

FISHER–SCHULTZ LECTURE:
GENERIC MACHINE LEARNING INFERENCE ON HETEROGENEOUS
TREATMENT EFFECTS IN RANDOMIZED EXPERIMENTS, WITH AN
APPLICATION TO IMMUNIZATION IN INDIA

VICTOR CHERNOZHUKOV

Department of Economics and Center for Statistics and Data Sciences, MIT

MERT DEMIRER

Sloan School of Management, MIT

ESTHER DUFLO

Department of Economics, MIT

IVÁN FERNÁNDEZ-VAL

Department of Economics, Boston University

We propose strategies to estimate and make inference on key features of heterogeneous effects in randomized experiments. These key features include *best linear predictors of the effects* using machine learning proxies, *average effects sorted by impact groups*, and *average characteristics of most and least impacted units*. The approach is valid in high-dimensional settings, where the effects are proxied (but not necessarily consistently estimated) by predictive and causal machine learning methods. We post-process these proxies into estimates of the key features. Our approach is generic; it can be used in conjunction with penalized methods, neural networks, random forests, boosted trees, and ensemble methods, both predictive and causal. Estimation and inference are based on repeated data splitting to avoid overfitting and achieve validity. We use quantile aggregation of the results across many potential splits, in particular taking medians of p -values and medians and other quantiles of confidence intervals. We show that quantile aggregation lowers estimation risks over a single split procedure, and establish its principal inferential properties. Finally, our analysis reveals ways to build provably better machine learning proxies through causal learning: we can use the objective functions that we develop to construct the best linear predictors of the effects, to obtain better machine learning proxies in the initial step. We illustrate the use of both inferential tools and causal learners with a randomized field experiment that evaluates a combination of nudges to stimulate demand for immunization in India.

Victor Chernozhukov: vchern@mit.edu

Mert Demirer: mdemirer@mit.edu

Esther Duflo: eduflo@mit.edu

Iván Fernández-Val: ivanf@bu.edu

This paper was delivered (virtually) by Esther Duflo at the Fisher–Schultz Lecture of the Econometric Society World Congress, 2020. We thank three anonymous referees, Susan Athey, Moshe Buchinsky, Denis Chetverikov, Carlos Cineli, Matt Hong, Stella Hong, Steven Lehrer, Siyi Luo, Max Kasy, Sylvia Klosin, Susan Murphy, Whitney Newey, Patrick Power, Victor Quintas-Martinez, Suhas Vijaykumar, and seminar participants at ASSA 2018, Barcelona GSE Summer Forum 2019, Brazilian Econometric Society Meeting 2019, BU, Lancaster, NBER summer institute 2018, NYU, UCLA, Whitney Newey’s Contributions to Econometrics conference, and York for valuable comments. Anirudh Sankar provided us with excellent research assistance. We gratefully acknowledge research support from the National Science Foundation, AFD, USAID, and 3ie. An R package that implements the methods in this paper, GenericML, is available on GitHub at <https://github.com/mwetz/GenericML>. Funding Statement: Open Access funding enabled and organized by MIT Hybrid 2025. WOA Institution: Massachusetts Institute of Technology. Consortia Name: MIT Hybrid 2025.

© 2025 The Authors. *Econometrica* published by John Wiley & Sons Ltd on behalf of The Econometric Society. Esther Duflo is the corresponding author on this paper. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

KEYWORDS: Agnostic inference, causal machine learning, confidence intervals, quantification of uncertainty, multiple sample splitting, heterogeneous effects, immunization incentives, nudges.

1. INTRODUCTION

RANDOMIZED CONTROLLED TRIALS (RCT) and Machine Learning (ML) are arguably two of the most important developments in data analysis methods for applied researchers. RCTs play an important role in the evaluation of social and economic programs, medical treatments, and marketing (e.g., [Duflo, Glennerster, and Kremer \(2007\)](#), [Imbens and Rubin \(2015\)](#)). ML is a name attached to a variety of constantly evolving statistical learning methods including Random Forest, Boosted Trees, Neural Networks, Penalized Regression, Ensembles, and Hybrids; see, for example, [Wasserman \(2016\)](#), [Friedman, Hastie, and Tibshirani \(2001\)](#), [Bishop and Nasrabadi \(2006\)](#), [Murphy \(2012\)](#), [Hastie, Tibshirani, and Wainwright \(2015\)](#), [Goodfellow, Bengio, and Courville \(2016\)](#), and [James, Witten, Hastie, and Tibshirani \(2021\)](#) for prominent textbook treatments. ML has become a key tool for prediction and pattern recognition problems, surpassing classical methods in high-dimensional settings.

At first blush, those two sets of methods may seem to have very different applications: in the most basic randomized controlled experiment, there is a sample with a single treatment and a single outcome. Covariates are not necessary and even linear regression is not the best way to analyze the data ([Freedman \(2008\)](#), [Imbens and Rubin \(2015\)](#)). In practice, however, applied researchers are often confronted with more complex experiments. For example, there might be accidental imbalances in the sample, which require selecting control variables in a principled way. ML tools, such as the lasso method proposed in [Belloni, Chernozhukov, and Hansen \(2014\)](#), [Belloni, Chernozhukov, Fernández-Val, and Hansen \(2017\)](#) or the double machine learning method proposed in [Chernozhukov et al. \(2017\)](#), have proven useful for this purpose. Moreover, some complex RCT designs have so many treatment combinations that ML methods may be useful to select the few treatments that actually work and pool the rest with the control groups for statistical power ([Banerjee et al. \(2019\)](#)). Finally, researchers and policy makers are often interested in features of the impact of the treatment that go beyond the simple average treatment effect. In particular, very often, they want to know whether the treatment effect depends on covariates, such as gender, age, etc. This heterogeneity is essential to assess if the impact of the program would generalize to a population with different characteristics, and, for economists, to better understand the driving mechanism behind the effects of a particular program. In a review of 189 RCTs published in top economic journals since 2006, we found that 76 (40%) report at least one subgroup analysis, wherein they report treatment effects in subgroups formed by baseline covariates.¹

One issue with reporting treatment effects split by subgroups, however, is that there might be a large number of potential ways to form subgroups. Often researchers collect rich baseline surveys, which give them access to a large number of covariates: choosing subgroups *ex post* opens the possibility of overfitting. To solve this problem, medical journals and the FDA require pre-registering the subsample of interest in medical trials *in advance*. In economics, this approach has gained some traction with the adoption of

¹The papers were published in *Quarterly Journal of Economics*, *American Economic Review*, *Review of Economics Studies*, *Econometrica*, and *Journal of Political Economy*. We thank Karthik Mularidharan, Mauricio Romero, and Kaspar Wüthrich for sharing the list of papers they computed for another project.

pre-analysis plans, which can be filed in the AEA registry for randomized experiments. However, restricting the heterogeneity analysis to pre-registered subgroups amounts to throwing away a large amount of potentially valuable information, especially now that many researchers collect large baseline data sets. It should be possible to use the data to discover *ex post* whether there is any relevant heterogeneity in treatment effect by covariates.

To do this in a disciplined fashion and avoid the risk of overfitting, scholars have recently proposed using ML tools. Indeed, ML tools seem ideal for exploring heterogeneity of treatment effects when researchers have access to a potentially large array of baseline variables to form subgroups and few guiding principles on which of those are likely to be relevant. Several recent papers, which we review below, develop methods for detecting heterogeneity in treatment effects. Empirical researchers have taken notice.²

This paper develops a generic approach to using any of the available ML tools to predict and make inference on heterogeneous treatment or policy effects. A core difficulty of applying ML tools to the estimation of heterogeneous causal effects is that, while they are successful in prediction empirically, it is much more difficult to obtain uniformly valid inference, that is, inference that remains valid under a large class of data generating processes. In fact, in high-dimensional settings, absent strong assumptions, generic ML tools may not even produce consistent estimators of the *conditional average treatment effect* (CATE), the difference in the expected potential outcomes between treated and control states conditional on covariates. Previous attempts to solve this problem focused either on specific tools (e.g., the method proposed by [Athey and Imbens \(2016\)](#), which has become popular with applied researchers, and uses trees), or on situations where those assumptions might be satisfied. Our approach to resolving the fundamental impossibilities in non-parametric inference is different. Motivated by [Genovese and Wasserman \(2008\)](#), instead of attempting to get consistent estimation and uniformly valid inference on the CATE itself, we focus on providing valid estimation and inference on *features* of CATE.

We start by building an ML proxy predictor of CATE, and then target features of the CATE based on this proxy predictor. In particular, we consider three objects, which are likely to be of interest to applied researchers and policy makers: (1) *Best Linear Predictor* (BLP) of the CATE on the ML proxy predictor; (2) *Sorted Group Average Treatment Effects* (GATES) or average treatment effect by heterogeneity groups induced by the ML proxy predictor; and (3) *Classification Analysis* (CLAN) or the average characteristics of the most and least affected units defined in terms of the ML proxy predictor. Thus, we can find out if there is detectable heterogeneity in the treatment effect based on observables, and if there is any, what the treatment effect is for different bins. And finally, we can describe which of the covariates are associated with this heterogeneity.

There is a trade-off between more restrictive assumptions or tools and a more ambitious estimation. We address this trade-off by focusing on coarser objects of the function rather than the function itself, but make as little assumptions as possible. This seems to be a worthwhile sacrifice: the objects for which we have developed inference appear to

²In the recent past, several new empirical papers in economics used ML methods to estimate heterogeneous effects. For example, [Hussam, Rigol, and Roth \(2022\)](#) showed that villagers outperform the machine learning tools when they predict heterogeneity in returns to capital. [Davis and Heller \(2020\)](#) predicted who benefits the most from summer internship projects. [Deryugina, Heutel, Miller, Molitor, and Reif \(2019\)](#) used the methods developed in the present paper to evaluate the heterogeneity in the effect of air pollution on mortality. [Crepon, Duflo, Elisa, Pariente, Seban, and Veillon \(2021\)](#) also built on the present paper to develop a methodology to determine if the impact of two different programs can be accounted for by different selection. The methodological papers reviewed later also contain a number of empirical applications.

us at this point to be the most relevant, but in the future, one could easily use the same approach to develop methods to estimate other objects of interest. For example, [Crepon et al. \(2021\)](#) used the same technique to construct and estimate a specific form of heterogeneity at the post-processing stage. Even then, as we will see, getting robust and conservative standard errors for heterogeneity requires a larger sample size than just estimating average treatment effects. This reflects a different trade-off: if we do not assume that we can predict *ex ante* where the heterogeneity might be (in which case we can write it down in a pre-analysis plan), power will be lower, and detecting heterogeneity will require a larger sample. This is a consideration that applied researchers will need to keep in mind when designing and powering their experiments, and when writing pre-analysis plans: if a particular dimension of heterogeneity is deemed important, it should be pre-specified.

Another trade-off between our generic approach and the alternative approaches in the literature is that the parameters of interest are estimated conditional on a particular split of the data, since our estimation and inference methods rely on sample splitting to avoid overfitting and other inferential non-regularities. Conditional on a single data split into a training and a hold-out sample, statistical inference is conceptually straightforward and appealing. Indeed, in this case, statistical inference reduces to the classical inference for linear regression and sample means. Theoretically, if a researcher can credibly pre-commit to a single data split, this gives one clean solution to the inferential problem. However, this introduces additional variability in the results, stemming from the random draw of the particular split of the data. Therefore, researchers often consider multiple sample splits, and aggregate the results. This reduces the probability that two researchers working with the same data will arrive at different conclusions. To formalize and support this approach, we propose quantile-aggregated inference—which aggregates inferential results by taking medians of estimates and medians and other quantiles of upper and lower confidence intervals obtained from different splits. We show that quantile aggregation formally lowers estimation (reporting) risks over a single-split procedure, and we establish its inferential properties.

The proposed approach is generic in that it can be applied in conjunction with any ML method. To compare and select among ML methods, we develop goodness-of-fit measures for the BLP and GATES. We also take one step backward and use these goodness-of-fit measures to build ML proxies that better target the CATE through causal learning. We show that these causal machines produce provably better proxies of the CATE than generic (predictive) ML methods. Moreover, by designing the ML to target CATE directly, the post-processing methods that we develop can focus on providing valid inference, rather than correcting biases.

We apply our method to a large-scale RCT of nudges to encourage immunization in the state of Haryana, Northern India. This experiment, an important practical application in its own right, is designed and discussed in [Banerjee et al. \(2019\)](#). Immunization is generally recognized as one of the most effective and cost-effective ways to prevent illness, disability, and diseases. Yet, worldwide, close to 20 million children every year do not receive critical immunizations ([Unicef \(2019\)](#)). While early policy efforts have focused mainly on improving the infrastructure for immunization services, a more recent literature suggests that “nudges” (such as small incentives, leveraging the social network, SMS reminders, social signaling, etc.) may have large effects on the use of those services.³ This project was

³See, for example, [Banerjee, Duflo, Glennerster, and Kothari \(2010\)](#), [Bassani, Arora, Wazny, Gaffey, Lenters, and Bhutta \(2013\)](#), [Wakadha et al. \(2013\)](#), [Johri et al. \(2015\)](#), [Oyo-Ita, Wiysonge, Oringanje, Nwachukwu, Oduwole, and Meremikwu \(2016\)](#), [Gibson et al. \(2017\)](#), [Karing \(2018\)](#), [Domek et al. \(2016\)](#),

a collaboration with the government of Haryana, which was willing to experiment with a combination of nudges, with the goal of choosing the most effective policy and implement it at scale. It built a custom vaccination platform, and ran a large-scale experiment covering seven districts, 140 Primary health centers, 2360 villages involved in the experiment (including 915 at risk for all the treatments), and 295,038 children in the resulting database. Immunization was very low at baseline: in every single village of the district, the fraction of children whose parents had reported they received the measles vaccine (the last in the sequence) was 39%, and only 19.4% had received the vaccine before the age of 15 months, whereas the full sequence is supposed to be completed in one year. The experiment was a village-level, cross-randomized design of three main nudges: providing incentives, sending SMS reminders, and seeding ambassadors. It included several variants for each policy: the level and schedule of the incentives, the number of people receiving reminders, and the mode of selection of the ambassadors, leading to a large number (75) of finely differentiated bundles.

Banerjee et al. (2019) developed a methodology to identify the most effective and cost-effective bundle of policies, based on an application of LASSO to a marginal effects specification that imposes some structure on the bundles, and in particular, the idea that policy variants (e.g., level of incentives, or level of coverage of SMS reminders) may be indistinguishable in practice. They found that the most cost-effective policy is to combine “information hubs” (people identified by others as good at diffusing information) and SMS reminders. This is cheap and can be done everywhere. In fact, they showed that this policy is the only one among those tested that would actually save money to the government for each measles shot, while increasing immunization. But the most *effective* policy, that is, the policy that increases immunization the most, is the combination of incentives, immunization ambassadors, and SMS reminders, which is much more expensive. Yet, while this policy increases the cost per immunization, the effects are important: the number of monthly measles shots (the last vaccine in the schedule, and thus a marker for full immunization) delivered increases by 3.26, corresponding to 44% of the mean vaccination rate in the control group that got neither SMS nor increasing incentives and information hubs. The government was therefore interested in finding out where the program would be most effective, to implement it only in those places even at the higher cost per immunization.

The pre-analysis plan specified to look for heterogeneity by gender and by “Village-level baseline/national census variables, including assets, beliefs, knowledge, and attitudes towards immunization” but did not identify one or two specific baseline variables to look at. This reflected genuine uncertainty (as is often the case). Many factors can influence policy impact, from attitudes to implementation capabilities to baseline levels, and we did not have a specific theory of where to look. It is precisely the type of context that requires a principled approach to avoid overfitting, and provide a policy-relevant recommendation.⁴

The rest of the paper is organized as follows. Section 2 formalizes the framework, describes our approach, and compares it with the existing literature. Section 3 presents identification and estimation strategies for the key features of CATE of interest. Section 4 introduces our inference method that accounts for uncertainty coming from parameter estimation and sample splitting. Section 5 presents the construction of causal machines

Uddin et al. (2016), Regan, Bloomfield, Peters, and Effler (2017), Alatas, Chandrasekhar, Mobius, Olken, and Paladines (2019), Banerjee et al. (2021).

⁴This approach of finding the best treatment and then looking at where it works the best gets closer to the idea of “personalized medicine.” Using the same data, Agarwal, Alomar, Cosson, Shah, and Shen (2020) go one step further and use a “synthetic intervention” approach to look for the policy that works the best for each kind of village.

that can learn CATE better than purely predictive approaches or some existing proposals for causal approaches. Section 6 reports the results of the empirical application and provides detailed implementation algorithms. Section 7 concludes with some remarks. The Supplemental Appendix gathers proofs of the main theoretical results and additional technical results (Chernozhukov, Demirer, Duflo, and Fernández-Val (2025)).

2. OUR AGNOSTIC APPROACH

This section presents our framework and approach. We observe $\text{Data} := (Y_i, Z_i, D_i)_{i=1}^N$, consisting of i.i.d. copies of the random vector (Y, Z, D) having probability law P , where Y is the outcome of interest, D is a binary treatment indicator, and Z is a possibly high-dimensional vector of covariates that characterize the observational units. The data are defined on an underlying probability space with measure P . The expectation operator is denoted by E . When we need to emphasize the dependence of P and E on P , we use the notation P_P and E_P .

2.1. Model and Key Causal Functions

Let $Y(1)$ and $Y(0)$ be the potential outcomes in the treatment state 1 and the non-treatment state 0 (see [Neyman \(1923\)](#), [Rubin \(1974\)](#)). The main causal functions are the baseline conditional average (BCA):

$$b_0(Z) := E[Y(0)|Z], \quad (2.1)$$

and the conditional average treatment effect (CATE):

$$s_0(Z) := E[Y(1) - Y(0)|Z] = E[Y(1)|Z] - E[Y(0)|Z]. \quad (2.2)$$

Suppose D is randomly assigned conditional on Z , with probability of assignment depending on a subvector of stratifying variables $Z_1 \subseteq Z$, namely,

$$D \perp\!\!\!\perp (Y(1), Y(0))|Z, \quad (2.3)$$

and the propensity score is known and is given by

$$p(Z) := P[D = 1|Z] = P[D = 1|Z_1], \quad (2.4)$$

which we assume is bounded away from zero or 1:

$$p(Z) \in [p_0, p_1] \subset (0, 1) \quad \text{a.s.} \quad (2.5)$$

This setup is similar to [Rosenbaum and Rubin \(1983\)](#).

