is that nearest-neighbor non-parametric estimators typically suffer from bias at the boundaries of the support of the feature space. A systematic approach to trimming at the boundaries, and possibly correcting for bias, would improve the coverage of the confidence intervals. In general, work can be done to identify methods that produce accurate variance estimates even

in more challenging circumstances, e.g., with small samples or a large number of covariates, or to identify when variance estimates are unlikely to be reliable.

Acknowledgment

We are grateful for helpful feedback from Brad Efron, Trevor Hastie, Guido Imbens, Guenther Walther, and two anonymous referees. S. W. was partially supported by a B. C. and E. J. Eaves Stanford Graduate Fellowship.

References

- Susan F Assmann, Stuart J Pocock, Laura E Enos, and Linda E Kasten. Subgroup analysis and other (mis) uses of baseline data in clinical trials. *The Lancet*, 355(9209):1064–1069, 2000.
- Susan Athey and Guido Imbens. Machine learning methods for estimating heterogeneous causal effects. *arXiv preprint arXiv:1504.01132*, 2015.
- Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, and Linda Zhao. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- Alina Beygelzimer and John Langford. The offset tree for learning with partial labels. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 129–138. ACM, 2009.
- Alina Beygelzimer, Sham Kakadet, John Langford, Sunil Arya, David Mount, and Shengqiao Li. *FNN: Fast Nearest Neighbor Search Algorithms and Applications*, 2013. URL <http://CRAN.R-project.org/package=FNN>. R package version 1.1.
- Debopam Bhattacharya and Pascaline Dupas. Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, 167(1):168–196, 2012.
- G erard Biau. Analysis of a random forests model. *The Journal of Machine Learning Research*, 13:1063–1095, 2012.
- G erard Biau, Luc Devroye, and G abor Lugosi. Consistency of random forests and other averaging classifiers. *The Journal of Machine Learning Research*, 9:2015–2033, 2008.
- G erard Biau, Fr ed eric C erou, and Arnaud Guyader. On the rate of convergence of the bagged nearest neighbor estimate. *The Journal of Machine Learning Research*, 11:687–712, 2010.
- Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008.
- Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001a.
- Leo Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199–231, 2001b.
- Leo Breiman. Consistency for a simple model of random forests. *Statistical Department, University of California at Berkeley. Technical Report*, (670), 2004.
- Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. *Classification and Regression Trees*. CRC press, 1984.
- Peter B uhlmann and Bin Yu. Analyzing bagging. *The Annals of Statistics*, 30(4):927–961, 2002.
- Andreas Buja and Werner Stuetzle. Observations on bagging. *Statistica Sinica*, 16(2):323, 2006.

- Song Xi Chen and Peter Hall. Effects of bagging and bias correction on estimators defined by estimating equations. *Statistica Sinica*, 13(1):97–110, 2003.
- Victor Chernozhukov, Christian Hansen, and Martin Spindler. Valid post-selection and post-regularization inference: An elementary, general approach. *Annual Review of Economics*, 7(1):649–688, 2015.
- Hugh A Chipman, Edward I George, and Robert E McCulloch. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.
- David I Cook, Val J Gebski, and Anthony C Keech. Subgroup analysis in clinical trials. *Medical Journal of Australia*, 180(6):289–292, 2004.
- Richard K Crump, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics*, 90(3):389–405, 2008.
- Rajeev H Dehejia. Program evaluation as a decision problem. *Journal of Econometrics*, 125(1):141–173, 2005.
- Misha Denil, David Matheson, and Nando De Freitas. Narrowing the gap: Random forests in theory and in practice. In *Proceedings of The 31st International Conference on Machine Learning*, pages 665–673, 2014.
- Jiangtao Duan. *Bootstrap-Based Variance Estimators for a Bagging Predictor*. PhD thesis, North Carolina State University, 2011.
- Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1097–1104, 2011.
- Bradley Efron. Estimation and accuracy after model selection (with discussion). *Journal of the American Statistical Association*, 109(507), 2014.
- Bradley Efron and Charles Stein. The jackknife estimate of variance. *The Annals of Statistics*, 9(3):586–596, 1981.
- Jared C Foster, Jeremy MG Taylor, and Stephen J Ruberg. Subgroup identification from randomized clinical trial data. *Statistics in medicine*, 30(24):2867–2880, 2011.
- Jerome H Friedman and Peter Hall. On bagging and nonlinear estimation. *Journal of Statistical Planning and Inference*, 137(3):669–683, 2007.
- Donald P Green and Holger L Kern. Modeling heterogeneous treatment effects in survey experiments with bayesian additive regression trees. *Public opinion quarterly*, 76(3):491–511, 2012.
- Jaroslav Hájek. Asymptotic normality of simple linear rank statistics under alternatives. *The Annals of Mathematical Statistics*, 39(2):325–346, 1968.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. New York: Springer, 2009.
- Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1), 2011.
- Keisuke Hirano and Jack R Porter. Asymptotics for statistical treatment rules. *Econometrica*, 77(5):1683–1701, 2009.
- Keisuke Hirano, Guido W Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948.
- Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.

- Kosuke Imai and Marc Ratkovic. Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443–470, 2013.
- Guido W Imbens and Donald B Rubin. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015.
- Louis A Jaeckel. *The Infinitesimal Jackknife*. 1972.
- Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. Prediction policy problems. *American Economic Review*, 105(5):491–95, 2015.
- Myoung-jae Lee. Non-parametric tests for distributional treatment effect for randomly censored responses. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(1):243–264, 2009.
- Andy Liaw and Matthew Wiener. Classification and regression by randomForest. *R News*, 2(3):18–22, 2002. URL <http://CRAN.R-project.org/doc/Rnews/>.
- Yi Lin and Yongho Jeon. Random forests and adaptive nearest neighbors. *Journal of the American Statistical Association*, 101(474):578–590, 2006.
- Charles F Manski. Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246, 2004.
- Daniel F McCaffrey, Greg Ridgeway, and Andrew R Morral. Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychological methods*, 9(4):403, 2004.
- Nicolai Meinshausen. Quantile regression forests. *The Journal of Machine Learning Research*, 7:983–999, 2006.
- Lucas Mentch and Giles Hooker. Ensemble trees and CLTs: Statistical inference for supervised learning. *arXiv preprint arXiv:1404.6473*, 2014.
- Dimitris N. Politis, Joseph P. Romano, and Michael Wolf. *Subsampling*. Springer Series in Statistics. Springer New York, 1999.
- Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- Richard J Samworth. Optimal weighted nearest neighbour classifiers. *The Annals of Statistics*, 40(5):2733–2763, 2012.
- Erwan Scornet, Gérard Biau, and Jean-Philippe Vert. Consistency of random forests. *The Annals of Statistics*, 43(4):1716–1741, 2015.
- Joseph Sexton and Petter Laake. Standard errors for bagged and random forest estimators. *Computational Statistics & Data Analysis*, 53(3):801–811, 2009.
- James Edward Signorovitch. *Identifying Informative Biological Markers in High-Dimensional Genomic Data and Clinical Trials*. PhD thesis, Harvard University, 2007.
- Carolin Strobl, Anne-Laure Boulesteix, Achim Zeileis, and Torsten Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics*, 8(1):25, 2007.
- Xiaogang Su, Chih-Ling Tsai, Hansheng Wang, David M Nickerson, and Bogong Li. Subgroup analysis via recursive partitioning. *The Journal of Machine Learning Research*, 10: 141–158, 2009.
- Matt Taddy, Matt Gardner, Liyun Chen, and David Draper. Heterogeneous treatment effects in digital experimentation. *arXiv preprint arXiv:1412.8563*, 2014.
- Jonathan Taylor and Robert J Tibshirani. Statistical learning and selective inference. *Proceedings of the National Academy of Sciences*, 112(25):7629–7634, 2015.