The observed outcome is $Y = DY(1) + (1 - D)Y(0)$. Under the stated assumption, the causal functions are identified by the components of the regression function of Y given D, Z :

$$Y = b_0(Z) + Ds_0(Z) + U, \quad E[U|Z, D] = 0, \quad (2.6)$$

that is, $b_0(Z) = E[Y|D = 0, Z]$, and

$$s_0(Z) = E[Y|D = 1, Z] - E[Y|D = 0, Z]. \quad (2.7)$$

This regression underlies the use of predictive ML methods that learn $E[Y|D, Z]$ and then estimate CATE using the formula.

Alternatively, one can identify CATE using the following two equivalent “causal” regressions:

$$s_0(Z) = E[HY|Z] = \text{Cov}[H, Y|Z], \quad (2.8)$$

where H is the residualized treatment scaled by its variance:

$$H = H(D, Z) := \frac{D - p(Z)}{p(Z)(1 - p(Z))}, \quad (2.9)$$

also known as the Horvitz–Thompson transform. We mention these alternative strategies here, because as shown in Section 5, they can lead to better ways of approximating $s_0(Z)$ than through the predictive regression (2.6), and our inference tools equally apply to ML methods that try to learn $s_0(Z)$ through either of these relations. In fact, in our empirical analysis, the strategies based on (2.8) measurably outperform the strategies based on (2.7).

2.2. Estimation and Inference Challenges

Regardless of the way we try to learn $s_0(Z)$, estimation and inference are challenging in modern high-dimensional settings, because the target function $z \mapsto s_0(z)$ can live in a very complex class. ML methods effectively explore various forms of sparsity to yield “good” approximations to $s_0(z)$. In its simplest form, sparsity reduces the complexity of $z \mapsto s_0(z)$ by assuming that it can be well-approximated by a function that only depends on a low-dimensional subset of z , making consistent estimation possible. As a result, these methods can perform much better than classical methods in high-dimensional settings under sparsity. However, sparsity or, more generally, low complexity of the CATE function s_0 , are untestable assumptions that must be used with caution.

Without some form of sparsity, it is hard, if not impossible, to obtain consistent estimators of $z \mapsto s_0(z)$. There are several fundamental reasons as well as large gaps between theory and practice that are responsible for this. One fundamental reason is that ML methods might not even produce consistent estimators of $z \mapsto s_0(z)$ in high-dimensional settings. For example, if z has dimension d and the target function $z \mapsto s_0(z)$ is assumed to have p continuous and bounded derivatives, then the worst case (minimax) lower bound on the rate of learning this function from a random sample of size N cannot be better than $N^{-p/(2p+d)}$ as $N \rightarrow \infty$, as shown by Stone (1982). Hence if p is fixed and d is also small, but slowly increasing with N , such as $d \geq \log N$, then there exists no consistent estimator of $z \mapsto s_0(z)$ generally. Hence, generic ML estimators cannot be regarded as consistent, unless further assumptions are made. Examples of such assumptions include structured forms of linear and non-linear sparsity and super-smoothness.⁵ The problem of obtaining uniformly valid inference on $z \mapsto s_0(z)$ using generic ML methods is even more difficult.⁶

⁵The function $z \mapsto s_0(z)$ is super-smooth if it has continuous and bounded derivatives of all orders.

⁶While the previous assumptions make consistent adaptive estimation possible (e.g., Bickel, Ritov, and Tsybakov (2009)), confidence sets that adapt to unknown regularity (smoothness or sparsity) do not exist even for low-dimensional nonparametric problems (Low et al. (1997), Genovese and Wasserman (2008)). Let $z \mapsto s_0(z)$ be a target function that lives in an infinite-dimensional class with unknown regularity s (e.g., smoothness or degree of sparsity). Adaptive consistent estimation (resp. inference) for $z \mapsto s_0(z)$ with respect to s is possible

In this paper, we take an agnostic view. We neither rely on any sparsity or low-complexity assumptions to make the ML estimators consistent, nor impose other stronger conditions to make “traditional” confidence intervals valid. We simply treat ML as providing proxy predictors for the objects of interest.

2.3. Our Approach

To address the previous challenges, we propose strategies for estimation and inference on *key features* of $s_0(Z)$ rather than on $s_0(Z)$ itself. Because of this difference in focus, and by relying on sample splitting, we can loosen the restrictions about the properties of the ML estimators.

Let (M, A) denote a random partition of the set of indices $\{1, \dots, N\}$. The strategies that we consider rely on random splitting of $\text{Data} = (Y_i, D_i, Z_i)_{i=1}^N$ into a main sample, denoted by $\text{Data}_M = (Y_i, D_i, Z_i)_{i \in M}$, and an auxiliary sample, denoted by $\text{Data}_A = (Y_i, D_i, Z_i)_{i \in A}$. We will sometimes refer to these samples as M and A . After splitting the sample, we carry out two stages:

Stage 1: From the auxiliary sample A , we obtain ML estimators of the baseline functions and treatment effects, which we call the ML proxy predictors,

$$z \mapsto B(z) \quad \text{and} \quad z \mapsto S(z).$$

Here $S(Z)$ is a possibly biased and noisy predictor of $s_0(z)$ and $B(Z)$ is a possibly biased and noisy predictor of $b_0(Z)$ (or other technical “baseline” functions, as we discuss in Section 5). We do not require these predictors to be consistent for the true functions.

Stage 2: We post-process the proxies from Stage 1 to estimate and make inference on features of the CATE function $z \mapsto s_0(z)$ in the main sample M . The key features that we target include:

- (1) **Best Linear Predictor** (BLP) of the CATE $s_0(Z)$ on the ML proxy predictor $S(Z)$;
- (2) **Sorted Group Average Treatment Effects** (GATES): average of $s_0(Z)$ (ATE) by heterogeneity groups induced by the ML proxy predictor $S(Z)$;
- (3) **Classification Analysis** (CLAN): average characteristics of the most and least affected units defined in terms of the ML proxy predictor $S(Z)$.

Our approach is *generic* with respect to the ML method being used, and is *agnostic* about its formal properties. However, it relies on sample splitting, which introduces a specific source of uncertainty. To account for this, we use many data splits into main and auxiliary samples to produce robust estimators, and we employ quantile aggregation of inference to combine results across splits. Specifically, for point estimation, we report the median of the estimated key features over different random splits of the data. We take medians and other quantiles of many random conditional confidence sets for interval estimation. Finally, we construct p -values by taking medians of many random conditional p -values. We establish the formal inferential properties of this procedure.

2.4. Relationship to the Literature

We focus the review strictly on the literatures about estimation and inference on heterogeneous effects and inference using sample splitting.

if there exists a consistent estimator (resp. valid confidence set) with a rate of convergence (resp. diameter) that changes with s in a (nearly) rate-optimal way. Construction of adaptive confidence bands then requires making additional untestable assumptions. See, for example, [Giné and Nickl \(2010\)](#), where self-similarity conditions are used in low-dimensional nonparametric problems.

This work is related to the literature that uses linear and semiparametric regression methods for estimation and inference on heterogeneous effects. Crump, Hotz, Imbens, and Mitnik (2008) developed tests of treatment effect homogeneity for low-dimensional settings based on traditional series estimators of the CATE. A semiparametric inference method for characterizing heterogeneity, called the sorted effects method, was given in Chernozhukov, Fernandez-Val, and Luo (2015). This approach does provide a full set of inference tools, including simultaneous bands for percentiles of the CATE, but is strictly limited to the traditional semiparametric estimators of the regression and causal functions. Hansen, Kozbur, and Misra (2017) proposed a sparsity-based method called “targeted undersmoothing” to perform inference on heterogeneous effects. This approach does allow for high-dimensional settings, but imposes sparsity as well as additional assumptions that enable the targeted undersmoothing. A related approach, which allows for simultaneous inference on many coefficients (e.g., inference on the coefficients corresponding to the interaction of the treatment with other variables), is proposed in Belloni, Chernozhukov, and Kato (2013) using a Z-estimation framework, where the number of interactions can be very large; see also Dezeure, Bühlmann, and Zhang (2016) for a more recent effort in this direction, focusing on de-biased lasso in mean regression problems. This approach, however, still relies on a strong form of sparsity assumptions. Zhao, Small, and Ertefaie (2017) proposed a post-selection inference framework within high-dimensional linear sparse models for the heterogeneous effects. The approach is attractive because it allows for some misspecification of the model.

Another approach is to use tree-based and other methods. Imai and Ratkovic (2013) discussed the use of a heuristic support-vector-machine method with lasso penalization for classification of heterogeneous treatments into positive and negative ones. They used the Horvitz–Thompson transformation of the outcome (e.g., as in Hirano, Imbens, and Ridder (2003), Abadie (2005)) such that the new outcome becomes an unbiased, noisy version of CATE.⁷ Athey and Imbens (2016) made use of the Horvitz–Thompson transformation of the outcome to inform the process of building causal trees, with the main goal of predicting CATE. They also provided a valid inference result on average treatment effects for groups defined by the tree leaves, conditional on the data split into two subsamples: one used to build the tree leaves and one to estimate the predicted values given the leaves. Like our methods, this approach is essentially assumption-free. Our paper is a complement, in that our approach can be used with any ML method. Wager and Athey (2018) proposed a subsampling-based construction of a causal random forest, providing valid pointwise inference for CATE (see also the review in Wager and Athey (2018) on prior uses of random forests in causal settings) for the case when covariates are very low-dimensional (and essentially uniformly distributed).⁸ This condition rules out the typical high-dimensional settings that arise in many empirical problems, especially in current RCTs, where the number of baseline covariates is potentially very large.

Several other studies look at model-based strategies for performing inference on CATE. Semenova and Chernozhukov (2021) used ML to perform inference on the “partial” CATE, $E[s_0(Z)|X]$, where X is a prespecified low-dimensional set of covariates. Specifically, they constructed an estimator of a denoised HT transform of the outcome

⁷Note that using Horvitz–Thompson (HT) transform of outcome, in this and other references, typically gives very noisy signal. One can improve the approach by either including the HT transform interacted with some baseline covariates as regressors in a regression model, as we do in the present paper, or using residualized outcomes in conjunction with HT, as in Semenova and Chernozhukov (2021).

⁸The dimension d is fixed in Wager and Athey (2018); the analysis relies on the Stone’s model with smoothness index $\beta = 1$, in which no consistent estimator exists once $d \geq \log n$ in the minimax sense.

and projected it using a nonparametric series estimator on the set of low-dimensional prespecified covariates of interest X , whose dimension is much lower than the dimension of Z .⁹ The main advantage of this approach is that it delivers familiar nonparametric inference on partial CATE (even though inference on the full CATE remains intractable). Fan, Hsu, Lieli, and Zhang (2022), Zimmert and Lechner (2019), Chernozhukov, Newey, and Singh (2018), and Chernozhukov, Newey, and Singh (2021) developed kernel versions of this procedure. Related ideas but based on partialing-out (using residualized outcomes and treatment) appear in Semenova, Goldman, Chernozhukov, and Taddy (2017), Nekipelov, Semenova, and Syrgkanis (2022), Nie and Wager (2020), Foster and Syrgkanis (2019), and Kennedy (2020). Relative to the approach taken here, the assumptions made in these papers are more restrictive, but deliver stronger results. For example, in the approach of Semenova and Chernozhukov (2021), one has to specify the baseline covariates X for the partial CATE analysis (which is exactly what we are trying to avoid in our approach). Second, the methods critically rely on the consistency of ML to estimate the nuisance components well, which can be restrictive in high-dimensional settings, as discussed above.

The idea of using a “hold-out” sample to validate the result of an ML procedure to discover heterogeneity was suggested in Davis and Heller (2020), who used the method proposed in Wager and Athey (2018) and compared their results to the heterogeneity in a holdout sample. Our inference approach is different because it calls for multiple splits. This procedure itself is also of independent interest and could be applied to many problems, where sample splitting is used to produce ML predictions (e.g., Abadie, Chingos, and West (2017)). Related references include Wasserman and Roeder (2009), and Meinshausen, Meier, and Bühlmann (2009), where the ideas are related, but the details are quite different, as we shall explain below. The premise is the same; however, as in Meinshausen, Meier, and Bühlmann (2009) and Rinaldo, Wasserman, G’Sell, Lei, and Tibshirani (2016), we should not rely on a single random split of the data and should adjust inference in some way. Our construction of p -values builds upon ideas in Meinshausen, Meier, and Bühlmann (2009), though what we propose is simpler, and our confidence intervals appear to be new. Of course, sample splitting ideas are classical, going back to Hartigan (1969), Kish and Frankel (1974), Barnard (1974), Cox (1975), Mosteller and Tukey (1977), though having been mostly underdeveloped and overlooked for inference, as characterized by Rinaldo et al. (2016). Finally, our inference method shares with the literature on post-selection inference in statistics that the target estimands are random functions depending on an ML proxy (e.g., Fithian, Sun, and Taylor (2014), Lee, Sun, Sun, and Taylor (2016)).

3. MAIN IDENTIFICATION RESULTS AND ESTIMATION STRATEGIES

In this section, we condition on Data_A (the auxiliary sample in a random split) and therefore consider the functions

$$z \mapsto B_A(z) := B(z; \text{Data}_A) \quad \text{and} \quad z \mapsto S_A(z) := S(z; \text{Data}_A)$$

as fixed functions. To lighten the notation, we keep the dependence on Data_A implicit in all the expectations, objects, and parameters; that is, for example, we use $E[\cdot]$ and $S(Z)$

⁹Specifically, the denoised HT transform of outcome is $\tilde{Y} = g(1, Z) - g(0, Z) + H(Y - g(D, Z))$, where $H = (D - p(Z))/[p(Z)(1 - p(Z))]$ and $g(D, Z) = E(Y|D, Z)$. Semenova and Chernozhukov (2021) used ML to estimate $g(D, Z)$ and the propensity score $p(Z)$, in case the latter is unknown.

instead of $E[\cdot | \text{Data}_A]$ and $S_A(Z)$. We shall make the dependence on Data_A explicit when we discuss estimation and inference in Section 4.

3.1. Best Linear Predictor of CATE

The first inferential target is the best linear predictor of the CATE using the proxy $S(Z)$.

DEFINITION 3.1—BLP: The best linear predictor of $s_0(Z)$ by $S(Z)$ is the solution to

$$\min_{b_1, b_2} E[s_0(Z) - b_1 - b_2 S(Z)]^2,$$

which, if exists, is defined as

$$\text{BLP}[s_0(Z) | S(Z)] := \beta_1 + \beta_2 (S(Z) - ES(Z)),$$

where $\beta_1 = Es_0(Z)$ and $\beta_2 = \text{Cov}[s_0(Z), S(Z)] / \text{Var}[S(Z)]$.

By construction, $\text{BLP}[s_0(Z) | S(Z)]$ is an unbiased predictor of $s_0(Z)$, which improves over $S(Z)$ in the mean squared error sense, that is,

$$E\{s_0(Z) - \text{BLP}[s_0(Z) | S(Z)]\}^2 \leq E[s_0(Z) - S(Z)]^2.$$

Indeed, we can quantify the improvement by

$$\begin{aligned} & E[s_0(Z) - S(Z)]^2 - E\{s_0(Z) - \text{BLP}[s_0(Z) | S(Z)]\}^2 \\ &= (1 - \beta_2)^2 \text{Var}[S(Z)] + [ES(Z) - Es_0(Z)]^2, \end{aligned}$$

which is positive unless $S(Z)$ is an unbiased predictor and, either $\beta_2 = 1$ or $\text{Var}[S(Z)] = 0$.¹⁰ Accordingly, compared to the ML proxy, the BLP can be seen as a refined predictor of the individual CATE, $s_0(Z)$. If $S(Z)$ is a perfect proxy for $s_0(Z)$, then $\beta_2 = 1$. In general, $\beta_2 \neq 1$, correcting for noise in $S(Z)$. If $S(Z)$ is complete noise, uncorrelated to $s_0(Z)$, then $\beta_2 = 0$. Furthermore, if there is no heterogeneity, that is, $s_0(Z) = s$, then $\beta_2 = 0$. Rejecting the hypothesis $\beta_2 = 0$ therefore means both that there is heterogeneity in $s_0(Z)$ and $S(Z)$ is a relevant predictor.

COMMENT 3.1—Conditioning on Data_A : As we mention above, all the objects and parameters are conditional on the auxiliary sample Data_A . If we make the dependence explicit,

$$\text{BLP}_A[s_0(Z) | S_A(Z)] = \beta_{A,1} + \beta_{A,2} (S_A(Z) - E[S_A(Z) | \text{Data}_A]).$$

The BLP_A can be interpreted as the best linear predictor of $s_0(Z)$ given the proxy $S_A(Z)$ obtained from a random partition of the data. Therefore, BLP_A is random with respect to A , but converges to a fixed number as the size of the auxiliary sample grows under

¹⁰The previous expression follows from the decompositions $E[s_0(Z) - S(Z)]^2 = E[\{s_0(Z) - Es_0(Z)\} - \{S(Z) - ES(Z)\} + \{Es_0(Z) - ES(Z)\}]^2$ and $E\{s_0(Z) - \text{BLP}[s_0(Z) | S(Z)]\}^2 = E[\{s_0(Z) - Es_0(Z)\} - \{\text{BLP}[s_0(Z) | S(Z)] - Es_0(Z)\}]^2$, using that $E\text{BLP}[s_0(Z) | S(Z)] = Es_0(Z)$ and $\beta_2 = \text{Cov}(s_0(Z), S(Z)) / \text{Var}(S(Z))$.

suitable stability conditions on $S_A(Z)$. This is discussed in Section B of the Supplemental Appendix. If, in addition, $S_A(z)$ is a consistent estimator of $s_0(z)$, then $\beta_{1,A} \rightarrow_P \text{E}s_0(Z)$ and $\beta_{2,A} \rightarrow_P 1$ as the size of the auxiliary sample grows.

We provide two strategies for identifying and estimating $\text{BLP}[s_0(Z)|S(Z)]$.

Strategy A: Weighted Residual BLP

Consider the weighted linear projection:

$$Y = \alpha'_0 X_1 + \alpha_1(D - p(Z)) + \alpha_2(D - p(Z))(S - \text{ES}) + \epsilon, \quad \text{E}[w(Z)\epsilon X] = 0, \quad (3.1)$$

where $S := S(Z)$, $w(Z) := \{p(Z)(1 - p(Z))\}^{-1}$, $X := (X'_1, X'_2)'$,

$$X_1 = [1, B(Z), p(Z), p(Z)S(Z)]', \quad X_2 := [D - p(Z), (D - p(Z))(S - \text{ES})]'$$

The term $B(Z)$ could be replaced by any “noise-reducing” proxy function. For example, the algorithms of [Semenova et al. \(2017\)](#) and [Nie and Wager \(2020\)](#), targeting the delivery of $S(Z)$, also construct $B(Z)$ that are meant to approximate $\text{E}[Y|Z]$ but not $b_0(Z)$; see also Section 5 for other examples of such algorithms. We include X_1 to reduce finite sample noise of the estimators of the BLP parameter based on this strategy.¹¹

Note that α_1 and α_2 are identified under weak assumptions. Further, we note that the interaction $(D - p(Z))(S - \text{ES})$ is orthogonal to $D - p(Z)$ under the weight $w(Z)$, and to all functions of Z such as X_1 . Consequently, we obtain the following result that shows that the linear projection (3.1) identifies the BLP.

THEOREM 3.1—BLP Identification A: *Consider $z \mapsto S(z)$ and $z \mapsto B(z)$ as fixed maps. Assume that Y and X have finite second moments, and $\text{E}XX'$ is finite and full rank, which requires $\text{Var}(S(Z)) > 0$. Then, (α_1, α_2) defined in (3.1) identifies the coefficients of the BLP,*

$$\alpha_1 = \beta_1, \quad \alpha_2 = \beta_2.$$

COMMENT 3.2—Why Not Classical OLS of Y on Proxies?: It is tempting and perhaps more natural to consider the projection equation

$$Y = \underbrace{\tilde{\alpha}_1 + \tilde{\alpha}_2 B + \tilde{\beta}_1 D + \tilde{\beta}_2 D(S - \text{ES})}_{\text{BLP of CEF}} + \epsilon, \quad \text{E}[\epsilon \tilde{X}] = 0,$$

where $\tilde{X} = [1, B, D, D(S - \text{ES})]'$. The idea here is the classical one: the ordinary least squares method with Y as the outcome provides the Best Linear Predictor or Approximation to the CEF $\text{E}[Y|D, Z]$, even if the latter is nonlinear. [Angrist and Pischke \(2008\)](#) discuss the importance and practical relevance of this property. However, this property does not translate into providing the BLP of CATE $s_0(Z)$. Indeed, even in pure RCTs, while $\tilde{\beta}_1 = \beta_1$ is true, we have that $\tilde{\beta}_2 \neq \beta_2$ in general, and therefore

$$\tilde{\beta}_1 + \tilde{\beta}_2(S - \text{ES}) \neq \text{BLP}(s_0(Z)|S). \quad (3.2)$$

See Appendix A of the Supplemental Appendix for further discussion and proof.

¹¹Note that X_1 can include other functions of Z . In our experiments, the use of $B(Z)$ strongly improves the precision of estimating BLP (and other quantities such as GATEs and CLAN introduced below).

The identification result in Theorem 3.1 is constructive. We can base a corresponding estimation strategy on the empirical analog:

$$Y_i = \hat{\alpha}_0 X_{1i} + \hat{\alpha}_1 (D_i - p(Z_i)) + \hat{\alpha}_2 (D_i - p(Z_i))(S_i - \mathbb{E}_{N,M} S_i) + \hat{\epsilon}_i, \quad i \in M, \\ \mathbb{E}_{N,M} [w(Z_i) \hat{\epsilon}_i X_i] = 0, \quad (3.3)$$

where $X_i = [X'_{1i}, X'_{2i}]'$, $X_{1i} := [1, B(Z_i), p(Z_i), p(Z_i)S(Z_i)]'$, $X_{2i} := [D_i - p(Z_i), (D_i - p(Z_i))(S_i - \mathbb{E}_{N,M} S_i)]'$, and $\mathbb{E}_{N,M}$ denotes the empirical expectation with respect to the main sample, that is,

$$\mathbb{E}_{N,M} h(Y_i, D_i, Z_i) := |M|^{-1} \sum_{i \in M} h(Y_i, D_i, Z_i).$$

Figure 1 provides two examples. The left panel shows a case where $s_0(Z) = 0$ with zero effect and zero heterogeneity in the CATE, whereas the right panel shows a case where $s_0(Z) = Z$ with strong heterogeneity in the CATE. In both cases, we evenly split 1000 observations between the auxiliary and main samples, Z follows uniform distribution on $(-1, 1)$, $b_0(Z) = 3Z$, U is standard normal, independently of Z , and Y is generated by (2.6). We obtain the proxy predictor $S(Z)$ by Breiman's random forest, using the ranger implementation in R (Wright and Ziegler (2017)).

In the first example, the ML proxy is pure noise by construction, and the BLP post-processor correctly eliminates the noise, producing a CATE prediction that is roughly a 68% better approximation to the CATE under the RMSE metric. Furthermore, using our inferential methods of Section 4, we cannot reject the null hypothesis that the BLP is zero. In the second case, under strong heterogeneity, the signal in the ML proxy dominates the noise component. As a result, the BLP does not change the ML proxy drastically,

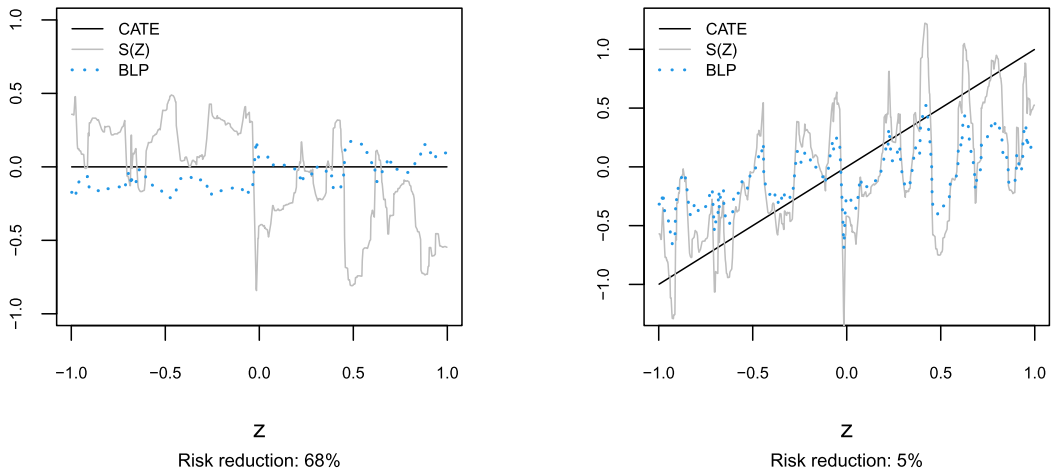


FIGURE 1.—BLP using ML proxy versus the ML proxy. NOTES: The CATE is plotted with the solid black line; the proxy predictor $S(Z)$, produced by Random Forest, is plotted with the solid gray (light) line; and the BLP is plotted with the dotted line. The left panel corresponds to the no heterogeneity example, $s_0(z) = 0$, and the right panel to the strong heterogeneity example, $s_0(z) = z$. In both panels, the BLP is less noisy than the ML proxy, reducing the RMSE by 68% and 5%.

but still gives a meaningful improvement to the CATE under the RMSE metric. These improvements agree with the theoretical arguments given above.¹²

COMMENT 3.3—Significance for RCTs: The first example has implications for the empirical analysis of RCTs. Here we see that one of the best ML algorithms, as per [Friedman, Hastie, and Tibshirani \(2001\)](#), can easily suggest a heterogeneous CATE when the treatment is, in fact, a placebo. Placebos (ineffective treatments) are common occurrences in real-world experiments, and our methodology provides a simple way to confirm that the CATE (and not just ATE) is indeed zero in such cases.

Strategy B: HT BLP

This strategy makes use of the Horvitz–Thompson transform H defined in (2.9). It is well known that the transformed response YH provides an unbiased signal about CATE:

$$E[YH|Z] = s_0(Z),$$

and it follows by the properties of the best linear predictor that

$$\text{BLP}[s_0(Z)|S(Z)] = \text{BLP}[YH|S(Z)].$$

Note that $\text{BLP}[s_0(Z)|S(Z)]$ is a more precise unbiased predictor of $s_0(Z)$ than YH because, by construction,

$$\text{Var}(\text{BLP}[s_0(Z)|S(Z)]) = \text{Var}(\text{BLP}[YH|S(Z)]) \leq \text{Var}(YH).$$

The R-squared of $\text{BLP}[YH|S(Z)]$ quantifies the percent reduction in variance of the BLP relative to YH .

The simple linear projection $\text{BLP}[YH|S(Z)]$ is completely fine for identification purposes, but can severely underperform in estimation and inference due to lack of precision. We can repair the deficiencies by considering, instead, the linear projection

$$YH = \mu'_0 X_1 H + \mu_1 + \mu_2(S - ES) + \epsilon, \quad E\epsilon\tilde{X} = 0, \quad (3.4)$$

where $\tilde{X} := (X'_1 H, \tilde{X}'_2)'$, $\tilde{X}_2 := (1, S - ES)'$, and $X_1 := [1, B(Z), p(Z), p(Z)S(Z)]'$ as before. The term X_1 could contain other functions of Z . We include $X_1 H$ in order to *reduce noise*.

The following theorem shows that the linear projection (3.4) also identifies the BLP.

THEOREM 3.2—BLP Identification B: *Consider $z \mapsto S(z)$ and $z \mapsto B(z)$ as fixed maps. Assume that Y has finite second moments, \tilde{X} is such that $E\tilde{X}\tilde{X}'$ is finite and full rank, which requires $\text{Var}(S(Z)) > 0$. Then, (μ_1, μ_2) defined in (3.4) identifies the coefficients of the BLP,*

$$\mu_1 = \beta_1, \quad \mu_2 = \beta_2.$$

¹²To show that the risk reductions are not a fluke, we repeated the calculations in 1000 simulations. We found average risk reductions of 65% and 18% in the no heterogeneity and strong heterogeneity examples, respectively.

Theorem 3.2 leads to an estimator defined through the empirical analog:

$$Y_i H_i = \hat{\mu}'_0 X_{1i} H_i + \hat{\mu}_1 + \hat{\mu}_2 (S_i - \mathbb{E}_{N,M} S_i) + \hat{\epsilon}_i, \quad i \in M, \mathbb{E}_{N,M} \hat{\epsilon}_i \tilde{X}_i = 0, \quad (3.5)$$

where $\tilde{X}_i := (X'_{1i} H_i, \tilde{X}'_{2i})'$, $\tilde{X}_{2i} := (1, S_i - \mathbb{E}_{N,M} S_i)'$, and $X_{1i} := [1, B(Z_i), p(Z_i), p(Z_i) \times S(Z_i)]'$.

COMMENT 3.4—BLP Estimators: The properties of the estimators in (3.3) and (3.5), conditional on the auxiliary data, are given in Lemma D.1 of the Supplemental Appendix.

COMMENT 3.5—Comparison A vs. B: While the identification strategies A and B are natural, one may wonder whether the two corresponding estimation strategies can be ranked in terms of asymptotic efficiency. We show in Appendix A of the Supplemental Appendix that they produce estimators that are first-order equivalent in large main samples.

3.2. Sorted Group Average Treatment Effects

The second inferential target is the set of group average treatment effects, where the groups are induced by $S(Z)$.

DEFINITION 3.2—GATES: The Sorted Group Average Treatment Effects (GATES) are

$$\gamma_k := \mathbb{E}[s_0(Z) | G_k], \quad k = 1, \dots, K,$$

where $G_k := \{S \in I_k\}$, with $I_k := [\ell_{k-1}, \ell_k)$ and $-\infty = \ell_0 < \ell_1 < \dots < \ell_K = +\infty$.

COMMENT 3.6—Choice of Groups: We build the groups to explain as much variation in $s_0(Z)$ as possible. There are many alternatives for creating groups based upon ML tools applied to the auxiliary data. For example, one can group or cluster based upon predicted baseline response as in the “endogenous stratification” analysis (Abadie, Chingos, and West (2017)), or based upon actual predicted treatment effect S . We focus on the latter approach for defining groups, although our identification and inference ideas immediately apply to other ways of defining groups, and could be helpful in these contexts. The causal tree approach of Athey and Imbens (2016) can also be viewed as a GATES analysis, with a specific way of forming groups via recursive partitioning.¹³

COMMENT 3.7—GATES as Predictors of CATE: The GATES can also be used as non-linear predictors of the CATE based on the proxy S , in a similar fashion to the BLP. Indeed, the GATES provide the BLP of CATE using the group indicators $G_k, k = 1, \dots, K$.

We provide two strategies for identifying and estimating the GATES.

¹³Another strand of the literature related to the GATES is the learning policy problem, where an ML method is trained to assign units to treatment and control based on their covariates (e.g., Kitagawa and Tetenov (2018), Athey and Wager (2021)). This problem can be seen as a GATES analysis with two groups chosen to maximize some function of the CATES.

Strategy A: Weighted Residual GATES

Consider the weighted linear projection equation

$$Y = \alpha'_0 X_1 + \sum_{k=1}^K \alpha_k \cdot [D - p(Z)] \cdot 1(G_k) + \nu, \quad \mathbb{E}[w(Z)\nu W] = 0, \quad (3.6)$$

where $W := (X'_1, W'_2)'$, X_1 contains a vector of functions of Z , for example, $X_1 = (B(Z), p(Z)\{1(G_k)\}_{k=1}^K)'$ and $W_2 := (\{[D - p(Z)] \cdot 1(G_k)\}_{k=1}^K)'$. The presence of $D - p(Z)$ in the interaction $[D - p(Z)] \cdot 1(G_k)$ *orthogonalizes* this regressor relative to all other regressors that are functions of Z , such as X_1 . The controls in X_1 , as in the BLP, are included to reduce noise in estimation.

Theorem 3.3 below shows that the linear projection (3.6) identifies the GATES. We can therefore base an estimation strategy on the empirical analog

$$Y_i = \hat{\alpha}'_0 X_{1i} + \hat{\alpha}' W_{2i} + \hat{\nu}_i, \quad i \in M, \quad \mathbb{E}_{N,M}[w(Z_i)\hat{\nu}_i W_i] = 0, \quad (3.7)$$

where $\hat{\alpha} = (\hat{\alpha}_1, \dots, \hat{\alpha}_K)'$.

Strategy B: HT GATES

Here we employ a linear projection on Horvitz–Thompson transformed variables:

$$YH = \mu'_0 X_1 H + \sum_{k=1}^K \mu_k \cdot 1(G_k) + \nu, \quad \mathbb{E}[\nu \tilde{W}] = 0, \quad (3.8)$$

where $\tilde{W} := (X'_1 H, \tilde{W}'_2)'$, X_1 includes functions of Z , for example, X_1 the same as above, and $\tilde{W}_2 := [\{1(G_k)\}_{k=1}^K]'$.

Theorem 3.3 shows that the linear projection (3.8) also identifies the GATES. We can therefore base an estimation strategy on the empirical analog:

$$Y_i H_i = \hat{\mu}'_0 X_{1i} H_i + \hat{\mu}' \tilde{W}_{2i} + \hat{\nu}_i, \quad i \in M, \quad \mathbb{E}_{N,M}[\hat{\nu}_i \tilde{W}_i] = 0, \quad (3.9)$$

where $\hat{\mu} = (\hat{\mu}_1, \dots, \hat{\mu}_K)'$. The resulting estimator has similar performance to the estimator in (3.7), and under some conditions their first-order properties coincide.

We now provide a formal statement of the identification results.

THEOREM 3.3—GATES: *Consider $z \mapsto S(z)$ and $z \mapsto B(z)$ as fixed maps. Assume that Y has finite second moments and W and \tilde{W} are such that $EW W'$ and $E\tilde{W} \tilde{W}'$ are finite and have full rank. Consider $\alpha = (\alpha_k)_{k=1}^K$ defined by the weighted regression equation (3.6) and $\mu = (\mu_k)_{k=1}^K$ defined by the regression equation (3.8). These coefficients are equal and identify the GATES:*

$$\alpha_k = \mu_k = \gamma_k = \mathbb{E}[s_0(Z)|G_k], \quad k = 1, \dots, K.$$

COMMENT 3.8—GATES Estimators: The properties of the estimators in (3.7) and (3.9), conditional on the auxiliary data, are given in Lemma D.1 of the Supplemental Appendix.

3.3. Classification Analysis

When the BLP and GATES analyses reveal substantial heterogeneity, it is interesting to know the properties of the subpopulations that are the most and least affected. Here we focus on the “least affected group” G_1 and “most affected group” G_K , where the labels “most” and “least” can be swapped depending on the context.