- Terry Therneau, Beth Atkinson, and Brian Ripley. *rpart: Recursive Partitioning and Regression Trees*, 2014. URL <http://CRAN.R-project.org/package=rpart>. R package version 4.1-8.
- Lu Tian, Ash A Alizadeh, Andrew J Gentles, and Robert Tibshirani. A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association*, 109(508):1517–1532, 2014.
- Aad W. Van der Vaart. *Asymptotic Statistics*. Number 3. Cambridge Univ Pr, 2000.
- Stefan Wager. Asymptotic theory for random forests. *arXiv preprint arXiv:1405.0352*, 2014.
- Stefan Wager and Guenther Walther. Uniform convergence of random forests via adaptive concentration. *arXiv preprint arXiv:1503.06388*, 2015.
- 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, 2014.
- Herbert I Weisberg and Victor P Pontes. Post hoc subgroups in clinical trials: Anathema or analytics? *Clinical Trials*, 12(4):357–364, 2015.
- Daniel Westreich, Justin Lessler, and Michele J Funk. Propensity score estimation: neural networks, support vector machines, decision trees (CART), and meta-classifiers as alternatives to logistic regression. *Journal of Clinical Epidemiology*, 63(8):826–833, 2010.
- Richard J Willke, Zhiyuan Zheng, Prasun Subedi, Rikard Althin, and C Daniel Mullins. From concepts, theory, and evidence of heterogeneity of treatment effects to methodological approaches: a primer. *BMC medical research methodology*, 12(1):185, 2012.
- Achim Zeileis, Torsten Hothorn, and Kurt Hornik. Model-based recursive partitioning. *Journal of Computational and Graphical Statistics*, 17(2):492–514, 2008.

n	d	s	MSE	coverage	variance	n	d	s	MSE	coverage	variance	n	d	s	MSE	coverage	variance
1000	2	100	0.10 (1)	0.48 (4)	0.02 (0)	2000	4	200	0.08 (1)	0.47 (5)	0.01 (0)	5000	6	500	0.03 (0)	0.75 (2)	0.01 (0)
1000	2	200	0.06 (1)	0.86 (2)	0.03 (0)	2000	4	400	0.03 (0)	0.90 (2)	0.02 (0)	5000	6	1000	0.02 (0)	0.86 (1)	0.01 (0)
1000	2	250	0.06 (0)	0.90 (1)	0.04 (0)	2000	4	500	0.03 (0)	0.93 (1)	0.03 (0)	5000	6	1250	0.02 (0)	0.88 (1)	0.02 (0)
1000	2	333	0.06 (0)	0.94 (1)	0.05 (0)	2000	4	667	0.04 (0)	0.95 (1)	0.04 (0)	5000	6	1667	0.03 (0)	0.91 (1)	0.02 (0)
1000	2	500	0.07 (1)	0.95 (1)	0.08 (0)	2000	4	1000	0.04 (0)	0.97 (1)	0.05 (0)	5000	6	2500	0.03 (0)	0.95 (1)	0.03 (0)
1000	2	667	0.08 (1)	0.92 (4)	0.09 (1)	2000	4	1333	0.05 (0)	0.98 (1)	0.07 (0)	5000	6	3333	0.03 (0)	0.96 (1)	0.04 (0)
1000	3	100	0.15 (1)	0.31 (5)	0.02 (0)	2000	5	200	0.14 (1)	0.26 (3)	0.01 (0)	5000	8	500	0.04 (0)	0.70 (1)	0.01 (0)
1000	3	200	0.08 (1)	0.77 (4)	0.03 (0)	2000	5	400	0.06 (0)	0.74 (3)	0.02 (0)	5000	8	1000	0.02 (0)	0.84 (1)	0.01 (0)
1000	3	250	0.07 (1)	0.89 (3)	0.04 (0)	2000	5	500	0.06 (0)	0.82 (2)	0.03 (0)	5000	8	1250	0.02 (0)	0.87 (1)	0.02 (0)
1000	3	333	0.06 (1)	0.92 (3)	0.05 (0)	2000	5	667	0.05 (0)	0.88 (1)	0.03 (0)	5000	8	1667	0.03 (0)	0.90 (1)	0.02 (0)
1000	3	500	0.07 (1)	0.95 (2)	0.07 (0)	2000	5	1000	0.05 (0)	0.93 (1)	0.05 (0)	5000	8	2500	0.03 (0)	0.93 (1)	0.03 (0)
1000	3	667	0.07 (1)	0.84 (6)	0.06 (1)	2000	5	1333	0.06 (0)	0.89 (2)	0.05 (0)	5000	8	3333	0.03 (0)	0.95 (1)	0.03 (0)
1000	4	100	0.23 (2)	0.18 (2)	0.02 (0)	2000	6	200	0.15 (1)	0.24 (3)	0.01 (0)	10000	2	1000	0.02 (0)	0.90 (1)	0.01 (0)
1000	4	200	0.12 (1)	0.59 (4)	0.03 (0)	2000	6	400	0.06 (1)	0.75 (3)	0.02 (0)	10000	2	2000	0.03 (0)	0.93 (1)	0.02 (0)
1000	4	250	0.10 (1)	0.72 (4)	0.04 (0)	2000	6	500	0.05 (1)	0.83 (2)	0.03 (0)	10000	2	2500	0.03 (0)	0.93 (1)	0.03 (0)
1000	4	333	0.09 (1)	0.85 (2)	0.05 (0)	2000	6	667	0.05 (0)	0.88 (1)	0.03 (0)	10000	2	3333	0.04 (0)	0.95 (1)	0.04 (0)
1000	4	500	0.08 (0)	0.93 (1)	0.07 (0)	2000	6	1000	0.05 (0)	0.92 (1)	0.05 (0)	10000	2	5000	0.05 (0)	0.95 (0)	0.05 (0)
1000	4	667	0.08 (0)	0.80 (6)	0.06 (1)	2000	6	1333	0.06 (0)	0.93 (1)	0.05 (0)	10000	2	6667	0.06 (0)	0.97 (0)	0.07 (0)
1000	5	100	0.25 (1)	0.17 (2)	0.02 (0)	2000	8	200	0.17 (2)	0.17 (3)	0.01 (0)	10000	3	1000	0.02 (0)	0.87 (2)	0.01 (0)
1000	5	200	0.12 (1)	0.61 (3)	0.03 (0)	2000	8	400	0.06 (1)	0.74 (4)	0.02 (0)	10000	3	2000	0.02 (0)	0.91 (1)	0.02 (0)
1000	5	250	0.10 (1)	0.75 (2)	0.04 (0)	2000	8	500	0.05 (1)	0.84 (2)	0.03 (0)	10000	3	2500	0.02 (0)	0.92 (1)	0.02 (0)
1000	5	333	0.07 (0)	0.90 (2)	0.05 (0)	2000	8	667	0.05 (1)	0.90 (2)	0.03 (0)	10000	3	3333	0.03 (0)	0.94 (0)	0.02 (0)
1000	5	500	0.06 (0)	0.95 (1)	0.07 (0)	2000	8	1000	0.05 (0)	0.93 (2)	0.04 (0)	10000	3	5000	0.03 (0)	0.95 (0)	0.04 (0)
1000	5	667	0.07 (1)	0.85 (4)	0.06 (1)	2000	8	1333	0.05 (0)	0.93 (2)	0.05 (0)	10000	3	6667	0.04 (0)	0.97 (0)	0.05 (0)
1000	6	100	0.31 (1)	0.11 (1)	0.01 (0)	5000	2	500	0.02 (0)	0.93 (1)	0.01 (0)	10000	4	1000	0.02 (0)	0.83 (2)	0.01 (0)
1000	6	200	0.15 (1)	0.46 (5)	0.03 (0)	5000	2	1000	0.03 (0)	0.93 (1)	0.02 (0)	10000	4	2000	0.02 (0)	0.89 (1)	0.01 (0)
1000	6	250	0.12 (1)	0.68 (5)	0.04 (0)	5000	2	1250	0.03 (0)	0.94 (1)	0.03 (0)	10000	4	2500	0.02 (0)	0.90 (1)	0.02 (0)
1000	6	333	0.08 (1)	0.86 (2)	0.05 (0)	5000	2	1667	0.04 (0)	0.95 (0)	0.04 (0)	10000	4	3333	0.02 (0)	0.93 (1)	0.02 (0)
1000	6	500	0.06 (1)	0.95 (1)	0.06 (0)	5000	2	2500	0.05 (0)	0.95 (0)	0.06 (0)	10000	4	5000	0.03 (0)	0.95 (0)	0.03 (0)
1000	6	667	0.06 (0)	0.75 (4)	0.04 (0)	5000	2	3333	0.06 (0)	0.97 (0)	0.08 (0)	10000	4	6667	0.03 (0)	0.97 (0)	0.04 (0)
1000	8	100	0.40 (2)	0.08 (1)	0.01 (0)	5000	3	500	0.02 (0)	0.86 (3)	0.01 (0)	10000	5	1000	0.02 (0)	0.78 (2)	0.01 (0)
1000	8	200	0.21 (2)	0.34 (3)	0.03 (0)	5000	3	1000	0.02 (0)	0.91 (2)	0.02 (0)	10000	5	2000	0.02 (0)	0.85 (1)	0.01 (0)
1000	8	250	0.17 (2)	0.51 (5)	0.04 (0)	5000	3	1250	0.03 (0)	0.93 (1)	0.02 (0)	10000	5	2500	0.02 (0)	0.87 (1)	0.01 (0)
1000	8	333	0.12 (1)	0.76 (4)	0.04 (0)	5000	3	1667	0.03 (0)	0.94 (1)	0.03 (0)	10000	5	3333	0.02 (0)	0.90 (1)	0.02 (0)
1000	8	500	0.08 (1)	0.91 (1)	0.06 (0)	5000	3	2500	0.04 (0)	0.96 (1)	0.04 (0)	10000	5	5000	0.03 (0)	0.93 (1)	0.02 (0)
1000	8	667	0.09 (1)	0.71 (5)	0.04 (1)	5000	3	3333	0.05 (0)	0.97 (0)	0.06 (0)	10000	5	6667	0.03 (0)	0.95 (1)	0.03 (0)
2000	2	200	0.04 (0)	0.82 (3)	0.02 (0)	5000	4	500	0.02 (0)	0.84 (2)	0.01 (0)	10000	6	1000	0.02 (0)	0.76 (2)	0.01 (0)
2000	2	400	0.04 (0)	0.93 (1)	0.03 (0)	5000	4	1000	0.03 (0)	0.89 (2)	0.02 (0)	10000	6	2000	0.02 (0)	0.83 (2)	0.01 (0)
2000	2	500	0.04 (0)	0.94 (1)	0.04 (0)	5000	4	1250	0.03 (0)	0.91 (2)	0.02 (0)	10000	6	2500	0.02 (0)	0.86 (1)	0.01 (0)
2000	2	667	0.04 (0)	0.96 (1)	0.05 (0)	5000	4	1667	0.03 (0)	0.92 (1)	0.03 (0)	10000	6	3333	0.02 (0)	0.89 (1)	0.02 (0)
2000	2	1000	0.06 (0)	0.96 (1)	0.07 (0)	5000	4	2500	0.04 (0)	0.95 (1)	0.04 (0)	10000	6	5000	0.02 (0)	0.93 (1)	0.02 (0)
2000	2	1333	0.07 (0)	0.97 (0)	0.09 (0)	5000	4	3333	0.04 (0)	0.96 (1)	0.05 (0)	10000	6	6667	0.03 (0)	0.94 (1)	0.03 (0)
2000	3	200	0.07 (1)	0.59 (4)	0.02 (0)	5000	5	500	0.02 (0)	0.80 (2)	0.01 (0)	10000	8	1000	0.02 (0)	0.71 (3)	0.01 (0)
2000	3	400	0.04 (0)	0.87 (2)	0.03 (0)	5000	5	1000	0.02 (0)	0.88 (2)	0.02 (0)	10000	8	2000	0.02 (0)	0.80 (1)	0.01 (0)
2000	3	500	0.04 (0)	0.90 (2)	0.03 (0)	5000	5	1250	0.02 (0)	0.90 (2)	0.02 (0)	10000	8	2500	0.02 (0)	0.83 (1)	0.01 (0)
2000	3	667	0.05 (0)	0.94 (1)	0.04 (0)	5000	5	1667	0.03 (0)	0.92 (1)	0.02 (0)	10000	8	3333	0.02 (0)	0.87 (1)	0.01 (0)
2000	3	1000	0.05 (0)	0.96 (1)	0.06 (0)	5000	5	2500	0.03 (0)	0.95 (1)	0.03 (0)	10000	8	5000	0.02 (0)	0.92 (1)	0.02 (0)
2000	3	1333	0.06 (0)	0.97 (0)	0.08 (0)	5000	5	3333	0.03 (0)	0.96 (1)	0.04 (0)	10000	8	6667	0.02 (0)	0.94 (1)	0.02 (0)