DEFINITION 3.3—CLAN: Let $g(Y, D, Z)$ be a vector of characteristics of an observational unit. The classification analysis (CLAN) is the comparison of the average characteristics of the most and least affected groups:

$$\delta_1 := E[g(Y, D, Z)|G_1] \quad \text{and} \quad \delta_K := E[g(Y, D, Z)|G_K].$$

The parameters δ_1 and δ_K are identified because they are averages of variables that are directly observed. The CLAN quantifies the differences between the most and least affected groups and singles out the covariates that are associated with the heterogeneity in the CATE. The CLAN can be extended to comparisons of features other than averages, such as variances, covariances, or distributions. In the empirical analysis, we estimate the CLAN parameters by taking averages in M :

$$\hat{\delta}_1 = \frac{\mathbb{E}_{N,M}[g(Y_i, D_i, Z_i)G_{1,i}]}{\mathbb{E}_{N,M}G_{1,i}} \quad \text{and} \quad \hat{\delta}_K = \frac{\mathbb{E}_{N,M}[g(Y_i, D_i, Z_i)G_{K,i}]}{\mathbb{E}_{N,M}G_{K,i}}, \quad (3.10)$$

using $G_{k,i} = 1\{S(Z_i) \in I_k\}$, where $I_k = [\ell_{k-1}, \ell_k)$ and ℓ_k is the (k/K) -quantile of $\{S_i\}_{i \in M}$.

3.4. Goodness-of-Fit Measures for Fitting CATE

In practical applications, it is useful to have goodness-of-fit measures to guide the selection of ML proxies.

For the analysis based on the BLP of CATE, we propose to use

$$\Lambda := |\beta_2|^2 \text{Var}(S(Z)) = \text{Corr}(s_0(Z), S(Z))^2 \text{Var}(s_0(Z)). \quad (3.11)$$

Maximizing Λ is equivalent to maximizing the correlation between the ML proxy predictor $S(Z)$ and the CATE $s_0(Z)$, or equivalent to maximizing the R^2 in the regression of $s_0(Z)$ on $S(Z)$. Therefore, an ML method that attains a higher Λ is a preferred method.

Analogously, for the GATES analysis, we propose to use

$$\bar{\Lambda} = E \left(\sum_{k=1}^K \gamma_k 1(S \in I_k) \right)^2 = \sum_{k=1}^K \gamma_k^2 P(S \in I_k). \quad (3.12)$$

This is the part of variation of $s_0(z)$, $E s_0(Z)^2$, explained by $\bar{S}(Z) = \sum_{k=1}^K \gamma_k 1(S(Z) \in I_k)$. Hence, choosing the ML proxy $S(Z)$ to maximize $\bar{\Lambda}$ is equivalent to maximizing the R^2 in the regression of $s_0(Z)$ on $\bar{S}(Z)$ (without a constant). If the groups $G_k = \{S \in I_k\}$ have equal size, namely, $P(S(Z) \in I_k) = 1/K$ for each $k = 1, \dots, K$, then

$$\bar{\Lambda} = \frac{1}{K} \sum_{k=1}^K \gamma_k^2.$$

Therefore, an ML method that attains a higher $\bar{\Lambda}$ is a preferred method. The empirical versions of the parameters above are

$$\widehat{\Lambda} = |\widehat{\beta}_2|^2 \mathbb{E}_{N,M}(S_i - \mathbb{E}_{N,M} S_i)^2, \quad \widehat{\bar{\Lambda}} = \sum_{k=1}^K \widehat{\gamma}_k^2 \mathbb{E}_{N,M} 1\{S_i \in I_k\}. \quad (3.13)$$

The choice of the ML method using goodness-of-fit measures does not pose any additional inferential challenge, when there is clearly an ML method that dominates the others—so that we select the best method with probability approaching 1. This means that the inferential methods of the next section would not need any further adjustment. When this is not the case, there are several possibilities depending on the scientific reporting objectives. For example, suppose there are two (near) winners and we want to construct a $(1 - \alpha)$ -confidence set, as in our empirical analysis. Then, in the spirit of sensitivity analysis, we can report the union of the $(1 - \alpha)$ -confidence sets. This ensures the inferential coverage guarantee of $1 - \alpha$ continues to apply, if the reader of the empirical report chooses one or the other winner at random. On the other hand, the inferential guarantee needs to be discounted to $1 - 2\alpha$ via Bonferroni adjustment, if the reader of the empirical report chooses one or the other report depending on the empirical results themselves. We use the first approach because the readers of our empirical analysis are not likely to follow the latter approach.

Finally, if the data sets are big, we could use additional splitting to choose the best-performing ML method, before taking the resulting ML proxies to the main sample.¹⁴

4. SPLIT-SAMPLE ROBUST ESTIMATION AND INFERENCE METHODS

4.1. *Estimation and Inference: The Generic Targets*

Let θ denote a generic target parameter or functional. For example,

- $\theta = \beta_2$ is the BLP slope, the heterogeneity loading parameter;
- $\theta = \text{BLP}[s_0(Z)|S(z)] = \beta_1 + \beta_2(S(z) - \text{ES})$ is the “personalized” predictor of CATE $s_0(z)$;
- $\theta = \gamma_k$ is the GATES for the group G_k ;
- $\theta = \gamma_K - \gamma_1$ is the difference in the GATES between the most and least affected groups;
- $\theta = \delta_K - \delta_1$ is the difference in the expectation of the characteristics of the most and least impacted groups in CLAN.

Let (a, m) denote a fixed partition of $\{1, \dots, N\}$. In this section, we make explicit when the probabilities and expectations are conditional on the auxiliary sample a ,

$$\text{Data}_a := \{(Y_i, D_i, X_i)\}_{i \in a},$$

and also index the ML proxies $B = B_a$ and $S = S_a$ and estimands $\theta = \theta_a$ by a to make explicit the dependence on Data_a . As we mention in Comment 3.1, this dependence vanishes as the size of the auxiliary sample becomes large under suitable conditions.

¹⁴We also refer to Section 5 which discusses other, more exploratory ideas for building the best ML algorithms for targeting CATE already in the first stage.

Single Split

We begin the discussion of inference conditional on a single split of data induced by the partition $\{(a, m)\}$ of $\{1, \dots, N\}$ into sets of cardinality $(N - n, n)$. All of the examples admit an estimator $\widehat{\theta}_a$ that is approximately Gaussian, conditionally on Data_a , namely, as $(N - n, n) \rightarrow \infty$ and for any z ,

$$P(\widehat{\sigma}_a^{-1}(\widehat{\theta}_a - \theta_a) < z | \text{Data}_a) \rightarrow_P \Phi(z). \quad (4.1)$$

We provide sufficient regularity conditions for (4.1) in Lemma D.1, which may be of independent interest (in single-split inference context).¹⁵

As a consequence of (4.1), the standard p -values

$$p_a^+ := 1 - \Phi(\widehat{\sigma}_a^{-1}(\widehat{\theta}_a - \theta_0)), \quad p_a^- := \Phi(\widehat{\sigma}_a^{-1}(\widehat{\theta}_a - \theta_0)),$$

for testing the null hypothesis $\theta_a = \theta_0$ against the alternatives $\theta_a > \theta_0$ and $\theta_a < \theta_0$, respectively, are approximately uniform under the null, namely,

$$P(p_a^\pm < \alpha | \text{Data}_a) = \alpha + o_P(1).$$

As another consequence of (4.1), the standard confidence interval (CI)

$$[L_a, U_a] := [\widehat{\theta}_a \pm \Phi^{-1}(1 - \alpha/2)\widehat{\sigma}_a]$$

covers θ_a with approximate probability $1 - \alpha$ conditional on Data_a :

$$P[\theta_a \in [L_a, U_a] | \text{Data}_a] = 1 - \alpha - o_P(1).$$

Thus, we have straightforward inference conditional on a single data split.

Multiple Splits

In practice, researchers often prefer using multiple splits (a, m) 's to reduce estimation risk and demonstrate that the results are robust to how they split the data. Therefore, we need a way to aggregate the results across different splits, and propose quantile aggregation methods and analyze their properties.

DEFINITION 4.1—Collection of Splits: Consider the collection $\{(a, m), a \in \mathcal{A}\}$ of partitions of $[N] = \{1, \dots, N\}$ into auxiliary sets a of cardinality $N - n$ and main sets m of cardinality n . We generate the collection independently of

$$\text{Data} := (Y_i, D_i, X_i)_{i=1}^N.$$

Different partitions of $[N]$ yield different estimands and estimators. To formalize this randomness, we consider \mathcal{A} as a uniform random variable taking values $a \in \mathcal{A}$, that is, $\mathcal{A} \sim U(\mathcal{A})$. Therefore, conditional on Data , the estimand $\theta_{\mathcal{A}}$ is a random variable. In

¹⁵The conventional asymptotic normality of the OLS estimator does not imply validity of the single-split inference procedure automatically, because the distribution of the estimator is random. We need to guarantee that this random distribution converges in probability to a normal (non-random), which requires conditions for splits being “regular” with high probability. We refer to Lemma D.1 in the Supplemental Appendix for the formal result.

what follows, we will mainly target our inference on the median value of θ_A , instead of θ_a , the estimand for a specific partition. Furthermore, different partitions yield different estimators $\hat{\theta}_A$ and approximate distributions for these estimators. Therefore, conditional on Data, estimators $\hat{\theta}_A$, p -values p_A , and intervals $[L_A, U_A]$ are all random variables. We will use quantile aggregation methods to summarize them.

It is useful to recall some definitions of quantiles for discrete variables. For a random variable X with law P_X and k points of support, and index $u \in (0, 1)$, the lower and upper u -quantiles are $\underline{Q}_u(X) := \inf\{x \in \mathbb{R} : P_X(X \leq x) \geq u\}$ and $\overline{Q}_u(X) := \sup\{x \in \mathbb{R} : P_X(X \geq x) \geq 1 - u\}$, respectively. For $w_u := \lfloor uk \rfloor / (\lfloor uk \rfloor + \lceil uk \rceil)$, we define

$$Q_u(X) := w_u \underline{Q}_u(X) + (1 - w_u) \overline{Q}_u(X)$$

as the central quantile.¹⁶ If X is continuous, all three definitions coincide. To define upper, lower, and central medians, we use $u = 1/2$ in the definitions above and M instead of Q_u .

We now formally define the median-aggregated p -value.

DEFINITION 4.2—Median-Aggregated p -Value: The median p -values for testing one-sided alternative hypotheses are

$$p^\pm = M(p_A^\pm | \text{Data}).$$

The two-sided median p -value is $\bar{p} = 2 \min(p^+, p^-)$.

Aggregation using lower median p -values was first proposed by [Meinshausen, Meier, and Bühlmann \(2009\)](#) in the context of split-sample hypothesis testing in linear regression with selection. Here we take the central medians, since they are more likely to behave like regular p -values.¹⁷

We next define the quantile-aggregated point and interval estimators.

DEFINITION 4.3—Quantile-Aggregated Point and Interval Estimators: The median point estimator is

$$\hat{\theta} := M[\hat{\theta}_A | \text{Data}].$$

The β -quantile confidence interval is $[L, U]$, where

$$L := Q_\beta(L_A | \text{Data}), \quad U := Q_{1-\beta}(U_A | \text{Data}), \quad \beta \leq 1/2.$$

We can interpret these definitions as risk-reducing inferential summaries.

LEMMA 4.1—Risk Contraction: Consider any fixed target value $\theta' \in \mathbb{R}$. Then $\hat{\theta}$ is more concentrated near θ' than any single-split generated $\hat{\theta}_A$:

$$E|\hat{\theta} - \theta'| \leq E|\hat{\theta}_A - \theta'|. \quad (4.2)$$

Set $\beta = 1/2$. Then the confidence set $[L, U]$ has the same concentration property:

$$E|U - \theta'| \vee E|L - \theta'| \leq E|U_A - \theta'| \vee E|L_A - \theta'|. \quad (4.3)$$

¹⁶For example, the quantile function in R uses this definition ([R Core Team \(2022\)](#)).

¹⁷For example, the sample lower median of $\{U, 1 - U\}$, $U \sim U(0, 1)$, obeys $P(\underline{M} < \alpha) = 2\alpha$ for $\alpha < 1/2$. In contrast, the central median obeys $P(M < \alpha) < \alpha$ for $\alpha < 1/2$.

Moreover, for any $\beta \in (0, 1/2]$, the width of $[L, U]$ is weakly smaller than the worst-case width of the sets across splits:

$$|U - L| \leq \sup_{a \in \mathcal{A}} |U_a - L_a|. \quad (4.4)$$

Analogous risk contraction properties hold for the mean aggregation, but we focus on medians for robustness reasons.

In what follows, we study the formal inferential guarantees of $[L, U]$. The default choice of β is $1/2$, but we obtain useful theoretical guarantees for other choices $\beta < 1/2$ as well.

Principal Regularity Condition

As the main regularity condition, we assume approximate normality of the split-sample t-statistics:¹⁸

(R1) There exist a sequence of positive constants $\gamma'_N \searrow 0$ as $(n, N - n) \rightarrow \infty$, such that

$$\sup_{z \in \mathbb{R}} |\mathbb{P}\{\widehat{\sigma}_A^{-1}(\widehat{\theta}_A - \theta_A) < z\} - \Phi(z)| \leq \gamma'_N. \quad (4.5)$$

Suppose that the data $\{(Y_i, Z_i, D_i)\}_{i=1}^N$ are generated as i.i.d. copies of (Y, Z, D) . In this case, for any $a \in \mathcal{A}$,

$$\mathbb{P}\{\widehat{\sigma}_A^{-1}(\widehat{\theta}_A - \theta_A) < z\} = \mathbb{E}\left[\frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} 1(\widehat{\sigma}_a^{-1}(\widehat{\theta}_a - \theta_a) < z)\right] = \mathbb{P}\{\widehat{\sigma}_a^{-1}(\widehat{\theta}_a - \theta_a) < z\},$$

because the expression on the right does not depend on $a \in \mathcal{A}$ under the i.i.d. sampling. This observation simplifies the verification of (R1) for least squares type estimators; see Lemma D.1. While the i.i.d. case is our main focus, the main results in this section rely only on the conditions labeled as R, which are likely to hold more generally.

Below, we give various theoretical guarantees for our inferential summaries using this condition and adding more conditions to get stronger results.

4.2. Hypothesis Testing With Multiple Splits

We start the analysis by testing homogeneous hypotheses $\theta_A = \theta_0$, which imply that θ_a does not vary with a . Suppose, for example, that we want to test that the slope of the BLP is zero with probability 1, $\beta_{2A} = 0$, against the alternative $\beta_{2A} > 0$ with positive probability. This problem amounts to both testing the heterogeneity in CATE and the relevance of the ML score S_A as a predictor. Another interesting hypothesis is whether $\beta_{2A} = 1$, with probability 1, that is, whether S_A is well-calibrated and needs no post-processing.

More generally, suppose we are testing the hypothesis

$$H_0 : \theta_A = \theta_0, \quad (4.6)$$

with probability 1, against $H_1^+ : \theta_A > \theta_0$ with positive probability. Testing using the median p -value p^+ will have power against the null when the majority of θ_a 's violate the null, so

¹⁸Here and below, we use the standard error $\widehat{\sigma}_A$ instead of the theoretical standard deviation σ_A in all statements, but we can exchange the two if $\widehat{\sigma}_A/\sigma_A \rightarrow_P 1$, which holds under typical conditions, for example, Lemma D.1.

we can interpret the rejection accordingly. Similarly, we can test against $H_1^- : \theta_A < \theta_0$ or $H_1 : \theta_A \neq \theta_0$ with positive probability using the median p -values p^- or \bar{p} .

Below, we establish the properties of the median p -values under (R1). To get the sharpest results, we can invoke a concentration condition for approximate medians:

(R2) For all $z = \Phi^{-1}(\alpha)$, where the nominal level of interest α is in some closed subinterval of $(0, 1/4)$, and some sequences of positive constants $\gamma_N'' \searrow 0$ and non-negative constants $\varepsilon_N \searrow 0$,

$$\begin{aligned} P(Q_{0.5-\varepsilon_N}(\hat{\sigma}_A^{-1}(\theta_A - \hat{\theta}_A)|\text{Data}) < z) &\leq P(\hat{\sigma}_A^{-1}(\theta_A - \hat{\theta}_A) < z) + \gamma_N'', \\ P(Q_{0.5-\varepsilon_N}(\hat{\sigma}_A^{-1}(\hat{\theta}_A - \theta_A)|\text{Data}) < z) &\leq P(\hat{\sigma}_A^{-1}(\hat{\theta}_A - \theta_A) < z) + \gamma_N''. \end{aligned} \quad (4.7)$$

This condition states that the approximate median over-the-splits t -statistics tend to concentrate more than any single-split t -statistic. This condition is intuitive, but it is hard to give general primitive conditions for it.¹⁹ We believe (4.7) is quite plausible. We verified it for typical values of $\alpha < 1/4$ numerically in various experiments that mimic empirical applications, and were unable to find any counterexample.

THEOREM 4.1—Uniform Validity of Median-Aggregated p -Value: *Suppose that the null hypothesis H_0 in (4.6) holds with probability 1. Let p be either of $\{p^+, p^-, \bar{p}\}$. (i) Suppose that approximate normality (R1) holds; then*

$$P(2p < \alpha) \leq \alpha + o(1),$$

where the $o(1)$ depends only on γ_N'' . (ii) Suppose in addition that the median concentration condition (R2) holds with $\varepsilon_N = 0$; then

$$P(p < \alpha) \leq \alpha + o(1),$$

where the $o(1)$ depends only on γ_N'' and γ_N'' .

Therefore, under the median concentration condition, the median p -values have the standard property. Without the median concentration condition, the median p -values need to be multiplied by 2. However, based on our computational experiments, median p -values are conservative even for the nominal level α (once $\alpha < 1/4$), mainly due to the concentration property holding with $\gamma_N'' < 0$. We therefore do not recommend multiplying by 2; see also DiCiccio, DiCiccio, and Romano (2020) for a similar point. The exact form of $o(1)$, given in the proof, allows us to convert the results into those holding uniformly in a set of probability measures P . The proof of the first result partly relies on the idea of Meinshausen, Meier, and Bühlmann (2009) to use Markov inequality to bound quantiles of an arbitrary collection of marginally uniform random variables.