Table 4: Simulations for the generative model described in (30), while varying s and n . The results presented in Table 2 are a subset of these results. The numbers in parentheses indicate the (rounded) estimates for standard sampling error for the last printed digit, obtained by aggregating performance over 10 simulation replications. Mean-squared error (MSE) and coverage denote performance for estimating τ on a random test set; the variance column denotes the mean variance estimate obtained by the infinitesimal jackknife on the test set. Target coverage is 0.95. We always grew $B = n$ trees.

A Additional Simulations

In Table 4, we expand on the simulation results given in Table 2, and present results for data generated according to (30) all while varying s and n . As expected, we observe a bias-variance trade-off: the causal forest is more affected by bias when s is small relative to n , and by variance when s is larger relative to n . Reassuringly, we observe that our confidence intervals obtain close-to-nominal coverage when the mean-squared error matches the average variance estimate $\hat{\sigma}^2(X)$ generated by the infinitesimal jackknife, corroborating the hypothesis that failures in coverage mostly arise when the causal forest is bias- instead of variance-dominated.

B Proofs

Notation. Throughout the appendix we use the following notation to describe asymptotic scalings: $f(s) \sim g(s)$ means that $\lim_{s \rightarrow \infty} f(s)/g(s) = 1$, $f(s) \gtrsim g(s)$ means that $\liminf_{s \rightarrow \infty} f(s)/g(s) \geq 1$ and $f(s) \lesssim g(s)$ is analogous, $f(s) = \mathcal{O}(g(s))$ means that $f(s) \lesssim C g(s)$ for some $C > 0$, $f(s) = \Omega(g(s))$ means that $f(s) \gtrsim c g(s)$ for some $c > 0$, and finally $f(s) = o(g(s))$ means that $\limsup_{s \rightarrow \infty} f(s)/g(s) = 0$.

Proof of Theorem 1. Given the statements of Theorem 8 and Theorem 9, it only remains to show that (15) holds with $\mathbb{E}[\hat{\mu}_n(x)]$ replaced with $\mu(x)$. To do so, it suffices to show that $|\mathbb{E}[\hat{\mu}_n(x)] - \mu(x)|/\sigma_n(x) \rightarrow 0$; the rest follows from Slutsky's lemma. Now, recall that by Theorem 3,

$$|\mathbb{E}[\hat{\mu}_n(x)] - \mu(x)| = \mathcal{O}\left(n^{-\frac{\beta}{2}} \frac{\log((1-\alpha)^{-1})}{\pi^{-1}d \log(\alpha^{-1})}\right).$$

Meanwhile, from Theorem 5 and the proof of Theorem 8, we see that

$$\sigma_n^2(x) \gtrsim C_{f,d} \frac{s}{n} \frac{\text{Var}[T]}{\log(s)^d}.$$

By honesty of T , $\text{Var}[T] \gtrsim \text{Var}[Y | X = x] / |\{i : X_i \in L(x)\}| \geq \text{Var}[Y | X = x] / (2k)$, and so

$$\sigma_n^2(x) \gtrsim \frac{C_{f,d}}{2k} \frac{s}{n} \frac{\text{Var}[Y | X = x]}{\log(s)^d} = \Omega(n^{\beta-1-\varepsilon})$$

for any $\varepsilon > 0$. It follows that

$$\frac{|\mathbb{E}[\hat{\mu}_n(x)] - \mu(x)|}{\sigma_n(x)} = \mathcal{O}\left(n^{\frac{1}{2}} \left(1 + \varepsilon - \beta \left(1 + \frac{\log((1-\alpha)^{-1})}{\pi^{-1}d \log(\alpha^{-1})}\right)\right)\right).$$

The right-hand-side bound converges to 0 for some small enough $\varepsilon > 0$ provided that

$$\beta > \left(1 + \frac{\log((1-\alpha)^{-1})}{\pi^{-1}d \log(\alpha^{-1})}\right)^{-1} = 1 - \left(1 + \frac{d}{\pi} \frac{\log(\alpha^{-1})}{\log((1-\alpha)^{-1})}\right)^{-1} = \beta_{\min}.$$

□

B.1 Bounding the Bias of Regression Trees

Proof of Lemma 2. Let $c(x)$ be the number of splits leading to the leaf $L(x)$, and let $c_j(x)$ be the number of these splits along the j -th coordinate. By regularity, we know that $s\alpha^{c(x)} \leq 2k - 1$, and so $c(x) \geq \log(s/(2k - 1))/\log(\alpha^{-1})$. Thus, because the tree is a random split tree, $c_j(x)$ we have the following stochastic lower bound for $c_j(x)$:

$$c_j(x) \stackrel{d}{\geq} \text{Binom} \left(\frac{\log(s/(2k - 1))}{\log(\alpha^{-1})}, \frac{\pi}{d} \right). \quad (32)$$

By Chernoff's inequality, it follows that

$$\begin{aligned} \mathbb{P} \left[c_j(x) \geq \frac{\pi}{d} \frac{\log(s/(2k - 1))}{\log(\alpha^{-1})} (1 - \eta) \right] &\leq \exp \left[-\frac{\eta^2}{2} \frac{\log(s/(2k - 1))}{\pi^{-1} d \log(\alpha^{-1})} \right] \\ &= \left(\frac{s}{2k - 1} \right)^{-\frac{\eta^2}{2} \frac{1}{\pi^{-1} d \log(\alpha^{-1})}}. \end{aligned} \quad (33)$$

Meanwhile, again by regularity, we might expect that $\text{diam}_j(L(x))$ should be less than $(1 - \alpha)^{c_j(x)}$, at least for large n . This condition would hold directly if the regularity condition from Definition 4 were framed in terms of Lebesgue measure instead of the number of training examples in the leaf; our task is to show that it still holds approximately in our current setup.

Using the methods developed in Wager and Walther [2015], and in particular their Lemma 12, we can verify that, with high probability and simultaneously for all but the last $\mathcal{O}(\log \log n)$ parent nodes above $L(x)$, the number of training examples inside the node divided by n is within a factor $1 + o(1)$ of the Lebesgue measure of the node. From this, we conclude that, with high probability

$$\text{diam}_j(L(x)) \leq (1 - \alpha + o(1))^{(1+o(1))c_j(x)},$$

or, more prosaically, that

$$\text{diam}_j(L(x)) \leq (1 - \alpha)^{0.99 c_j(x)}.$$

Combining this result with the above Chernoff bound yields the desired inequality.

Finally, we note that with double-sample trees, all the “ s ” terms in the above argument need to be replaced by “ $s/2$ ”; this additional factor 2, however, does not affect the final result. \square

Proof of Theorem 3. We begin with two observations. First, by honesty,

$$\mathbb{E}[T(x; Z)] - \mathbb{E}[Y | X = x] = \mathbb{E}[\mathbb{E}[Y | X \in L(x)] - \mathbb{E}[Y | X = x]].$$

Second, by Lipschitz continuity of the conditional mean function,

$$|\mathbb{E}[Y | X \in L(x)] - \mathbb{E}[Y | X = x]| \leq C \text{diam}(L(x)),$$

where C is the Lipschitz constant. Thus, in order to bound the bias under both Lipschitz and honesty assumptions, it suffices to bound the average diameter of the leaf $L(x)$.

To do so, we start by plugging in $\eta = \sqrt{\log((1-\alpha)^{-1})}$ in the bound from Lemma 2. Thanks our assumption that $\alpha \leq 0.2$, we see that $\eta \leq 0.48$ and so $0.99 \cdot (1-\eta) \geq 0.51$; thus, a union bound gives us that, for large enough s ,

$$\mathbb{P} \left[\text{diam}(L(x)) \geq \sqrt{d} \left(\frac{s}{2k-1} \right)^{-0.51 \frac{\log((1-\alpha)^{-1})}{\log(\alpha^{-1})} \frac{\pi}{d}} \right] \leq d \left(\frac{s}{2k-1} \right)^{-\frac{1}{2} \frac{\log((1-\alpha)^{-1})}{\log(\alpha^{-1})} \frac{\pi}{d}}.$$

The Lipschitz assumption lets us bound the bias on the event that that $\text{diam}(L(x))$ satisfies this bound, while the boundedness of the response lets us bound the bias on the rare failure occurrence. We see that, for large s , the second term dominates and so, asymptotically

$$|\mathbb{E}[T(x; Z)] - \mathbb{E}[Y | X = x]| \lesssim 2M d \left(\frac{s}{2k-1} \right)^{-\frac{1}{2} \frac{\log((1-\alpha)^{-1})}{\log(\alpha^{-1})} \frac{\pi}{d}}.$$

Since a forest is just an average of trees, the above result also holds for $\hat{\mu}(x)$. \square

B.2 Bounding the Incrementality of Regression Trees

Proof of Lemma 4. First, we focus on the case where f is constant, i.e., the features X_i have a uniform distribution over $[0, 1]^d$. To begin, recall that the S_i denote selection weights

$$T(x; Z) = \sum_{i=1}^s S_i Y_i \quad \text{where} \quad S_i = \begin{cases} |\{i : X_i \in L(x; Z)\}|^{-1} & \text{if } X_i \in L(x; Z), \\ 0 & \text{else,} \end{cases}$$

where $L(x; Z)$ denotes the leaf containing x . We also define the quantities

$$P_i = 1(\{X_i \text{ is a } k\text{-PNN of } x\}).$$

Because T is a k -PNN predictor, $P_i = 0$ implies that $S_i = 0$, and, moreover, $|\{i : X_i \in L(x; Z)\}| \geq k$; thus, we can verify that

$$\mathbb{E}[S_1 | Z_1] \leq \frac{1}{k} \mathbb{E}[P_1 | Z_1].$$

The bulk of the proof involves showing that

$$\mathbb{P} \left[\mathbb{E}[P_1 | Z_1] \geq \frac{1}{s^2} \right] \lesssim k \frac{2^{d+1} \log(s)^d}{(d-1)!} \frac{1}{s}; \quad (34)$$

by the above argument, this immediately implies that

$$\mathbb{P} \left[\mathbb{E}[S_1 | Z_1] \geq \frac{1}{k s^2} \right] \lesssim k \frac{2^{d+1} \log(s)^d}{(d-1)!} \frac{1}{s}. \quad (35)$$

Now, by exchangeability of the indices i , we know that

$$\mathbb{E}[\mathbb{E}[S_1 | Z_1]] = \mathbb{E}[S_1] = \frac{1}{s} \mathbb{E} \left[\sum_{i=1}^s S_i \right] = \frac{1}{s},$$

moreover, we can verify that

$$\mathbb{P} \left[\mathbb{E} [S_1 | Z_1] \geq \frac{1}{ks^2} \right] \mathbb{E} \left[\mathbb{E} [S_1 | Z_1] \mid \mathbb{E} [S_1 | Z_1] \geq \frac{1}{ks^2} \right] \sim \frac{1}{s}.$$

By Jensen's inequality, we then see that

$$\begin{aligned} \mathbb{E} \left[\mathbb{E} [S_1 | Z_1]^2 \right] &\geq \mathbb{P} \left[\mathbb{E} [S_1 | Z_1] \geq \frac{1}{ks^2} \right] \mathbb{E} \left[\mathbb{E} [S_1 | Z_1] \mid \mathbb{E} [S_1 | Z_1] \geq \frac{1}{ks^2} \right]^2 \\ &\sim \frac{s^{-2}}{\mathbb{P} \left[\mathbb{E} [S_1 | Z_1] \geq \frac{1}{ks^2} \right]} \end{aligned}$$

which, paired with (35), implies that

$$\mathbb{E} \left[\mathbb{E} [S_1 | Z_1]^2 \right] \gtrsim \frac{(d-1)!}{2^{d+1} \log(s)^d} \frac{1}{ks}.$$

This is equivalent to (21) because $\mathbb{E} \left[\mathbb{E} [S_1 | Z_1]^2 \right] = 1/s^2$ is negligibly small.

We now return to establishing (34). Recall that X_1, \dots, X_s are independently and uniformly distributed over $[0, 1]^d$, and that we are trying to find points that are $(2k-1)$ -PNNs a prediction point x . For now, suppose that $x = 0$. We know that X_1 is a PNN of 0 if and only if there are at most $2k-2$ other points X_i such that $X_{ij} \leq X_{1j}$ for all $j = 1, \dots, d$ (because the X have a continuous density, there will almost surely be no ties). Thus,

$$\begin{aligned} \mathbb{E}_{x=0} [P_1 | Z_1] &= \mathbb{P} \left[\text{Binomial} \left(s-1; \prod_{j=1}^d X_{1j} \right) \leq 2k-2 \right] \\ &\leq \binom{s-1}{2k-2} \left(1 - \prod_{j=1}^d X_{1j} \right)^{s-k} \\ &\leq s^{2k-2} \left(1 - \prod_{j=1}^d X_{1j} \right)^{s-k}. \end{aligned} \tag{36}$$