¹⁹When the t -stats are independent, then their median concentrates in a fixed interval around $1/2$ with probability approaching 1 exponentially fast; see DiCiccio, DiCiccio, and Romano (2020). Therefore, the approximate median concentration condition holds. This happens when m 's are non-overlapping and some further homogeneity conditions hold. On the other hand, suppose the p -values are the same asymptotically; then the inequality in the concentration condition binds, but does not fail. This situation is not common in our context, though.

4.3. Prediction Intervals With Multiple Splits

Outside of the settings with homogeneity, the estimand θ_A is a random variable, and we might be interested in characterizing its typical values. Our first approach serves this purpose, and is connected to conformal/permutation inference that is commonly used for predicting unobserved outcomes (Hoeffding (1952), Vovk, Gammerman, and Shafer (2005), Barber, Candes, Ramdas, and Tibshirani (2022)).

Here our goal is to have a prediction interval for θ_A , and the challenge we face is that θ_A is not observed directly, which places us outside the standard conformal setting. However, for each $a \in \mathcal{A}$, we have a (random) confidence interval $[L_a, U_a]$ that has the covering property:

$$P(\theta_a < L_a) \leq \alpha/2 + o(1), \quad P(\theta_a > U_a) \leq \alpha/2 + o(1). \quad (4.8)$$

This condition is implied by the basic regularity condition (R1) in our context.

We take the quantile-aggregated confidence interval $[L, U]$ as our prediction interval for θ_A .

THEOREM 4.2—Properties of the Prediction Interval: *Suppose that (4.8) holds. Then*

$$P(\theta_A < L) \leq \beta + \alpha/2 + o(1), \quad P(\theta_A > U) \leq \beta + \alpha/2 + o(1),$$

where the $o(1)$ terms are the same as in (4.8). Therefore, $P(\theta_A \in [L, U]) \geq 1 - 2\beta - \alpha - 2o(1)$.

We can use the prediction interval $[L, U]$ to characterize the “majority” of the central values of the random target θ_A that one could get from sample splitting. For this, we set $\beta = 0.25$ and “small” $\alpha = o(1)$; then $MP = [L, U]$ has the property

$$P(\theta_A \in MP) \geq 0.5 - o(1). \quad (4.9)$$

That is, MP contains majority of central values of θ_A . On the other hand, if we set $\beta = 1/2$, we get a median aggregated interval

$$MI = [L, U].$$

In this setting, we can think of MI as predicting the median of θ_A with small margin of error, if $\alpha = o(1)$:

$$P(\theta_A < L) \leq 1/2 + o(1) \quad \text{and} \quad P(\theta_A > U) \leq 1/2 + o(1). \quad (4.10)$$

The latter gives us a useful interpretation of median confidence intervals, and notably this interpretation applies under the weakest possible regularity condition (R1) in our setting.

4.4. Confidence Intervals for Median Parameter With Multiple Splits

Instead of predicting the “majority” of θ_A , we may focus the inference on a single target.

DEFINITION 4.4—Inferential Target: Our inferential target is the median estimand:

$$\theta^* = M[\theta_A | \text{Data}].$$

The choice of the target has the intuitive appeal of representing a typical value. Moreover, in many cases, θ_A will concentrate around its median value θ^* , making it an even more natural target. What follows is the principal regularity condition that covers this concentration scenario.

(R3) For some positive sequences of constants $r_N \searrow 0$ and $\gamma_N''' \searrow 0$ as $(n, N - n) \rightarrow \infty$,

$$P(\widehat{\sigma}_A^{-1} |\theta^* - \theta_A| > r_N) \leq \gamma_N'''. \quad (4.11)$$

The concentration condition above is a convenient property for the interpretability of the inference. Here the rate of concentration of θ_A around θ^* should be faster than the rate $\widehat{\sigma}_A$ of $\widehat{\theta}_A$ estimating θ_A . Thus, the concentration condition implicitly requires the auxiliary set a to be large and the main set m to be small compared to a ; so that randomness in the inferential target is small compared to the size of the estimation error $\widehat{\sigma}_a$ in the main sample, which typically is proportional to $n^{-1/2}$. Otherwise, if n is large relative to $N - n$, then each θ_a would be estimated accurately, but there would be a lot of variation of θ_a across $a \in \mathcal{A}$.²⁰ Condition (R3) is high-level; we demonstrate the plausibility of this condition for the BLP parameter in Appendix B of the Supplemental Appendix using notions of estimation and algorithmic stability.²¹

The following results summarize the properties of the proposed median confidence interval under various conditions.

THEOREM 4.3—Properties of the Confidence Interval for θ^* : *Let $\beta = 1/2$. (i) Suppose that (R1) and (R3) hold. Then,*

$$P(\theta^* \in [L, U]) \geq 1 - 2\alpha - o(1),$$

where $o(1)$ depends only on γ_N' , γ_N''' and r_N . (ii) Suppose in addition that (R2) holds with $\epsilon_N = 2\sqrt{\gamma_N''}$. Then,

$$P(\theta^* \in [L, U]) \geq 1 - \alpha - o(1),$$

where $o(1)$ depends only on γ_N' , γ_N''' , γ_N'' , and r_N . (iii) In either case, the event $\theta^ \in [L, U]$ implies $|\theta^* - \widehat{\theta}| \leq |U - L|$.*

Under the strongest assumptions, the target θ^* is covered with a probability of at least $1 - \alpha - o(1)$. Under the minimal set of assumptions, the coverage probability is $1 - 2\alpha - o(1)$. In our numerical results, the confidence intervals tend to be conservative even under the minimal condition, with coverage exceeding $1 - \alpha$. Therefore, using $1 - \alpha$ as the nominal level is our recommended choice based on this evidence.

²⁰As a practical diagnostic, we recommend the researchers to report the variation of $\widehat{\theta}_a$ across sample splits, in addition to the confidence interval for θ^* . We are grateful to Guido Imbens for this suggestion.

²¹Estimation stability implies that θ_A concentrates around a fixed value θ_* , in which case the median also concentrates around θ_* . Estimation stability follows from the ML proxy S_A being consistent for some fixed proxy function s_* , but not necessarily consistent for the true CATE. This condition can be readily verified using statistical learning theory, as we do in Section 5 for causal learners of CATE. The algorithmic stability condition is strictly weaker than estimation stability, though not as readily available. Our use of these stability criteria is inspired by similar ideas in Wager, Du, Taylor, and Tibshirani (2016), Chernozhukov, Wüthrich, and Zhu (2021), and Chen, Syrgkanis, and Austern (2022), applied to a different context. Appendix B of the Supplemental Appendix discusses all of this further.

5. FURTHER CONSIDERATION: CAUSAL MACHINES THAT LEARN CATE BETTER

Our main proposal so far is to take proxies from any first-stage black-box machine and post-process them to better target CATE and perform inference on functionals of CATE, such as the BLP and GATES. But, can we design the machines to target CATE directly in the first stage? If we can, then the post-processing methods of the previous section would mostly focus on providing inference, and less on correcting biases of the first-stage inputs. Building on [Athey and Imbens \(2016\)](#), we propose two types of such causal machines, and connect them to the emerging literature on orthogonal machine learning, such as [Nie and Wager \(2020\)](#), [Semenova et al. \(2017\)](#), and [Foster and Syrgkanis \(2019\)](#), among others.

5.1. Focusing ML Methods on CATE in Stage 1

We propose two options, taking ideas from our stage 2 analysis to stage 1. Specifically, we can train ML proxies in the auxiliary sample based on either:

- (A) minimizing $w(Z)$ -weighted square prediction errors of Y on B and $(D - p(Z))S$;
- (B) minimizing square prediction errors of YH on BH and S ;

where $B(Z)$ is now a technical “baseline” function of covariates Z , as described below, whose role is to reduce noise in the learning problem.

DEFINITION 5.1—Causal Learners for Stage 1: We can solve either of:

$$(B, S) \in \arg \min_{b \in \mathcal{B}, s \in \mathcal{S}} \sum_{i \in A} w(Z_i) [Y_i - b(Z_i) - \{D_i - p(Z_i)\} s(Z_i)]^2, \quad (\text{A})$$

$$(B, S) \in \arg \min_{b \in \mathcal{B}, s \in \mathcal{S}} \sum_{i \in A} [Y_i H_i - b(Z_i) H_i - s(Z_i)]^2, \quad (\text{B})$$

where $w(Z) = [p(Z)(1 - p(Z))]^{-1}$, and \mathcal{B} and \mathcal{S} are functional parameter spaces.

We can refer to the first causal learner as the weighted residual (WR) learner, and the second causal learner as the HT learner. Examples of parameter spaces include spaces of linear functions generated by a set of dictionary transformations of Z , reproducing kernels, linear combinations of decision trees, neural networks, and others. In (A), the parameter spaces are meant, but not required, to contain the functions $z \mapsto \tilde{b}_0(Z) := b_0(z) + p(z)s_0(z)$ and $z \mapsto s_0(Z)$. In (B), the parameter spaces are meant, but not required, to contain the functions $z \mapsto \bar{b}_0(Z) := b_0(Z) + (1 - p(z))s_0(z)$ and $z \mapsto s_0(Z)$.

Both (A) and (B) improve over the standard predictive learners that predict Y using the best approximation to $E[Y|D, Z]$ in a given class, but not necessarily the best approximation to the CATE $s_0(Z)$ itself, and may be of independent interest. Moreover, the loss functions in (A) and (B) are also helpful for validation purposes, and choosing the best or aggregating classes of ML methods for targeting the CATE function.

The proposal (B) generalizes and refines the strategy of [Athey and Imbens \(2016\)](#) of predicting YH using (a tree form of) S by introducing denoising by B . [Semenova and Chernozhukov \(2021\)](#) developed a related HT strategy that applies series/sieve learners to the denoised HT-transformed outcome, but it explicitly relies on consistent estimation of the regression function, unlike our approach. We further discuss connections of proposal

(A) to the unweighted residual learners of Semenova et al. (2017), Nie and Wager (2020), and others below.²²

THEOREM 5.1—Oracle Properties of the Population Objective Functions: *Suppose that Y , $b(Z)$, $s(Z)$, and $w(Z)$ are square integrable. (1) Then, the expectations of the loss functions in (A) and (B) are*

$$Ew(Z)[Y - b(Z) - (D - p(Z))s(Z)]^2 = E[s_0(Z) - s(Z)]^2 + C_{1b}, \quad (5.1)$$

$$E[YH - b(Z)H - s(Z)]^2 = E[s_0(Z) - s(Z)]^2 + C_{2b}, \quad (5.2)$$

where $C_{1b} := E[w(Z)(\tilde{b}_0(Z) - b(Z))^2] + C_1$ and $C_{2b} := E[w(Z)(\tilde{b}_0(Z) - b(Z))^2] + C_2$ for some constants C_1 and C_2 . (2) Therefore, the minimizers, say $s_\bullet(Z)$, of the left-hand sides of (5.1) and (5.2) over $s \in \mathcal{S}$, if exist, also minimize the oracle loss function $E[s_0(Z) - s(Z)]^2$ over the same set.

Theorem 5.1 shows that the minimizers of the two loss functions provide the best approximation in the mean square sense to the actual CATE function $s_0(Z)$ in the class \mathcal{S} . This property occurs even though we do not observe $s_0(Z)$, and such performance is usually qualified as “oracle.” A sufficient condition for the existence of minimizers is that the set \mathcal{S} be convex and closed in the $L^2(P)$ norm.

We illustrate the benefits of using the causal learning objectives (A) and (B) in Figure 2. In the two panels, we compare the CATE learners derived from the standard predictive Random Forest and Neural Network with the Causal Learners from Random Forests and

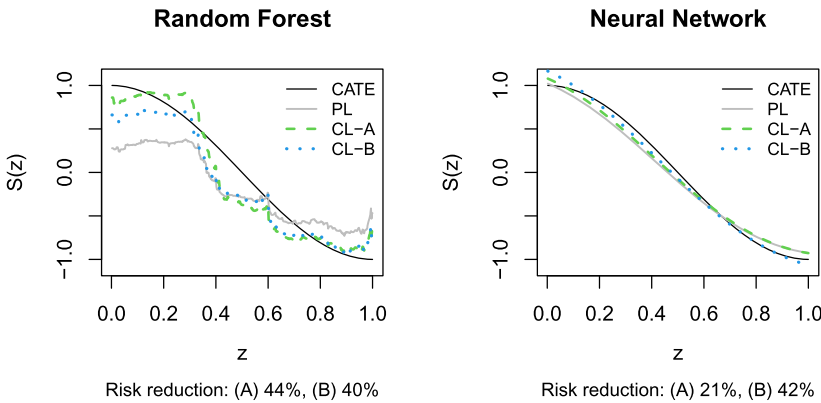


FIGURE 2.—Comparison of predictive machine learners versus causal learners based on (A) and (B). NOTES: The solid black curve is the CATE function $s_0(Z)$, and the solid gray (light) curve is the predictive learner $S(Z)$ (PL) obtained by random forest and neural network. The dashed and dotted curves are estimators $S(Z)$ produced by the causal learners of CATE based on solving objectives (A) and (B) (CL-A and CL-B). The underlying data are generated as $Y_i = b_0(Z_i) + s_0(Z_i)D_i + \xi_i$, where $\xi_i \sim N(0, 1/4)$, $Z_i \sim U(0, 1)$, D_i is Bernoulli with success probability $p(Z) = 0.1 \cdot 1\{Z < 1/2\} + 0.5 \cdot 1\{Z \geq 1/2\}$, $b_0(z) = z$, $s_0(z) = \cos(2\pi z)$, and $i = 1, \dots, 500$.

²²Both of our proposals appeared to be new around the first circulation of this paper as [arXiv:1712.04802](https://arxiv.org/abs/1712.04802). See links to the recent literature below, which proposed related, but different ideas. Relative to the initial version of the paper, the current version contributes with several formal learning properties of (A) and (B). We are grateful to the referees for suggesting that we develop these properties.

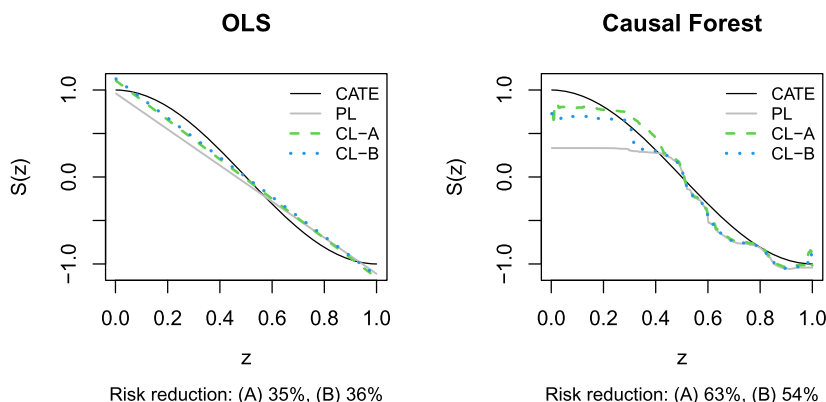


FIGURE 3.—Comparison of OLS and causal forest as predictive learners versus OLS causal learners and forest causal learners based on (A) and (B). NOTES: The solid black curve is the CATE function $s_0(Z)$, and the solid gray (light) curve is the predictive CATE learner $S(Z)$ (PL) obtained by OLS and Causal Random Forest. The dashed and dotted curves are estimators $S(Z)$ produced by the causal learners of CATE based on solving objectives (A) and (B) (CL-A and CL-B). The underlying data are generated in the same way as in Figure 2.

Neural Network that solve the objective functions (A) and (B). We find that the causal learners are better at approximating the CATE function, thereby providing better proxies for CATE. The improvements in the RMSE of approximating the CATE provided by the causal learners range from 21% to 44%.

Likewise, the left panel of Figure 3 shows that we can improve the standard predictive OLS by the Causal OLS that solves the objective functions (A) and (B). Here we estimate linear models in Z for the baseline function and CATE. The improvement in the RMSE of approximating the CATE provided by causal OLS is about 35%. This finding might be of interest to researchers using OLS in empirical work.

Finally, the right panel of Figure 3 shows that one can improve the Causal Forest by a causal boosting step that solves the objective functions (A) and (B) by looking for a shallow forest deviation away from the cross-fitted Causal Forest proxy. The improvements in the RMSE of approximating the true CATE provided by this step are 54%–63%. The explanation for this improvement is that the Causal Forest, while explicitly targeting CATE, actually solves a different objective than (A) or (B).²³ We also verified that this improvement only applies when the propensity score is not constant, like in this example.

COMMENT 5.1—Connections to the Literature: Residual learning like (A), but without weighting by $w(Z)$, appears in the debiased machine learning of Robinson’s partially linear model in Chernozhukov et al. (2017).²⁴ In the nonparametric setting, Nie and Wager (2020) and Semenova et al. (2017) also propose the unweighted version of type (A) objection function. Both of these papers target CATE learning in non-experimental settings. Their proposal does not use weighting by $w(Z)$ as ours and therefore does not provide the

²³In our understanding, it performs a residual learning approach, but not the weighted residual learning approach, which makes the method under-perform relative to the Forest Causal Learner based upon (A) or (B).

²⁴Also called “partialling out,” residual, and orthogonal learning, building upon classical econometric ideas due to Frisch–Waugh–Lovell and Robinson (1988). Note Chernozhukov et al. (2017) presents other strategies as well, with residual learning being just one of them.

best approximation property to the CATE in population. However, it is easy to verify that their proposal provides the best approximation to CATE weighted by $p(Z)(1 - p(Z))$:

$$\min_{s \in \mathcal{S}} \mathbb{E}[p(Z)(1 - p(Z))(s_0(Z) - s(Z))^2].$$

In contrast to our proposal, this weighting gives relatively less attention to units that are either more likely or less likely to be treated than units that are equally likely. While this property is not appealing in general, in pure RCTs with constant propensity score $p(Z)$, the objective function above reduces to the best approximation to CATE. Moreover, weighting by $w(Z)$ in (A) plays a less important role when s_0 can be estimated consistently, as in Nie and Wager (2020), Semenova et al. (2017), and Foster and Syrgkanis (2019). We are interested, however, in the high-dimensional settings in which consistent learning of s_0 might not be possible. Therefore, below we provide formal estimation results for causal learners (A) and (B) under this “agnostic” setting. Our inference methods of Section 4 also apply to the causal learners (A) and (B) used as first-stage proxies.