We can check that $X_{1j} \stackrel{d}{=} e^{-E_j}$ where E_j is a standard exponential random variable, and so

$$\mathbb{E}_{x=0} [P_1 | Z_1] \leq s^{2k-2} \left(1 - \exp \left[- \sum_{j=1}^d E_j \right] \right)^{s-k},$$

where $A \stackrel{d}{\leq} B$ means that $B - A \geq 0$ almost surely. Thus,

$$\begin{aligned}
& \mathbb{P}_{x=0} \left[\mathbb{E} [P_1 | Z_1] \geq \frac{1}{s^2} \right] \\
& \leq \mathbb{P} \left[s^{(2k-2)} \left(1 - \exp \left[- \sum_{j=1}^d E_j \right] \right)^{s-k} \geq \frac{1}{s^2} \right] \\
& = \mathbb{P} \left[\exp \left[- \sum_{j=1}^d E_j \right] \leq 1 - \left(\frac{1}{s^{2k}} \right)^{\frac{1}{s-k}} \right] \\
& = \mathbb{P} \left[\sum_{j=1}^d E_j \geq -\log \left(1 - \exp \left[-2k \frac{\log(s)}{s-k} \right] \right) \right].
\end{aligned}$$

Notice that this quantity goes to zero as s gets large. The sum of d standard exponential random variables has a gamma distribution with shape d and scale 1, and

$$\mathbb{P} \left[\sum_{j=1}^d E_j \geq c \right] = \frac{\Gamma(d, c)}{(d-1)!},$$

where Γ is the upper incomplete gamma function. It is well known that

$$\lim_{c \rightarrow \infty} \frac{\Gamma(d, c)}{c^{d-1} e^{-c}} = 1,$$

and so

$$\begin{aligned}
& \mathbb{P}_{x=0} \left[\mathbb{E} [P_1 | Z_1] \geq \frac{1}{s^2} \right] \\
& \lesssim \frac{\left(-\log \left(1 - \exp \left[-2k \frac{\log(s)}{s-k} \right] \right) \right)^{d-1} \left(1 - \exp \left[-2k \frac{\log(s)}{s-k} \right] \right)}{(d-1)!}.
\end{aligned}$$

We can check that

$$1 - \exp \left[-2k \frac{\log(s)}{s-k} \right] \sim 2k \frac{\log(s)}{s},$$

letting us simplify the above expression to

$$\mathbb{P}_{x=0} \left[\mathbb{E} [P_1 | Z_1] \geq \frac{1}{s^2} \right] \lesssim \frac{2k}{(d-1)!} \frac{\log(s)^d}{s}. \tag{37}$$

We thus have obtained a tight expression for our quantity of interested for a prediction point at $x = 0$.

In the case $x \neq 0$, the ambient space around x can be divided into 2^d quadrants. In order to check whether X_i is a PNN, we only need to consider other points in the same quadrant, as no point in a different quadrant can prevent X_i from being a PNN. Now, index the quadrants by $l = 1, \dots, 2^d$, and let v_l be the volume of the l -th quadrant. By applying

(37) on the l -th quadrant alone, we see that the probability of $\mathbb{E} [P_1 | Z_1] \geq \frac{1}{s^2}$ given that X_1 is in the l -th quadrant is asymptotically bounded on the order of

$$\frac{k+1}{(d-1)!} \frac{\log(s)^d}{v_k s}.$$

Summing over all quadrants, we find that

$$\begin{aligned} \mathbb{P}_{x=0} \left[\mathbb{E} [P_1 | Z_1] \geq \frac{1}{s^2} \right] &\lesssim \sum_{\{k: v_k > 0\}} v_k \frac{2k}{(d-1)!} \frac{\log(s)^d}{v_k s} \\ &= |\{k : v_k > 0\}| \frac{2k}{(d-1)!} \frac{\log(s)^d}{s} \\ &\leq k \frac{2^{d+1}}{(d-1)!} \frac{\log(s)^d}{s}, \end{aligned}$$

thus establishing (34). Finally, to generalize to bounded densities f , we note that if $f(x) \leq C$ for all $x \in [0, 1]^d$, then

$$\mathbb{E}_{x=0} [P_1 | Z_1] \leq \mathbb{P} \left[\text{Binomial} \left(s-1; C \prod_{j=1}^d X_{1j} \right) \leq 2k-2 \right],$$

and the previous argument goes through. \square

Proof of Theorem 5. Our main task is to show that if T is a regular tree and $\text{Var} [Y | X = x] > 0$, then

$$\text{Var} [\mathbb{E} [T(x; Z) | Z_1]] = \text{Var} [\mathbb{E} [S_1 | Z_1]] \cdot \text{Var} [Y | X = x] (1 + o(1)) \quad (38)$$

Given this result, Lemma 4 then implies that

$$\text{Var} [\mathbb{E} [T(x; Z) | Z_1]] \gtrsim \frac{1}{k} \frac{\nu(s)}{s} \text{Var} [Y | X = x].$$

Moreover, by Theorem 3, we know that

$$\begin{aligned} \mathbb{E} [Y_i | X_i \in L(x; Z)] &\rightarrow_p \mathbb{E} [Y | X = x], \text{ and} \\ \mathbb{E} [Y_i^2 | X_i \in L(x; Z)] &\rightarrow_p \mathbb{E} [Y^2 | X = x], \end{aligned}$$

and so

$$k \text{Var} [T(x; Z)] \leq |\{i : X_i \in L(x; Z)\}| \cdot \text{Var} [T(x; Z)] \rightarrow_p \text{Var} [Y | X = x],$$

because k remains fixed while the leaf size gets smaller. Thus, we conclude that

$$\frac{\text{Var} [\mathring{T}(x; Z)]}{\text{Var} [T(x; Z)]} \gtrsim k \frac{s \text{Var} [\mathbb{E} [T(x; Z) | Z_1]]}{\text{Var} [Y | X = x]} \gtrsim \nu(s),$$

as claimed.

Now, in order to verify (38), we first recall that by Lemma 4

$$\text{Var} [\mathbb{E} [S_1 | Z_1]] = \Omega \left(\frac{1}{s \log(s)^d} \right). \quad (39)$$

Thus, any terms that decay faster than the right-hand-side rate can safely be ignored in establishing (38). We begin by verifying that we can take the leaf $L(x)$ containing x to have a small diameter $\text{diam}(L(x))$. Define the truncated tree predictor

$$T'(x; Z) = T(x; Z) \mathbf{1}(\{\text{diam}(L(x)) \leq s^{-\omega}\}), \text{ where } \omega = \frac{1}{2} \frac{\pi}{d} \frac{\log((1-\alpha)^{-1})}{\log(\alpha^{-1})},$$

and define similarly the truncated selection variables $S'_i = S_i \mathbf{1}(\{\text{diam}(L(x)) \leq s^{-\omega}\})$. Now, thanks to the ANOVA decomposition (3rd line), we see that

$$\begin{aligned} & \text{Var} [\mathbb{E} [T'(x; Z) | Z_1] - \mathbb{E} [T(x; Z) | Z_1]] \\ &= \text{Var} [\mathbb{E} [T(x; Z) \mathbf{1}(\{\text{diam}(L(x)) > s^{-\omega}\}) | Z_1]] \\ &\leq \frac{1}{s} \text{Var} [T(x; Z) \mathbf{1}(\{\text{diam}(L(x)) > s^{-\omega}\})] \\ &\leq \frac{M^2}{s} \mathbb{P} [\text{diam}(L(x)) > s^{-\omega}]. \end{aligned}$$

Thus, by Lemma 2, the variance of the difference between $\mathbb{E} [T' | Z_1]$ and $\mathbb{E} [T | Z_1]$ decays faster than the target rate (39), and so

$$\text{Var} [\mathbb{E} [T(x; Z) | Z_1]] \sim \text{Var} [\mathbb{E} [T'(x; Z) | Z_1]],$$

provided that T' satisfies (38), as we will see it does. By the same argument, we also note that

$$\text{Var} [\mathbb{E} [S'_1 | Z_1]] \sim \text{Var} [\mathbb{E} [S_1 | Z_1]].$$

We can now proceed to analyze T' instead of T . Because T' is symmetric, we know that

$$\begin{aligned} & \mathbb{E} [T'(x; Z) | Z_1] - \mu(x) \\ &= \mathbb{E} [S'_1 (Y_1 - \mu(x)) | Z_1] + \sum_{i=2}^n \mathbb{E} [S'_i (Y_i - \mu(x)) | Z_1] \\ &= \mathbb{E} [S'_1 (Y_1 - \mu(x)) | Z_1] + (s-1) \mathbb{E} [S'_2 (Y_2 - \mu(x)) | Z_1]. \end{aligned}$$

Noting that $\mathbb{E} [Y_1 | Z_1] = Y_1$ and that, by honesty, $\mathbb{E} [S'_1 | Z_1] = \mathbb{E} [S'_1 | X_1]$, we can re-write the first term as

$$\begin{aligned} & \text{Var} [\mathbb{E} [S'_1 (Y_1 - \mu(x)) | Z_1]] = \text{Var} [\mathbb{E} [S'_1 | X_1] (Y_1 - \mu(x))] \\ &= \mathbb{E} [\mathbb{E} [S'_1 | X_1]^2 \mathbb{E} [(Y_1 - \mu(x))^2 | X_1]] - \mathbb{E} [\mathbb{E} [S'_1 | X_1] \mathbb{E} [Y_1 - \mu(x) | X_1]]^2. \end{aligned} \quad (40)$$

Now, because the first two conditional moments of Y given X are Lipschitz, and since $\mathbb{E} [S'_1 | X_1]$ is 0 for $\|X_1 - x\|_2 > s^{-\omega}$ thanks to our truncating argument, we see that

$$\begin{aligned} & \mathbb{E} [\mathbb{E} [S'_1 | X_1]^2 \mathbb{E} [(Y_1 - \mu(x))^2 | X_1]] \sim \mathbb{E} [\mathbb{E} [S'_1 | X_1]^2] \text{Var} [Y | X = x] \\ & \sim \mathbb{E} [\mathbb{E} [S_1 | X_1]^2] \text{Var} [Y | X = x], \end{aligned}$$

which by the proof of Lemma 4 scales exactly like the quantity we were looking for in (38). Meanwhile, the second term in the expansion (40) is of order M^2/s^2 and thus negligible. Finally, we verify that

$$\begin{aligned}
& \text{Var} [(s-1) \mathbb{E} [S'_2 (Y_2 - \mu(x)) \mid Z_1]] \\
&= (s-1)^2 \text{Var} [\mathbb{E} [\mathbb{E} [S'_2 \mid X_2] \mathbb{E} [Y_2 - \mu(x) \mid X_2] \mid Z_1]] \\
&= (s-1)^2 \text{Var} [\mathbb{E} [\mathbb{E} [S'_2 \mid X_2] \mid Z_1]] \cdot \mathcal{O}(s^{-\omega}) \\
&= (s-1)^2 \text{Var} \left[\mathbb{E} \left[\frac{1 - \mathbb{E} [S_1 \mid Z_1]}{s-1} \mid Z_1 \right] \right] \cdot \mathcal{O}(s^{-\omega}) \\
&= \mathcal{O}(s^{-\omega} \text{Var} [\mathbb{E} [S_1 \mid X_1]]),
\end{aligned}$$

where $s^{-\omega}$ bounds the diameter of the leaf $L(x)$ and the $\mathcal{O}(\cdot)$ term hides the Lipschitz constant. \square

B.3 Properties of Subsampled Incremental Base Learners

The results presented in this section rely heavily on the Efron-Stein ANOVA decomposition, summarized here for convenience. Suppose we have any symmetric function $T : \Omega^n \rightarrow \mathbb{R}$, and suppose that Z_1, \dots, Z_n are independent and identically distributed on Ω such that $\text{Var} [T(Z_1, \dots, Z_n)] < \infty$. Then Efron and Stein [1981] show that there exist functions T_1, \dots, T_n such that

$$T(Z_1, \dots, Z_n) = \mathbb{E} [T] + \sum_{i=1}^n T_1(Z_i) + \sum_{i < j} T_2(Z_i, Z_j) + \dots + T_n(Z_1, \dots, Z_n), \quad (41)$$

and that all $2^n - 1$ random variables on the right side of the above expression are all mean-zero and uncorrelated. It immediately follows that

$$\text{Var} [T] = \sum_{k=1}^n \binom{n}{k} V_k, \text{ where } V_k = \text{Var} [T_k(Z_1, \dots, Z_k)]. \quad (42)$$

For our purposes, it is also useful to note that the Hájek projection \mathring{T} can be written as

$$\mathring{T}(Z_1, \dots, Z_n) = \mathbb{E} [T] + \sum_{i=1}^n T_1(Z_i), \text{ and } \text{Var} [\mathring{T}] = n V_1.$$

Thus, the ANOVA decomposition provides a convenient abstract framework for analyzing our quantities of interest.

Proof of Lemma 7. Applying the ANOVA decomposition to the individual trees T , we see that a random forest estimator $\hat{\mu}(x)$ of the form (12) can equivalently be written as

$$\begin{aligned}
\hat{\mu}(x; Z_1, \dots, Z_n) &= \mathbb{E} [T] + \binom{n}{s}^{-1} \left(\binom{n-1}{s-1} \sum_{i=1}^n T_1(Z_i) \right. \\
&\quad \left. + \binom{n-2}{s-2} \sum_{i < j} T_2(Z_i, Z_j) + \dots + \sum_{i_1 < \dots < i_s} T_s(Z_{i_1}, \dots, Z_{i_s}) \right).
\end{aligned}$$

The above formula holds because each training point Z_i appears in $\binom{n-1}{s-1}$ out of $\binom{n}{s}$ possible subsamples of size s , each pair (Z_i, Z_j) appears in $\binom{n-2}{s-2}$ subsets, etc.

Now, we can also show that the Hájek projection of $\hat{\mu}$ is

$$\mathring{\hat{\mu}}(x; Z_1, \dots, Z_n) = \mathbb{E}[T] + \frac{s}{n} \sum_{i=1}^n T_1(Z_i).$$

As with all projections [Van der Vaart, 2000],

$$\mathbb{E} \left[\left(\hat{\mu}(x) - \mathring{\hat{\mu}}(x) \right)^2 \right] = \text{Var} \left[\hat{\mu}(x) - \mathring{\hat{\mu}}(x) \right].$$

Recall that the $T_k(\cdot)$ are all pairwise uncorrelated. Thus, using the notation $s_k = s \cdot (s-1) \cdots (s-k)$ it follows that

$$\begin{aligned} \mathbb{E} \left[\left(\hat{\mu}(x) - \mathring{\hat{\mu}}(x) \right)^2 \right] &= \sum_{k=2}^s \left(\frac{s_k}{n_k} \right)^2 \binom{n}{k} V_k, \\ &= \sum_{k=2}^s \frac{s_k}{n_k} \binom{s}{k} V_k, \\ &\leq \frac{s_2}{n_2} \sum_{k=2}^s \binom{s}{k} V_k, \\ &\leq \frac{s_2}{n_2} \text{Var}[T], \end{aligned}$$

where on the last line we used (42). We recover the stated result by noticing that $s_2/n_2 \leq s^2/n^2$ for all $2 \leq s \leq n$. \square

Proof of Theorem 8. Using notation from the previous lemma, let $\sigma_n^2 = s^2/n V_1$ be the variance of $\mathring{\hat{\mu}}$. We know that

$$\sigma_n^2 = \frac{s^2}{n^2} n V_1 \leq \frac{s^2}{n^2} \text{Var}[T],$$

and so $\sigma_n \rightarrow 0$ as desired. Now, by Theorem 5 combined with Lemma 7, we find that

$$\begin{aligned} \frac{1}{\sigma_n^2} \mathbb{E} \left[\left(\hat{\mu}(x) - \mathring{\hat{\mu}}(x) \right)^2 \right] &\leq \left(\frac{s}{n} \right)^2 \frac{\text{Var}[T]}{\sigma_n^2} \\ &= \frac{s}{n} \text{Var}[T] / \text{Var}[\mathring{T}] \\ &\lesssim \frac{s \log(s)^d}{n C_{f,d}} \\ &\rightarrow 0 \end{aligned}$$

by hypothesis. Thus, by Slutsky's lemma, it suffices to show that (24) holds for the Hájek projection of the random forest $\mathring{\hat{\mu}}(x)$.

By our definition of σ_n , all we need to do is check that $\mathring{\hat{\mu}}$ is asymptotically normal. One way to do so is using the Lindeberg central limit theorem [e.g., Billingsley, 2008]. Since

$$\mathring{\hat{\mu}}(x) = \frac{s}{n} \sum_{i=1}^n (\mathbb{E}[T | Z_i] - \mathbb{E}[T])$$

is a sum of independent and identically distributed random variables, it suffices to show that for any $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} \left[(\mathbb{E} [T | Z_1] - \mathbb{E} [T])^2 \mid (\mathbb{E} [T | Z_1] - \mathbb{E} [T])^2 > \tau_n \right]}{\mathbb{E} \left[(\mathbb{E} [T | Z_1] - \mathbb{E} [T])^2 \right]} = 0, \quad (43)$$

where $\tau_n = \varepsilon n \text{Var} [\mathbb{E} [T | Z_1]]$.