5.2. Learning Guarantees

The learning guarantees of the causal learners (A) and (B) follow from state-of-the-art statistical learning theory (Liang, Rakhlin, and Sridharan (2015)), in particular through the use of the expected off-set Rademacher complexity (ORC).

Liang, Rakhlin, and Sridharan (2015) define the expected ORC of the function class \mathcal{H} as

$$\mathcal{R}^o(A, \mathcal{H}, c) := \mathbb{E} \sup_{h \in \mathcal{H}} \frac{1}{|A|} \sum_{i \in A} [e_i h(Z_i) - c h(Z_i)^2],$$

where $\{e_i\}$ are i.i.d. Rademacher variables taking values -1 and 1 with probability $1/2$, that are generated independently of the data $\{Z_i\}_{i \in A}$, and $c > 0$ is a positive constant. ORC is a statistical measure of complexity that captures the ability of the functional class to fit i.i.d. Rademacher noise. The more complex the class is, the higher the ORC. As shown in Liang, Rakhlin, and Sridharan (2015), the expected ORC naturally scales like $d/|A|$, where d is the effective dimension of the function class and $|A|$ is the sample size. For example, for linear classes, d is the actual dimension of the linear class; and for VC classes, d scales like the VC dimension. Liang, Rakhlin, and Sridharan (2015) show that ORC is the sharpest characterization of complexity: in particular, the ORC is upper bounded by the standard critical radii of function classes defined in terms of local Rademacher complexity (e.g., Wainwright (2019)) or in terms of the standard uniform covering entropy (Dudley (2000)). This makes the ORC bounds readily available for all function classes used in modern ML.

The following result establishes formal learning guarantees for causal learners in randomized experiments under general settings that do not assume we can learn s_0 consistently.

THEOREM 5.2—Near-Oracle Guarantees for Causal Learners: *Suppose that Y , the elements of \mathcal{B} and \mathcal{S} , and $w(Z)$ are bounded in absolute values by K , and \mathcal{B} and \mathcal{S} are closed, convex, and symmetric sets. The estimator S obtained as a solution of either (A) or (B) is as good as using the best in class approximation, say $s_\bullet(Z)$, up to an error expressed in terms of*

ORC:

$$\begin{aligned}
 0 &\leq \mathbb{E}[s_\bullet(Z) - S(Z)]^2 \\
 &\leq \underbrace{\mathbb{E}[s_0(Z) - S(Z)]^2}_{\text{excess}} - \underbrace{\mathbb{E}[s_0(Z) - s_\bullet(Z)]^2}_{\text{oracle risk}} \leq C_K \mathcal{R}^o(A, \mathcal{H}, c_K), \tag{5.3}
 \end{aligned}$$

where C_K and c_K are positive constants that only depend on K , $\mathcal{H} := 4(w(Z)^2\mathcal{B} + Hw(Z)\mathcal{S})$ for type (A) loss, and $\mathcal{H} := 4(H\mathcal{B} + \mathcal{S})$ for type (B) loss.

The result shows that if the ORC of the functional parameter spaces is small, the excess risk of this estimator relative to the oracle approximation to the CATE is small. Note that the lower bound also bounds the distance of S to the oracle (best-in-class) s_\bullet predictor of the CATE. Therefore, the bounds on the excess risk and distance readily follow from the existing characterization of the ORC. For example, if \mathcal{H} has VC type covering entropy with VC index d , then the ORC is of order $d/|A|$.

Since the “base” functions B ’s play only a noise-reducing rule, we can always select \mathcal{B} to be no more complex than \mathcal{S} . For example, we can use $\mathcal{B} \subseteq \mathcal{S}$ or pre-train B using a separate auxiliary sample, in which case \mathcal{B} is a singleton.²⁵ In either case, learning the technical baseline function does not affect the rate of learning the oracle prediction s_\bullet ; and the ORC is determined solely by the complexity of \mathcal{S} .

Further, we can use losses (A) and (B) for choosing the best ML method. We provide detailed discussion in Appendix C of the Supplemental Appendix.

6. APPLICATION: WHERE ARE NUDGES FOR IMMUNIZATION THE MOST EFFECTIVE?

We apply our methods to an RCT in India that was conducted to improve immunization and provide detailed implementation algorithms. We first describe the setting and results, and then provide the implementation details. Our main specification reports median intervals (MI) from causal learners via Boosting as described in Algorithm 6.2. We also estimate results using predictive ML methods and report them in Appendix G of the Supplemental Appendix as they perform worse than the causal learners. Finally, inferential results are robust to using prediction intervals for the majority values (MP, reported in Appendix H of the Supplemental Appendix). For the sample splitting, we allocate 1/3 of the sample to the main sample.²⁶

6.1. Setting

Immunization is widely believed to be one of the most cost-effective ways to save children’s lives. Much progress has been made in increasing immunization coverage since the 1990s. For example, according to the World Health Organization (WHO), global measles deaths have decreased by 73% from 536,000 estimated deaths in 2000 to 142,000 in 2018. In the last few years, however, global vaccination coverage has remained stuck at around 85% (until the COVID-19 epidemics, when they plummeted). In 2018, 19.7 million children under the age of one year did not receive basic vaccines. Around 60% of these

²⁵This also applies to cross-fitting, which is a better form of sample-splitting for practice.

²⁶We find similar results using 1/2 splits. These results are available upon request.

children lived in ten countries: Angola, Brazil, the Democratic Republic of the Congo, Ethiopia, India, Indonesia, Nigeria, Pakistan, the Philippines, and Vietnam. The WHO estimates that immunization saves 2–3 million deaths every year and that an additional 1.5 million deaths could be averted every year if global vaccination coverage improves (this is comparable to 689,000 deaths from COVID-19 between January and August 2020).²⁷

While most of the early efforts have been devoted to building an immunization infrastructure and ensuring that immunization is available close to people's homes, there is a growing recognition that it is important to also address the demand for immunization. Part of the low demand reflects deep-seated mistrust, but in many cases, parents seem to be perfectly willing to immunize their children. For example, in our data for Haryana, India, among the sample's older siblings who should all have completed their immunization course, 99% had received polio drops, and about 90% had an immunization card. Ninety percent of the parents claimed to believe immunization is beneficial, and 3% claimed to believe it is harmful. However, only 37% of the older children had completed the course and received the measles vaccine, according to their parents (which is likely to be an overestimate), and only 19.4% had done so before the fifteen month of life, when it is supposed to be done between the tenth and the twelfth month. It seems that parents lose steam over the course of the immunization sequence, and nudges could be helpful to boost demand. Indeed, recent literature cited in the [Introduction](#) suggests that “nudges,” such as small incentives, leveraging the social network, SMS, etc., may have a large effect on those services.

In 2017, Esther Duflo, one of the authors of this paper, led a team that conducted a large-scale experiment with the government of Haryana in North India to test various strategies to increase the takeup of immunization services. The government health system rolled out an e-health platform designed by a research team and programmed by an MIT group (SANA health), in which nurses collected data on which child was given which shot at each immunization camp. The platform was implemented in over 2000 villages in seven districts and provides excellent administrative data on immunization coverage.²⁸ From the individual data, we constructed the monthly sum of the number of children eligible for the program (i.e., age 12 months or younger at their first vaccines) who received each particular immunization at a program location. These children were aged between 0 and 15 months. This paper focuses on the number of children who received the measles shot, as it is the last vaccine in the sequence and thus a reliable marker for full immunization.

Before the launch of the interventions, survey data were collected in 912 of those villages using a sample of 15 households with children aged 1–3 per village. The baseline data cover demographic and socio-economic variables and the immunization history of these children, who were too old to be included in the intervention. In these 912 villages, three different interventions (and their variants) were cross-randomized at the village level:

- (1) Small incentives for immunization: parents/caregivers receive mobile phone credit upon bringing children for vaccinations.
- (2) Immunization ambassador intervention: information about immunization camps was diffused through key members of a social network.
- (3) Reminders: a fraction of parents/caregivers who had come at least one time received SMS reminders for pending vaccinations of the children.

²⁷See WHO “10 facts on immunization,” <https://www.who.int/features/factfiles/immunization/facts/en/index1.html>.

²⁸Banerjee et al. (2019) discuss validation data from random checks conducted by independent surveyors.

For each of these interventions, there were several possible variants: incentives were either low or high and either flat or increasing with each shot; the immunization ambassadors were either randomly selected or chosen to be information hubs, using the “gossip” methodology developed by Banerjee et al. (2021), a trusted person, or both; and reminders were sent to either 33% or 66% of the people concerned. Moreover, each intervention was cross-cut, generating 75 possible treatment combinations.

Banerjee et al. (2021) developed and implemented a two-step methodology to identify the most cost-effective and the most effective policy to increase the number of children completing the full course of immunization at the village level and estimate its effects (correcting for bias due to the fact that the policy is found to be the best). First, they used a specific version of LASSO to determine which policies are irrelevant and which policy variants can be pooled together. Second, they obtained consistent estimates of this restricted set of pooled policies using post-LASSO (Chernozhukov, Hansen, and Spindler (2015)). They found that the most cost-effective policy (and the only one to reduce the cost of each immunization compared to the status quo) is to combine information hub ambassadors (trusted or not) and SMS reminders. But the policy that increases immunization the most is the combination of information-hub ambassador, the presence of reminders, and increasing incentives (regardless of levels). This is also the most expensive package, so the government was interested in prioritizing villages: where should they scale up the full package? This is an excellent application of this methodology because there was no strong prior.

6.2. Results

We compare 25 treated villages where this particular policy bundle was implemented with 78 control villages that received neither sloped incentives and social network intervention nor reminder. Our data constitute an approximately balanced monthly panel of the 103 treated and control villages for 12 months (the duration of the intervention). The outcome variable, Y , is the number of children 15 months or younger in a given month in a given village who receive the measles shot. The treatment variable, D , is an indicator of the household being in a village that receives the policy. The covariates, Z , include 36 baseline village-level characteristics such as religion, caste, financial status, marriage and family status, education, and baseline immunization. The propensity score is constant.

Table I shows sample averages in the control and treated groups for some of the variables used in the analysis weighted by village population, as the rest of the analysis. Treatment and control villages have similar baseline characteristics (in particular, the immunization status of the older cohort was similar). The combined treatment was very effective on average. During the course of the intervention, on average, 7.30 children per month aged 15 months or less got the measles shot that completes the immunization sequence in control villages, and 10.08 did so in treatment villages. This is a raw difference of 2.77 or 38% of the control mean. Note that while these effects are not insignificant, we are far from reaching full immunization: The baseline survey suggests that about 38% of children aged 1–3 had received the measles shot at baseline, and 19.4% had received it before they turned 15 months. These estimates imply that the fraction getting their measles shot before 15 months would only go up to 26.7% ($19.4 + 0.38 * 19.4$).

The implementation details for the heterogeneity analysis follow Algorithm 6.1 below, with three characteristics due to the design: we weight village-level estimations by village population, include district–time fixed effects, and cluster standard errors at the village

TABLE I
SELECTED DESCRIPTIVE STATISTICS OF VILLAGES.

	All	Treated	Control
Outcome Variables (Village-Month Level)			
Number of children who completed the immunization schedule	8.239	10.071	7.303
Baseline Covariates–Demographic Variables (Village Level)			
Household financial status (on 1–10 scale)	3.474	3.163	3.633
Fraction Scheduled Caste-Scheduled Tribe (SC/ST)	0.19	0.196	0.187
Fraction Other Backward Caste (OBC)	0.229	0.216	0.236
Fraction Hindu	0.912	0.854	0.942
Fraction Muslim	0.06	0.11	0.034
Fraction Christian	0.001	0.002	0
Fraction Literate	0.797	0.788	0.801
Fraction Single	0.052	0.051	0.053
Fraction Married (living with spouse)	0.516	0.499	0.525
Fraction Married (not living with spouse)	0.003	0.003	0.003
Fraction Divorced or Separated	0.002	0.006	0
Fraction Widow or Widower	0.039	0.036	0.041
Fraction who received Nursery level educ. or less	0.152	0.152	0.151
Fraction who received Class 4 level educ.	0.081	0.08	0.082
Fraction who received Class 9 educ.	0.157	0.162	0.154
Fraction who received Class 12 educ.	0.246	0.223	0.257
Fraction who received Graduate or Other Diploma	0.085	0.078	0.088
Baseline Covariates–Immunization History of Older Cohort (Village Level)			
Number of vaccines administered to pregnant mother	2.265	2.194	2.301
Number of vaccines administered to child since birth	4.464	4.373	4.51
Fraction of children who received polio drops	0.999	1	0.999
Number of polio drops administered to child	2.982	2.984	2.981
Fraction of children who received an immunization card	0.91	0.869	0.931
Fraction of kids who received Measles vaccine by 15 months of age	0.191	0.17	0.201
Fraction of kids who received Measles vaccine at credible locations	0.381	0.362	0.391
Number of Observations			
Villages	103	25	78
Village-Months	843	204	639

level. Table II compares the four ML methods for producing proxy predictors $S(Z_i)$ using the criteria in (3.11) and (3.12). We find that Elastic Net and Neural Network outperform the other methods, with Elastic Net beating Neural Network by a smaller margin than

TABLE II
COMPARISON OF CAUSAL ML METHODS: IMMUNIZATION INCENTIVES.

	Elastic Net	Boosting	Neural Network	Random Forest
Best BLP (Λ)	67.750 [51.491, 82.368]	32.900 [23.246, 44.665]	53.420 [42.516, 67.647]	25.200 [18.328, 34.705]
Best GATES ($\tilde{\Lambda}$)	8.254 [7.329, 9.314]	5.104 [4.27, 6.079]	6.001 [5.087, 6.888]	4.492 [3.339, 5.507]

Note: Medians over 250 splits. Note that we used Neural Network Causal Boosting for all methods, using Algorithm 6.2. The brackets report interquartile ranges for goodness-of-fit statistics.

TABLE III
BLP OF IMMUNIZATION INCENTIVES USING CAUSAL PROXIES.

Elastic Net		Neural Network	
ATE (β_1)	HET (β_2)	ATE (β_1)	HET (β_2)
2.814 (1.087, 4.506) [0.004]	1.047 (0.826, 1.262) [0.000]	2.441 (0.846, 3.979) [0.004]	0.899 (0.685, 1.107) [0.000]

Note: Medians over 250 splits. Median Confidence Intervals ($\alpha = 0.05$) in parentheses. p -Values for the hypothesis that the parameter is equal to zero against the two-sided alternative in brackets.

the other methods. Accordingly, we shall focus on these two methods for the rest of the analysis.

Table III presents the results of the BLP of CATE on the ML proxies. We report estimates of the coefficients β_1 and β_2 , which correspond to the ATE and heterogeneity loading (HET) parameters in the BLP. The ATE estimates in columns 1 and 3 indicate that the package treatment increases the number of immunized children by 2.81 based on elastic net estimates and by 2.44 based on neural network estimates. Reassuringly, these estimates are on either side of the raw difference in means (2.77). Focusing on the HET estimates, we find strong heterogeneity in treatment effects, as indicated by the statistically significant estimates. Moreover, the estimates are close to 1, suggesting that the ML proxies are good predictors of the CATE.

Next, we estimate the GATES by quintiles of the ML proxies. Figure 4 presents the estimated GATES coefficients $\gamma_1 - \gamma_5$ along with joint confidence bands and the ATE estimates. In Table IV, we present the result from the hypothesis test that the difference of the ATE for the most and least affected groups is statistically significant. We find that this difference is 21.60 and 18.13 based on elastic net and neural network methods, respectively, and is statistically significant. Given that the ATE estimates in the whole population are about 2.5, these results suggest a large and potentially policy-relevant heterogeneity.

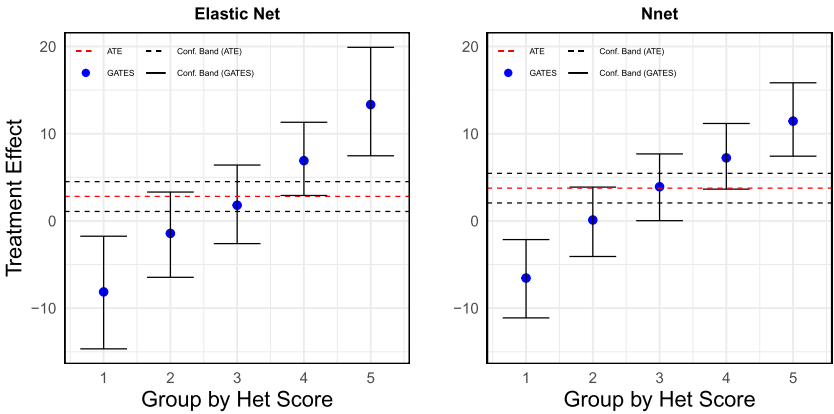


FIGURE 4.—GATES of immunization incentives. Notes: GATES of immunization incentives, based upon causal learners. Median point estimates and median confidence interval ($\alpha = 0.05$) in parentheses, over 250 splits.

TABLE IV
GATES OF 20% MOST AND LEAST AFFECTED GROUPS.