We know from Theorem 5 that

$$s \text{Var} [\mathbb{E} [T | Z_1]] \gtrsim \frac{C_{f,d}}{\log(s)^d}.$$

Thus, we can check that

$$\varepsilon n \text{Var} [\mathbb{E} [T | Z_1]] \gtrsim \varepsilon \frac{n}{s \log(s)^d},$$

the right-hand-side of which converges to infinity by hypothesis as $n \rightarrow \infty$. But, we assumed that Y —and thus also T —is bounded, and so for large enough n the quantity $(\mathbb{E} [T | Z_1] - \mathbb{E} [T])^2$ can never be greater than $n/s \log(s)^d$. Thus, Lindeberg's condition holds and $\hat{\mu}(x)$ makes asymptotically normal predictions. \square

Proof of Theorem 9. Let F denote the distribution from which we drew Z_1, \dots, Z_n . Then, the variance σ_n^2 of the Hájek projection of $\hat{\mu}(x)$ is

$$\begin{aligned} \sigma_n^2 &= \sum_{i=1}^n (\mathbb{E}_{Z \sim F} [\hat{\mu}(x) | Z_i] - \mathbb{E}_{Z \sim F} [\hat{\mu}(x)])^2 \\ &= \frac{s^2}{n^2} \sum_{i=1}^n (\mathbb{E}_{Z \sim F} [T | Z_i] - \mathbb{E}_{Z \sim F} [T])^2, \end{aligned}$$

whereas we can check that the infinitesimal jackknife as defined in (7) is equal to

$$\widehat{V}_{IJ} = \frac{n-1}{n} \left(\frac{n}{n-s} \right)^2 \frac{s^2}{n^2} \sum_{i=1}^n (\mathbb{E}_{Z^* \subset \widehat{F}} [T | Z_1^* = Z_i] - \mathbb{E}_{Z^* \subset \widehat{F}} [T])^2,$$

where \widehat{F} is the empirical distribution on $\{Z_1, \dots, Z_n\}$. Recall that we are sampling the Z^* from \widehat{F} without replacement.

It is useful to write our expression of interest \widehat{V}_{IJ} using the Hájek projection \mathring{T} of T :

$$\begin{aligned} \widehat{V}_{IJ} &= \frac{n-1}{n} \left(\frac{n}{n-s} \right)^2 \frac{s^2}{n^2} \sum_{i=1}^n (A_i + R_i)^2, \text{ where} \\ A_i &= \mathbb{E}_{Z^* \subset \widehat{F}} [\mathring{T} | Z_1^* = Z_i] - \mathbb{E}_{Z^* \subset \widehat{F}} [\mathring{T}] \text{ and} \\ R_i &= \mathbb{E}_{Z^* \subset \widehat{F}} [T - \mathring{T} | Z_1^* = Z_i] - \mathbb{E}_{Z^* \subset \widehat{F}} [T - \mathring{T}]. \end{aligned}$$

As we show in Lemma 12, the main effects A_i give us σ_n^2 , in that

$$\frac{1}{\sigma_n^2} \frac{s^2}{n^2} \sum_{i=1}^n A_i^2 \rightarrow_p 1. \quad (44)$$

Meanwhile, Lemma 13 establishes that the B_i all satisfy

$$\mathbb{E} [R_i^2] \lesssim \frac{2}{n} \text{Var} [T(x; , Z_1, \dots, Z_s)], \quad (45)$$

and so

$$\begin{aligned} \mathbb{E} \left[\frac{s^2}{n^2} \sum_{i=1}^n R_i^2 \right] &\lesssim \frac{2s^2}{n^2} \text{Var} [T(x; , Z_1, \dots, Z_s)] \\ &\lesssim \frac{2s}{n} \log(n)^d \sigma_n^2. \end{aligned}$$

Because all terms are positive and $s \log(n)^d / n$ goes to zero by hypothesis, Markov's inequality implies that

$$\frac{1}{\sigma_n^2} \frac{s^2}{n^2} \sum_{i=1}^n R_i^2 \rightarrow_p 0.$$

Using Cauchy-Schwarz to bound the cross terms of the form $A_i R_i$, and noting that $\lim_{n \rightarrow \infty} n(n-1)/(n-s)^2 = 1$, we can thus conclude that $\widehat{V}_{IJ}/\sigma_n^2$ converges in probability to 1. \square

Lemma 12. *Under the conditions of Theorem 9, (44) holds.*

Proof. We can write

$$\begin{aligned} A_i &= \mathbb{E}_{Z^* \subset \widehat{F}} [\overset{\circ}{T} | Z_1^* = Z_i] - \mathbb{E}_{Z^* \subset \widehat{F}} [\overset{\circ}{T}] \\ &= \left(1 - \frac{s}{n}\right) T_1(Z_i) + \left(\frac{s-1}{n-1} - \frac{s}{n}\right) \sum_{j \neq i} T_1(Z_j), \end{aligned}$$

and so our sum of interest is asymptotically unbiased for σ_n^2 :

$$\begin{aligned} \mathbb{E} \left[\frac{n}{n-1} \frac{s^2}{n^2} \sum_{i=1}^n A_i^2 \right] &\sim \frac{s^2}{n} \mathbb{E} [T_1(Z)^2] \\ &= \frac{s}{n} \text{Var} [\overset{\circ}{T}(Z_1, \dots, Z_s)] \\ &= \sigma_n^2. \end{aligned}$$

Thus, it only remains to show that this sum has a coefficient of variation that converges to zero. Since the A_i are only weakly correlated, it suffices to show that

$$\frac{1}{\sqrt{n}} \frac{\text{Var} [A_i^2]}{\mathbb{E} [A_i^2]} \rightarrow_p 0.$$

Because $|Y_i| \leq M$, we must have $0 \leq A_i^2 \leq 4M^2$. In the worst case

$$\frac{\text{Var} [A_i^2]}{\mathbb{E} [A_i^2]} \leq \frac{\mathbb{E} [4M^2 A_i^2]}{\mathbb{E} [A_i^2]} = 4M^2,$$

and so coefficient of variation in fact converges to zero. \square

Lemma 13. *Under the conditions of Theorem 9, (45) holds.*

Proof. Without loss of generality, we establish (45) for R_1 . Using the ANOVA decomposition (41), we can write our term of interest as

$$\begin{aligned}
R_1 &= \mathbb{E}_{Z^* \subset \hat{F}} [T - \mathring{T} \mid Z_1^* = Z_1] - \mathbb{E}_{Z^* \subset \hat{F}} [T - \mathring{T}] \\
&= \left(\frac{s-1}{n-1} - \binom{s}{2} / \binom{n}{2} \right) \sum_{i=2}^n T_2(Z_1, Z_i) \\
&\quad + \left(\binom{s-1}{2} / \binom{n-1}{2} - \binom{s}{2} / \binom{n}{2} \right) \sum_{2 \leq i < j \leq n} T_2(Z_i, Z_j) \\
&+ \left(\binom{s-1}{2} / \binom{n-1}{2} - \binom{s}{3} / \binom{n}{3} \right) \sum_{2 \leq i < j \leq n} T_3(Z_1, Z_i, Z_j) \\
&\quad + \left(\binom{s-1}{3} / \binom{n-1}{3} - \binom{s}{3} / \binom{n}{3} \right) \sum_{2 \leq i < j < k \leq n} T_3(Z_i, Z_j, Z_k) \\
&+ \dots
\end{aligned}$$

Because all the terms in the ANOVA expansion are mean-zero and uncorrelated, we see using notation from (42) that

$$\begin{aligned}
\mathbb{E} [R_1^2] &= (n-1) \left(\frac{s-1}{n-1} - \binom{s}{2} / \binom{n}{2} \right)^2 V_2 \\
&\quad + \binom{n-1}{2} \left(\binom{s-1}{2} / \binom{n-1}{2} - \binom{s}{2} / \binom{n}{2} \right)^2 V_2 \\
&+ \binom{n-1}{2} \left(\binom{s-1}{2} / \binom{n-1}{2} - \binom{s}{3} / \binom{n}{3} \right)^2 V_3 \\
&\quad + \binom{n-1}{3} \left(\binom{s-1}{3} / \binom{n-1}{3} - \binom{s}{3} / \binom{n}{3} \right)^2 V_3 \\
&+ \dots
\end{aligned}$$