	Elastic Net			Nnet		
	20% Most (G_5)	20% Least (G_1)	Difference	20% Most (G_5)	20% Least (G_1)	Difference
GATE $\gamma_k := \hat{E}[s_0(Z) G_k]$	13.230 (8.219, 18.67) [0.000]	−8.000 (−13.41, −2.574) [0.009]	21.60 (13.70, 29.74) [0.000]	11.210 (7.721, 14.47) [0.000]	−6.551 (−10.37, −2.786) [0.002]	18.13 (12.84, 23.52) [0.000]
Control Mean $:= \hat{E}[b_0(Z) G_k]$	2.431 (1.611, 3.216) [0.000]	12.580 (11.71, 13.50) [0.000]	−9.916 (−11.17, −8.679) [0.000]	1.058 (0.340, 1.750) [0.009]	9.576 (8.824, 10.31) [0.000]	−8.544 (−9.542, −7.589) [0.000]

Note: Medians over 250 splits. Median confidence interval ($\alpha = 0.05$) in parentheses. p -Values for the hypothesis that the parameter is equal to zero against the two-sided alternative in brackets.

The analysis so far reveals very large heterogeneity, with two striking results. First, the results are very large for the most affected villages. In these villages, an average of 13.23 extra children eligible for baseline incentives get the measles vaccines every month (starting from a mean of 2.19 in the elastic net estimation). Second, the impact is *negative* and significant in the least affected villages (an average decline of 8.00 immunizations per month, starting from 12.68 in the elastic net estimation). It looks like in some contexts, the combined package of small incentives, reminders, and persuasion by members of the social network put people off immunization.

Given these large differences, it is important to determine whether this heterogeneity seems to be associated with pre-existing characteristics. To answer this question, we ask what variables are associated with the heterogeneity detected in BLP and GATES via CLAN. Table V reports the CLAN estimates for a selected set of covariates and Tables 1–2 in Appendix F of the Supplemental Appendix do so for the rest of covariates. Regardless of the method used, the estimated differences in the means of most and least affected groups for the number of vaccines to child since birth, number of polio drops to child, the fraction of children receiving measles vaccines by 15 months of age, and fraction of children receiving measles vaccine at credible locations, are negative and statistically significant. Those are various measures of pretreatment immunization levels, all survey-based, that have nothing to do with our measure of impact. These results suggest that the villages with low levels of pretreatment immunization are the most affected by the incentives. These are, in fact, the only variables that consistently pop up from the CLAN. Thus, in this instance, the policy preferred *ex ante* by the government (since it is equality-enhancing) also happens to be the most effective.

While the heterogeneity associated with the baseline immunization rates cannot be causally interpreted (it could always be proxying for other things), it still sheds light on the negative effect we find for the least affected group. Note that this effect is *not* mechanical. Even in the least affected villages, there was a good number of children who did not receive the measles shot, and they were not close to reaching full immunization, where they could not have experienced an increase. It may be that it would have been difficult to vaccinate 13.23 extra children every month, but there was scope to experience an increase in immunization. We had no prior on that the effect would be larger in the villages with the lowest immunization rate. On the contrary, immunization rates could

TABLE V
CLAN OF IMMUNIZATION INCENTIVES.

	Elastic Net			Nnet		
	20% Most (δ_5)	20% Least (δ_1)	Difference ($\delta_5 - \delta_1$)	20% Most (δ_5)	20% Least (δ_1)	Difference ($\delta_5 - \delta_1$)
Number of vaccines to pregnant mother	2.187 (2.115, 2.259)	2.277 (2.212, 2.342)	-0.081 (-0.180, 0.015)	2.174 (2.111, 2.234)	2.285 (2.224, 2.345)	-0.112 (-0.202, -0.028)
	-	-	[0.190]	-	-	[0.019]
Number of vaccines to child since birth	4.077 (3.858, 4.304)	4.639 (4.444, 4.859)	-0.562 (-0.863, -0.260)	4.264 (4.091, 4.434)	4.734 (4.549, 4.900)	-0.490 (-0.739, -0.250)
	-	-	[0.001]	-	-	[0.000]
Fraction of children received polio drops	0.998 (0.995, 1.001)	1.000 (0.997, 1.003)	-0.002 (-0.006, 0.002)	1.000 (1.000, 1.000)	1.000 (1.000, 1.000)	0.000 (0.000, 0.000)
	-	-	[0.683]	-	-	[0.943]
Number of polio drops to child	2.955 (2.935, 2.974)	2.993 (2.976, 3.010)	-0.037 (-0.063, -0.010)	2.965 (2.953, 2.977)	2.998 (2.985, 3.010)	-0.032 (-0.049, -0.016)
	-	-	[0.013]	-	-	[0.000]
Fraction of children received immunization card	0.803 (0.754, 0.851)	0.926 (0.882, 0.969)	-0.121 (-0.187, -0.054)	0.908 (0.881, 0.932)	0.927 (0.898, 0.959)	-0.027 (-0.059, 0.007)
	-	-	[0.001]	-	-	[0.217]
Fraction of children received Measles vaccine by 15 months of age	0.133 (0.097, 0.169)	0.243 (0.209, 0.276)	-0.106 (-0.153, -0.056)	0.126 (0.095, 0.159)	0.260 (0.228, 0.291)	-0.131 (-0.176, -0.085)
	-	-	[0.000]	-	-	[0.000]
Fraction of children received Measles vaccine at credible locations	0.293 (0.246, 0.338)	0.399 (0.358, 0.444)	-0.110 (-0.174, -0.045)	0.289 (0.246, 0.331)	0.433 (0.391, 0.475)	-0.142 (-0.206, -0.084)
	-	-	[0.002]	-	-	[0.000]

Note: Medians over 250 splits. Median confidence interval ($\alpha = 0.05$) in parentheses. p -Values for the hypothesis that the parameter is equal to zero against the two-sided alternatives in brackets.

have been low precisely because parents were more doubtful about immunization. For example, immunization is particularly low in Muslim-majority villages, which is believed to reflect their lack of trust in the health system. There were, therefore, reasons to be genuinely uncertain about where immunization would have had the largest effect.

One possible interpretation of the negative impact in some villages is that villagers were intrinsically motivated to get immunized. The nudging with small incentives and mild social pressure may have backfired, crowding out intrinsic motivation without providing a strong enough extrinsic motivation to act, as in [Gneezy and Rustichini \(2000\)](#). A point estimate of 13.23 extra immunizations per month in the most affected group might seem high: a multiplication by 6.0 of the baseline level (based on the elastic net specification). This increase in immunizations is not inconsistent with the literature: in a set of villages with a very low immunization rate in Rajasthan, [Banerjee et al. \(2010\)](#) find that small incentives increase immunization from 18% to 39% (relative to a treatment that just improves infrastructures, and 6% relative to the control group) in a low immunization region (in the entire sample, not in the places where it is most effective), which was also a very

large increase. Given the restrictions imposed on the data set (only children 1 year or less at their first immunization were included), the data cover children who were 15 months or younger when getting the measles shot. Among the older cohort in the most affected group, only 12.7% of children were vaccinated before 15 months. Taking this as a benchmark for the control group, the estimate would still imply that only 64% of the treatment group was immunized before 15 months: a big improvement but not implausible.

Our last exercise is to compute the cost-effectiveness of the program in various groups. To do so, we compute in each village the average number of immunizations delivered per dollar spent in a month in each group. The dollar spent is the fixed cost to run an immunization program per month (nurse salaries, administrative overheads, etc.) plus the marginal cost of each vaccine multiplied by the number of vaccines administered (incentives distributed to local health workers, vaccines doses, syringes, etc.) in both treatment and control villages, plus the extra cost of running each particular treatment (the cost of the tablets used for recording in all the treatment villages, the cost of contacting and enrolling the ambassadors, and the cost of the incentives). We then estimate the cost-effectiveness in each GATES group as $E[X(1) - X(0)|G_k]$, where X is the immunizations per dollar. $E[X(1) - X(0)|G_k] = E[X|D = 1, G_k] - E[X|D = 0, G_k]$ by the randomization assumption, and we can estimate each of $E[X|D = 1, G_k]$ and $E[X|D = 0, G_k]$ analogously to CLAN, that is, by taking sample averages within treatment groups for each sample split and aggregating over sample splits.

The results are presented in Table VI. They highlight the crucial importance of treatment effect heterogeneity for policy decisions in this context. Overall, as shown in Banerjee et al. (2021), the treatment is not cost-effective compared to the control (the immunization per dollar spent goes down). This analysis reveals that this is driven (not surprisingly) by negative impacts on cost-effectiveness in the groups where it is least effective. However, in the fourth and fifth quintile of cost-effectiveness, we cannot reject that the immunization per dollar spent is the same in the control group and in the treatment group, despite the added marginal cost of the incentives and the vaccines: this is because the fixed cost of running the program is now spread over a larger number of immunizations.

We performed the following additional analysis. First, since the main result in Banerjee et al. (2021) is that the most cost-effective option on average is the combination of SMS plus information hubs, an alternative policy question may therefore be whether there are places where it may be more cost-effective to add the incentives to this cheaper treatment. We replicated the heterogeneity analysis comparing these two treatments and looked at the cost-effectiveness of GATES in these two options. There is also considerable heterogeneity in this comparison (see Figure 2 of the Supplemental Appendix). The results for cost-effectiveness are shown in Table 3 of the Supplemental Appendix. There, again, we find that in the two quintiles where adding incentives is most effective, it would be cost-effective, even compared to an alternative status quo of just having SMS and information hubs.

Second, we reproduce the analysis with the first vaccine in the series to be given at birth (Penta 1). The results are very similar to the results for the measles vaccine, with regions with the lowest immunization experiencing the largest effects. Third, still for the penta 1 vaccination, we apply the method to the comparison of villages with or without increasing incentives for vaccination. Once again, we find significant heterogeneity with the most affected regions being those with low immunization to start with (these results are available upon request).

TABLE VI
COST-EFFECTIVENESS IN GATE QUINTILES.

	Elastic Net			Nnet		
	Mean in Treatment ($\hat{E}[X D = 1, G_k]$)	Mean in Control ($\hat{E}[X D = 0, G_k]$)	Difference	Mean in Treatment ($\hat{E}[X D = 1, G_k]$)	Mean in Control ($\hat{E}[X D = 0, G_k]$)	Difference
Imm. per dollar (G_1)	0.033 (0.030, 0.037)	0.047 (0.045, 0.048)	−0.013 (−0.017, −0.009)	0.033 (0.029, 0.036)	0.047 (0.045, 0.049)	−0.015 (−0.019, −0.010)
	-	-	[0.000]	-	-	[0.000]
Imm. per dollar (G_2)	0.032 (0.027, 0.037)	0.044 (0.042, 0.046)	−0.012 (−0.017, −0.007)	0.035 (0.031, 0.039)	0.044 (0.042, 0.046)	−0.009 (−0.014, −0.005)
	-	-	[0.000]	-	-	[0.000]
Imm. per dollar (G_3)	0.037 (0.032, 0.041)	0.043 (0.040, 0.045)	−0.006 (−0.011, −0.002)	0.038 (0.034, 0.041)	0.043 (0.040, 0.045)	−0.005 (−0.009, −0.001)
	-	-	[0.016]	-	-	[0.028]
Imm. per dollar (G_4)	0.039 (0.036, 0.042)	0.039 (0.036, 0.042)	0.000 (−0.005, 0.004)	0.038 (0.034, 0.041)	0.041 (0.038, 0.044)	−0.003 (−0.008, 0.001)
	-	-	[1.000]	-	-	[0.253]
Imm. per dollar (G_5)	0.036 (0.033, 0.040)	0.035 (0.030, 0.039)	0.002 (−0.004, 0.007)	0.037 (0.033, 0.040)	0.034 (0.031, 0.038)	0.002 (−0.004, 0.007)
	-	-	[1.000]	-	-	[1.000]

Note: Medians over 250 splits. Median confidence interval ($\alpha = 0.05$) in parentheses. p -Values for the hypothesis that the parameter is equal to zero against two-sided alternative in brackets.

6.3. Implementation Details

We describe two general algorithms and provide some specific implementation details for the empirical example.

ALGORITHM 6.1—Inference Algorithm: *The inputs are given by the data $\{(Y_i, D_i, Z_i, p(Z_i))$ on units $i \in [N] = \{1, \dots, N\}$. Fix the number of splits N_s and the significance level α , for example, $N_s = 250$ and $\alpha = 0.05$. Fix a set of ML or Causal ML methods.*

1. *Generate N_s random splits of $[N]$ into the main sample, M , and the auxiliary sample, A . Over each split apply the following steps:*

- Using A , train each ML method and output predictions B and S for M .*
- Optionally, choose the best or aggregate ML methods using the results of Section 5.*
- Estimate the BLP parameters via WR BLP (3.3) or HT BLP (3.5) in M .*
- Estimate the GATES parameters by WR GATES (3.7) or HT GATES (3.9) in M .*
- Estimate the CLAN parameters by taking averages (3.10) in M .*
- Compute the goodness-of-fit measures via (3.13) in M .*

2. *If the winning ML methods were not chosen in Step 1b, median-aggregate the goodness-of-fit measures and choose the best ML methods.*

3. *Compute and report the quantile-aggregated point estimates, p -values, and confidence intervals of Section 4. If Step 2 is used, compute and report the union of these statistics for all winners.*

COMMENT 6.1—Choices: We choose N_S sufficiently large to get enough representative values of the estimates of the target parameter values. In our experience, 250 splits are more than sufficient to obtain stable results in the sense that the point and interval estimates are not sensitive to increasing the number of splits in the empirical application. We followed a version of the algorithm with Step 2 and without Step 1a. Note that it is also possible to choose the best methods using a hold-out sample using either the loss functions of Section 5 or the goodness-of-fit measures of Section 3. More research is needed to determine better practice for choosing the best ML methods.

We implemented our causal learners via a boosting approach that looks for relatively simple deviations from the initial predictive learner to improve CATE predictions. The reason is that this approach performed better in our simulation experiments than directly solving the objective functions (A) and (B) over large parameter spaces. We also observed that this boosting implementation performed better in the empirical example. The performance improvement occurs because the objective functions (A) or (B) tend to be much noisier than the objective functions in predictive learning and harder to tune. The following algorithm summarizes the implementation.

ALGORITHM 6.2—Causal Learner via Boosting: *The inputs are given by data $\{Y_i, D_i, Z_i, p(Z_i)\}$ on units $i \in A \subset \{1, \dots, N\}$. Fix a predictive learner. Fix deviation (or boosting) parameter spaces: \mathcal{B}_Δ and \mathcal{S}_Δ that contain functions $z \mapsto b_\Delta(z)$ and $z \mapsto s_\Delta(z)$, mapping the support of Z to the real line.*

1. *Train the predictive learner on the input, and output the base proxy function $z \mapsto B_A(z)$ and CATE proxy function $z \mapsto S_A(z)$.*
2. *Solve the objective function (A) or (B) for the parameter sets $\mathcal{B} = \{B_A + b_\Delta, b_\Delta \in \mathcal{B}_\Delta\}$ and $\mathcal{S} = \{S_A + s_\Delta, s_\Delta \in \mathcal{S}_\Delta\}$, and update $z \mapsto S_A(z)$ and $z \mapsto B_A(z)$ to be the solution.*
3. *Optionally, iterate on Step 2 a few times.*

Ideally, Step 1 of the algorithm should make use of cross-fitting, but we present a simplified version for clarity.

COMMENT 6.2—Choices for Computational Experiments: We use simple deviation parameter spaces with low complexity (small ORC). For example, in the computational experiments reported in Figures 2 and 3, we made the following choices: for Causal Neural Network (NN) Learner, we used shallow, regularized NN as deviation spaces when the predictive learner was NN; for Forest Causal Learner, we used the shallow forest as the deviation space when the predictive learner was Random Forest; for Causal OLS Learner, we used linear deviation space when the predictive learner was OLS. Finally, when we used Causal Forest as the predictive learner, we used the shallow forest as the deviation space to obtain the Forest Causal Learner.²⁹ We have also experimented with hybrid versions, for example, using one type of predictive learner and a different type of booster (e.g., RF plus NN as a causal boost or Elastic Net plus NN as a causal boost).

COMMENT 6.3—Choices for Empirical Example: We used simple, regularized neural networks as deviation spaces in all results. The number of neurons was kept less than or

²⁹As this may sound confusing, we note that Forest Causal Learners (FCL) differ from Causal Random Forests (CRF) in stratified experiments, as FCL are based on weighted residualization whereas CRF are based on unweighted residualization of Nie and Wager (2020). Therefore, we use a slightly different name “Forest Causal Learner” rather than “Causal Forest” to distinguish the proposal.

equal to 10 and was chosen based on cross-validating the objective function of type (A) over auxiliary data subsamples. We used this choice regardless of the predictive ML as the starter. Causal Learners constructed in this way improved the performance of each predictive ML method, raising the goodness-of-fit metrics by 5–10% in relative terms. However, they did not have any qualitative impact on empirical results (we report the results for predictive learners in Appendix G of the Supplemental Appendix).

COMMENT 6.4—Predictive ML Methods: We considered four ML methods to estimate the proxy predictors: elastic net, boosted trees, neural network with feature extraction, and random forest. The ML methods are implemented in R using the package *caret* (Kuhn (2008)). The names of the elastic net, boosted tree, neural network with feature extraction, and random forest methods in *caret* are *glmnet*, *gbm*, *pcaNNet*, and *rf*, respectively. For each split of the data, we choose the tuning parameters separately for $B(z)$ and $S(z)$ based on mean squared error estimates of repeated 2-fold cross-validation, except for random forest, for which we use the default tuning parameters to reduce the computational time.³⁰ In tuning and training the ML methods, we use only the auxiliary sample. In all the methods, we rescale the outcomes and covariates to be between 0 and 1 before training.

7. CONCLUSION AND EXTENSIONS

We propose to focus inference on key features of heterogeneous effects in randomized experiments, and develop the corresponding methods. These key features include best linear predictors of the effects and average effects sorted by groups, as well as average characteristics of most and least affected units. Our approach is valid in high-dimensional settings, where the effects are estimated by machine learning methods. The main advantage of our approach is its agnostic nature; it avoids making strong assumptions. Estimation and inference rely on data splitting, where the latter allows us to avoid overfitting and all kinds of non-regularities. Our inference aggregates the results across many splits, reducing the replication risks, and could be of independent interest. An empirical application illustrates the practical use of the approach.