Recall that

$$\sum_{k=1}^s \binom{s}{k} V_k = \text{Var} [T(Z_1, \dots, Z_s)].$$

The above sum is maximized when all the variance is contained in second-order terms, and $\binom{s}{2} V_2 = \text{Var} [T]$. This implies that

$$\begin{aligned}
\mathbb{E} [R_1^2] &\lesssim (n-1) \left(\frac{s-1}{n-1} - \binom{s}{2} / \binom{n}{2} \right)^2 \binom{s}{2}^{-1} \text{Var} [T(Z_1, \dots, Z_s)] \\
&\sim \frac{2}{n} \text{Var} [T(Z_1, \dots, Z_s)],
\end{aligned}$$

thus completing the proof. \square

Proof of Proposition 10. Let N_i^* denote whether or not the i -training example was used for a subsample, as in (13). For trivial trees

$$T(x; \xi, Z_{i_1}, \dots, Z_{i_s}) = \frac{1}{s} \sum_{j=1}^s Y_{i_j}$$

we can verify that for any $i = 1, \dots, n$, $\mathbb{E}_* [\hat{\mu}^*] \mathbb{E} [N_i^*] = s/n \bar{Y}$,

$$\begin{aligned} \mathbb{E}_* [\hat{\mu}^* N_i^*] &= \frac{s}{n} \left(\frac{Y_i}{s} + \frac{s-1}{s} \frac{n\bar{Y} - Y_i}{n-1} \right) = \frac{1}{n} \frac{n-s}{n-1} Y_i + \frac{s-1}{n-1} \bar{Y}, \text{ and} \\ \text{Cov}_* [\hat{\mu}^*, N_i^*] &= \frac{1}{n} \frac{n-s}{n-1} Y_i + \left(\frac{s-1}{n-1} - \frac{s}{n} \right) \bar{Y} = \frac{1}{n-1} \frac{n-s}{n} (Y_i - \bar{Y}). \end{aligned}$$

Thus, we find that

$$\begin{aligned} \hat{V}_{IJ} &= \frac{n-1}{n} \left(\frac{n}{n-s} \right)^2 \sum_{i=1}^n \text{Cov}_* [\hat{\mu}^*, N_i^*]^2 \\ &= \frac{1}{n(n-1)} \sum_{i=1}^n (Y_i - \bar{Y})^2 \\ &= \hat{V}_{\text{simple}}, \end{aligned}$$

as we sought to verify. \square

B.4 Extension to Causal Forests

Proof of Theorem 11. Our argument mirrors the proof of Theorem 1. The main steps involve bounding the bias of causal forests with an analogue to Theorem 3 and their incrementality using an analogue to Theorem 5. In general, we find that the same arguments as used with regression forests go through, but the constants in the results get worse by a factor ε depending on the amount of overlap (6). Given these results, the subsampling-based argument from Section 3.3.2 can be reproduced almost verbatim, and the final proof of Theorem 11 is identical to that of Theorem 1 presented at the beginning of Appendix B.

Bias. Under the conditions of Lemma 2, suppose that $\mathbb{E} [Y^{(0)} | X = x]$ and $\mathbb{E} [Y^{(1)} | X = x]$ are Lipschitz continuous, that $|Y^{(0)}|, |Y^{(1)}| \leq M$, that the trees Γ comprising the random forest are honest, and, moreover, that the overlap condition (6) holds for some $\varepsilon > 0$. Then, provided that $\alpha \leq 0.2$, the bias of the random forest at x is bounded by

$$|\mathbb{E} [\hat{\tau}(x)] - \tau(x)| \lesssim 2M d \left(\frac{\varepsilon s}{2k-1} \right)^{-\frac{1}{2} \frac{\log((1-\alpha)^{-1})}{\log(\alpha^{-1})} \frac{\pi}{d}}.$$

To establish this claim, we first seek with an analogue to Lemma 2, except now s in (32) is replaced by s_{\min} , i.e., the minimum of the number of cases (i.e., observations with $W_i = 1$) or controls (i.e., observations with $W_i = 0$) in the sample. A straight-forward computation then shows that $s_{\min}/s \gtrsim \varepsilon$, and that a variant of (33) where we replace s with εs still holds for large s . Next, to bound the bias itself, we start by applying unconfoundedness as in (27); then, the argument of Theorem 3 goes through without modifications, provided we replace every instance of “ s ” with “ εs ”.

Incrementality. Suppose that the conditions of Lemma 4 hold and that Γ is an honest k -regular causal tree in the sense of Definitions 2b and 4b. Suppose moreover that $\mathbb{E}[Y^{(0/1)} | X = x]$ and $\text{Var}[Y^{(0/1)} | X = x]$ are all Lipschitz continuous at x , that $|Y^{(0/1)}| \leq M$ almost surely, and that $\text{Var}[Y | X = x] > 0$. Suppose, finally, that the overlap condition (6) holds with $\varepsilon > 0$. Then T is $\nu(s)$ -incremental at x with

$$\nu(s) = \varepsilon C_{f,d} / \log(s)^d,$$

where $C_{f,d}$ is the constant from Lemma 4.

To prove this claim, we again focus on the case where $f(x) = 1$, in which case we use $C_{f,d} = 2^{-(d+1)}(d-1)!$. We begin by setting up notation as in the proof of Lemma 4. We write our causal tree as $\Gamma(x; Z) = \sum_{i=1}^s S_i Y_i$, where

$$S_i = \begin{cases} |\{i : X_i \in L(x; Z)\}, W_i = 1|^{-1} & \text{if } X_i \in L(x; Z) \text{ and } W_i = 1, \\ -|\{i : X_i \in L(x; Z)\}, W_i = 0|^{-1} & \text{if } X_i \in L(x; Z) \text{ and } W_i = 0, \\ 0 & \text{else,} \end{cases}$$

where $L(x; Z)$ denotes the leaf containing x , and let

$$P_i^W = 1(\{X_i \text{ is a } k\text{-PNN of } x \text{ among points with treatment status } W_i\}).$$

Finally, in a break from Lemma 4, define $w_{\min}(x; Z)$ as the minority class within the leaf $L(x; Z)$; more formally,

$$w_{\min} = 1(\{|\{i : X_i \in L(x; Z)\}, W_i = 1| \leq |\{i : X_i \in L(x; Z)\}, W_i = 0|\}).$$

By regularity of Γ , we know that the leaf $L(x; Z)$ can contain at most $2k-1$ examples from its minority class, and so $P_i^W = 0$ and $W = w_{\min}$ together imply that $S_i = 0$. Thus, we can verify that

$$\mathbb{E}[|S_1| 1(\{W_1 = w_{\min}\}) | Z_1] \leq \frac{1}{k} \mathbb{E}[P_1^W | Z_1].$$

We are now ready to use the same machinery as before. The random variables P_1^W now satisfy

$$\mathbb{P}\left[\mathbb{E}[P_1^W | Z_1] \geq \frac{1}{s^2 \mathbb{P}[W = W_1]^2}\right] \lesssim k \frac{2^{d+1} \log(s)^d}{(d-1)!} \frac{1}{s \mathbb{P}[W = W_1]};$$

by the above argument and ε -overlap (6), this immediately implies that

$$\mathbb{P}\left[\mathbb{E}[|S_1| 1(\{W_1 = w_{\min}\}) | Z_1] \geq \frac{1}{k \varepsilon^2 s^2}\right] \lesssim k \frac{2^{d+1} \log(s)^d}{(d-1)!} \frac{1}{\varepsilon s}.$$

By construction, we know that

$$\mathbb{E}[\mathbb{E}[|S_1| 1(\{W_1 = w_{\min}\}) | Z_1]] = \mathbb{E}[|S_1| 1(\{W_1 = w_{\min}\})] = \frac{1}{s},$$

which by the same argument as before implies that

$$\mathbb{E}\left[\mathbb{E}[|S_1| 1(\{W_1 = w_{\min}\}) | Z_1]^2\right] \gtrsim \frac{(d-1)!}{2^{d+1} \log(s)^d} \frac{\varepsilon}{k s}.$$

By monotonicity, we then conclude that

$$\mathbb{E} \left[\mathbb{E} [S_1 | Z_1]^2 \right] = \mathbb{E} \left[\mathbb{E} [|S_1| | Z_1]^2 \right] \gtrsim \frac{(d-1)!}{2^{d+1} \log(s)^d} \frac{\varepsilon}{k s}.$$

The second part of the proof revolves around showing that

$$\text{Var} \left[\mathbb{E} [\Gamma(x; Z) | Z_1] \right] \sim \mathbb{E} \left[\mathbb{E} [S_1 | Z_1]^2 \right] \cdot \text{Var} [Y | X = x],$$

and is a straight-forward adaptation of the proof of Theorem 5. □