Our hope is that applied researchers use the method to discover whether there is heterogeneity in their data in a disciplined way. A researcher might be concerned about the application of our method due to the possible power loss induced by sample splitting. This power loss is the price to pay when the researcher is not certain or willing to fully specify the form of the heterogeneity prior to conducting the experiment. Thus, if the researcher has a well-defined pre-analysis plan that spells out a small number of heterogeneity groups in advance, then there is no need of splitting the sample.³¹ However, this situation is not common. In general, the researchers might not be able to fully specify the form of the heterogeneity due to lack of information, economic theory, or willingness to take a stand at the early stages of the analysis. They might also face data limitations

³⁰We have the following tuning parameters for each method: Elastic Net: alpha (Mixing Percentage), lambda (Regularization Parameter), Boosted trees: n.trees (Number of Boosting Iterations), interaction.depth (Max Tree Depth), shrinkage (Shrinkage), n.minobsinnode (Min. Terminal Node Size), size (Number of Hidden Units), decay (Weight Decay), mtry (Number of Randomly Selected Predictors).

³¹More generally, the plan needs to specify a parametric form for the heterogeneity as a low-dimensional function of pre-specified covariates (e.g., Chernozhukov, Fernandez-Val, and Luo (2015)). In this case, ML tools can still be used to efficiently estimate the CATEs in the presence of control variables but are not required to detect heterogeneity (Belloni et al. (2017), Chernozhukov et al. (2017)).

that preclude the availability of the desired covariates. Here we recommend the use of our method to avoid overfitting and p-hacking, and impose discipline to the heterogeneity analysis at the cost of some power loss due to sample splitting.³² If discovering and exploiting heterogeneity in treatment effect is a key goal of the research, the researcher should indeed plan for larger sample sizes (relative to just testing whether the treatment has an effect), but the required sample size remains within the realm of what is feasible in the field. In many applications we are aware of, there was apparent heterogeneity according to some covariates of interest, but the disciplined ML heterogeneity exercise found no systematic difference. This could be because this heterogeneity was a fluke, or because the method does not have the power to detect it in a small sample. In any case, what this experience suggests is that one should not rely on ex post heterogeneous effects in such cases.

The application to immunization in India is of substantive interest. Our findings suggest that a combination of small incentives, relay by information hub, and SMS reminders can have very large effect on vaccine takeup in some villages where immunization was low at baseline, and even be as cost-effective than the status quo, but can also backfire in other places. This suggests that these types of strategies need to be piloted in the relevant context before being rolled out, and that heterogeneity needs to be taken into account.

Our inference approach generalizes to any problem of the following sort, studied in [Semenova and Chernozhukov \(2021\)](#) using more conventional inference approaches. Suppose we can construct an *unbiased signal* \tilde{Y} such that

$$E[\tilde{Y}|Z] = s_0(Z),$$

where $s_0(Z)$ is now a generic target function. Let $S(Z)$ denote an ML proxy for $s_0(Z)$. This setting generalizes the CATE inference problem of this paper.³³ We can then apply our inferential framework to this setting. The setting covers a variety of causal effects of interest, for example, inference on conditional average causal derivatives, conditional average effects from transporting, and distributional shifts of covariates, as we further explain in Appendix E of the Supplemental Appendix.

REFERENCES

- ABADIE, ALBERTO (2005): “Semiparametric Difference-in-Differences Estimators,” *The Review of Economic Studies*, 72 (1), 1–19. [1129]
- ABADIE, ALBERTO, MATTHEW M. CHINGOS, AND MARTIN R. WEST (2017): “Endogenous Stratification in Randomized Experiments,” Technical report, National Bureau of Economic Research. [1130,1135]
- AGARWAL, ANISH, ABDULLAH ALOMAR, ROMAIN COSSON, DEVAVRAT SHAH, AND DENNIS SHEN (2020): *Synthetic Interventions*. [1125]
- ALATAS, VIVI, ARUN G. CHANDRASEKHAR, MARKUS MOBIUS, BENJAMIN A. OLKEN, AND CINDY PALADINES (2019): “When Celebrities Speak: A Nationwide Twitter Experiment Promoting Vaccination in Indonesia,” Technical report, National Bureau of Economic Research. [1125]

³²It is not clear whether this loss is real, though, as we are not aware of any alternative method that avoids sample splitting and that works at the same level of agnosticism as ours. In a previous version of the paper, we provided a numerical example using a simple parametric model where standard methods without sample splitting are available. We find that the extent of the power loss for not using the parametric form of the heterogeneity roughly corresponds to reducing the sample size by half in a test for the presence of heterogeneity, although the exact comparison depends on features of the data generating process.

³³The unbiased signals arise from multiplying an outcome with a Riesz representer for the effect of interest. In the CATE setting, the representer is the Horwitz–Thompson transform.

- ANGRIST, JOSHUA D., AND JORN-STEFFAN PISCHKE (2008): *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press. [1132]
- ATHEY, SUSAN, AND GUIDO IMBENS (2016): "Recursive Partitioning for Heterogeneous Causal Effects," *Proceedings of the National Academy of Sciences*, 113 (27), 7353–7360. [1123,1129,1135,1145]
- ATHEY, SUSAN, AND STEFAN WAGER (2021): "Policy Learning With Observational Data," *Econometrica*, 89 (1), 133–161. [1135]
- BANERJEE, ABHIJIT, ARUN CHANDRASEKHAR, ESTHER DUFLO, SURESH DALPATH, JOHN FLORETTA, MATTHEW JACKSON, LOZA FRANCINE, HARINI KANNAN, AND ANNA SCHRIMPF (2021): "Inference on Winners," Technical Report 28726, NBER Working Paper. [1125,1151,1156]
- BANERJEE, ABHIJIT, ARUN CHANDRASEKHAR, ESTHER DUFLO, SURESH DALPATH, JOHN FLORETTA, MATTHEW JACKSON, HARINI KANNAN, ANNA SCHRIMPF, AND MAHESH SHRESTHA (2019): "Leveraging the Social Network Amplifies the Effectiveness of Interventions to Stimulate Take up of Immunization." [1122, 1124,1125,1150]
- BANERJEE, ABHIJIT VINAYAK, ESTHER DUFLO, RACHEL GLENNERSTER, AND DHURVA KOTHARI (2010): "Improving Immunisation Coverage in Rural India: Clustered Randomised Controlled Evaluation of Immunisation Campaigns With and Without Incentives," *BMJ*, 340. ISSN 0959-8138. [1124,1155]
- BARBER, RINA FOYGEL, EMMANUEL J. CANDES, AADITYA RAMDAS, AND RYAN J. TIBSHIRANI (2022): "Conformal Prediction Beyond Exchangeability." [1143]
- BARNARD, GEORGE A. (1974): "Discussion of "Cross-Validatory Choice and Assessment of Statistical Predictions" by Stone," *Journal of the Royal Statistical Society. Series B (Methodological)*, 133–135. [1130]
- BASSANI, DIEGO G., PAUL ARORA, KERRI WAZNY, MICHELLE F. GAFFEY, LINDSEY LENTERS, AND ZULFIQAR A. BHUTTA (2013): "Financial Incentives and Coverage of Child Health Interventions: A Systematic Review and Meta-Analysis," *BMC Public Health*, 13 (S3), S30. [1124]
- BELLONI, ALEXANDRE, VICTOR CHERNOZHUKOV, IVÁN FERNÁNDEZ-VAL, AND CHRISTIAN HANSEN (2017): "Program Evaluation and Causal Inference With High-Dimensional Data," *Econometrica*, 85 (1), 233–298. [1122,1159]
- BELLONI, ALEXANDRE, VICTOR CHERNOZHUKOV, AND CHRISTIAN HANSEN (2014): "Inference on Treatment Effects After Selection Amongst High-Dimensional Controls," *Review of Economic Studies*, 81, 608–650. [1122]
- BELLONI, ALEXANDRE, VICTOR CHERNOZHUKOV, AND KENGO KATO (2013): "Uniform Post Selection Inference for l₁ Regression Models," arXiv preprint. arXiv:1304.0282. [1129]
- BICKEL, PETER J., YA'ACOV RITOV, AND ALEXANDRE B. TSYBAKOV (2009): "Simultaneous Analysis of Lasso and Dantzig Selector," *Annals of Statistics*, 37 (4), 1705–1732. [1127]
- BISHOP, CHRISTOPHER M., AND NASSER M. NASRABADI (2006): *Pattern Recognition and Machine Learning*, Vol. 4. Springer. [1122]
- CHEN, QIZHAO, VASILIS SYRGKANIS, AND MORGANE AUSTERN (2022): "Debiased Machine Learning Without Sample-Splitting for Stable Estimators," arXiv preprint. arXiv:2206.01825. [1144]
- CHERNOZHUKOV, VICTOR, DENIS CHETVERIKOV, MERT DEMIRER, ESTHER DUFLO, CHRISTIAN HANSEN, WHITNEY NEWAY, AND JAMES ROBINS (2017): "Double/Debiased Machine Learning for Treatment and Structural Parameters," *The Econometrics Journal*. [1122,1147,1159]
- CHERNOZHUKOV, VICTOR, MERT DEMIRER, ESTHER DUFLO, AND IVÁN FERNÁNDEZ-VAL (2025): "Supplement to 'Fisher-Schultz Lecture: Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, With an Application to Immunization in India'," *Econometrica Supplemental Material*, 93, <https://doi.org/10.3982/ECTA19303>. [1126]
- CHERNOZHUKOV, VICTOR, IVÁN FERNANDEZ-VAL, AND YE LUO (2015): "The Sorted Effects Method: Discovering Heterogeneous Effects Beyond Their Averages," ArXiv e-prints [1129,1159]
- CHERNOZHUKOV, VICTOR, CHRISTIAN HANSEN, AND MARTIN SPINDLER (2015): "Post-Selection and Post-Regularization Inference in Linear Models With Very Many Controls and Instruments," *American Economic Review: Papers and Proceedings*, 105 (5), 486–490. [1151]
- CHERNOZHUKOV, VICTOR, WHITNEY NEWAY, AND RAHUL SINGH (2018): "Debiased Machine Learning of Global and Local Parameters Using Regularized Riesz Representers," arXiv preprint. arXiv:1802.08667. [1130]
- CHERNOZHUKOV, VICTOR, WHITNEY K. NEWAY, AND RAHUL SINGH (2021): "A Simple and General Debiased Machine Learning Theorem With Finite Sample Guarantees," arXiv preprint. arXiv:2105.15197. [1130]
- CHERNOZHUKOV, VICTOR, KASPAR WÜTHRICH, AND YINCHU ZHU (2021): "An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls," *Journal of the American Statistical Association*, 116 (536), 1849–1864. [1144]
- COX, DAVID R. (1975): "A Note on Data-Splitting for the Evaluation of Significance Levels," *Biometrika*, 62 (2), 441–444. [1130]

- CREPON, BRUNO, ESTHER DUFLO, HUILLERY ELISA, WILLIAM PARIENTE, JULIETTE SEBAN, AND PAUL-ARMAND VEILLON (2021): “Cream Skimming and the Comparison Between Social Interventions Evidence From Entrepreneurship Programs for at-Risk Youth in France,” Report. [1123,1124]
- CRUMP, RICHARD K., V. JOSEPH HOTZ, GUIDO W. IMBENS, AND OSCAR A. MITNIK (2008): “Nonparametric Tests for Treatment Effect Heterogeneity,” *The Review of Economics and Statistics*, 90 (3), 389–405. [1129]
- DAVIS, JONATHAN MV, AND SARA B. HELLER (2020): “Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs,” *Review of Economics and Statistics*, 102 (4), 664–677. [1123,1130]
- DERYUGINA, TATYANA, GARTH HEUTEL, NOLAN H. MILLER, DAVID MOLITOR, AND JULIAN REIF (2019): “The Mortality and Medical Costs of Air Pollution: Evidence From Changes in Wind Direction,” *American Economic Review*, 109 (12), 4178–4219. [1123]
- DEZEURE, RUBEN, PETER BÜHLMANN, AND CUN-HUI ZHANG (2016): “High-Dimensional Simultaneous Inference With the Bootstrap,” arXiv preprint. arXiv:1606.03940. [1129]
- DI CICCIO, CYRUS J., THOMAS J. DI CICCIO, AND JOSEPH P. ROMANO (2020): “Exact Tests via Multiple Data Splitting,” *Statistics & Probability Letters*, 166, 108865. [1142]
- DOMEK, GRETCHEN J., INGRID L. CONTRERAS-ROLDAN, SEAN T. O’LEARY, SHEANA BULL, ANNA FURNISS, ALLISON KEMPE, AND EDWIN J. ASTURIAS (2016): “Sms Text Message Reminders to Improve Infant Vaccination Coverage in Guatemala: A Pilot Randomized Controlled Trial,” *Vaccine*, 34 (21), 2437–2443. [1124]
- DUDLEY, RICHARD M. (2000): *Uniform Central Limit Theorems*. Cambridge Studies in Advanced Mathematics. [1148]
- DUFLO, ESTHER, RACHEL GLENNERSTER, AND MICHAEL KREMER (2007): “Using Randomization in Development Economics Research: A Toolkit,” *Handbook of development economics*, 4, 3895–3962. [1122]
- FAN, QINGLIANG, YU-CHIN HSU, ROBERT P. LIEI, AND YICHONG ZHANG (2022): “Estimation of Conditional Average Treatment Effects With High-Dimensional Data,” *Journal of Business & Economic Statistics*, 40 (1), 313–327. [1130]
- FITHIAN, WILLIAM, DENNIS SUN, AND JONATHAN TAYLOR (2014): “Optimal Inference After Model Selection,” arXiv preprint. arXiv:1410.2597. [1130]
- FOSTER, DYLAN J., AND VASILIS SYRGKANIS (2019): “Orthogonal Statistical Learning,” arXiv preprint. arXiv:1901.09036. [1130,1145,1148]
- FREEDMAN, DAVID A. (2008): “On Regression Adjustments to Experimental Data,” *Advances in Applied Mathematics*, 40 (2), 180–193. ISSN 0196-8858. [1122]
- FRIEDMAN, JEROME, TREVOR HASTIE, AND ROBERT TIBSHIRANI (2001): *The Elements of Statistical Learning*, Vol. 1. New York: Springer series in statistics. [1122,1134]
- GENOVESE, CHRISTOPHER, AND LARRY WASSERMAN (2008): “Adaptive Confidence Bands,” *The Annals of Statistics*, 36 (4), 875–905. [1123,1127]
- GIBSON, DUSTIN G., BENARD OCHIENG, E. WANGECI KAGUCIA, JOYCE WERE, KYLA HAYFORD, LAWRENCE H. MOULTON, ORIN S. LEVINE, FRANK ODHIAMBO, KATHERINE L. O’BRIEN, AND DANIEL R. FEIKIN (2017): “Mobile Phone-Delivered Reminders and Incentives to Improve Childhood Immunisation Coverage and Timeliness in Kenya (m-Simu): A Cluster Randomised Controlled Trial,” *The Lancet Global Health*, 5 (4), e428–e438. [1124]
- GINÉ, EVARIST, AND RICHARD NICKL (2010): “Confidence Bands in Density Estimation,” *Ann. Statist.*, 38 (2), 1122–1170. ISSN 0090-5364. [1128]
- GNEEZY, URI, AND ALDO RUSTICHINI (2000): “A Fine Is a Price,” *The Journal of Legal Studies*, 29 (1), 1–17. [1155]
- GOODFELLOW, IAN, YOSHUA BENGIO, AND AARON COURVILLE (2016): *Deep Learning*. MIT press. [1122]
- HANSEN, CHRISTIAN, DAMIAN KOZBUR, AND SANJOG MISRA (2017): “Targeted Undersmoothing,” arXiv preprint. arXiv:1706.07328. [1129]
- HARTIGAN, JOHN A. (1969): “Using Subsample Values as Typical Values,” *Journal of the American Statistical Association*, 64 (328), 1303–1317. [1130]
- HASTIE, TREVOR, ROBERT TIBSHIRANI, AND MARTIN WAINWRIGHT (2015): *Statistical Learning With Sparsity*. Monographs on Statistics and Applied Probability, Vol. 143. Boca Raton, FL: CRC Press. The lasso and generalizations. ISBN 978-1-4987-1216-3. [1122]
- HIRANO, KEISUKE, GUIDO W. IMBENS, AND GEERT RIDDER (2003): “Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score,” *Econometrica*, 71 (4), 1161–1189. ISSN 0012-9682. [1129]
- HOEFFDING, WASSILY (1952): “The Large-Sample Power of Tests Based on Permutations of Observations,” *The Annals of Mathematical Statistics*, 169–192. [1143]
- HUSSAM, RESHMAAN, NATALIA RIGOL, AND BENJAMIN N. ROTH (2022): “Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field,” *American Economic Review*, 112 (3), 861–898. [1123]

- IMAI, KOSUKE, AND MARC RATKOVIC (2013): “Estimating Treatment Effect Heterogeneity in Randomized Program Evaluation,” *The Annals of Applied Statistics*, 7 (1), 443–470. [1129]
- IMBENS, GUIDO W., AND DONALD B. RUBIN (2015): *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press. [1122]
- JAMES, GARETH, DANIELA WITTEN, TREVOR HASTIE, AND ROBERT TIBSHIRANI (2021): *An Introduction to Statistical Learning—With Applications in r* (Second Ed.). Springer Texts in Statistics. New York: Springer. [1122]
- JOHRI, MIRA, MYRIAM CIELO PÉREZ, CATHERINE ARSENAULT, JITENDAR K. SHARMA, NITIKA PANT PAI, SMRITI PAHWA, AND MARIE-PIERRE SYLVESTRE (2015): “Strategies to Increase the Demand for Childhood Vaccination in Low-and Middle-Income Countries: A Systematic Review and Meta-Analysis,” *Bulletin of the World Health Organization*, 93, 339–346. [1124]
- KARING, ANNE (2018): “Social Signaling and Childhood Immunization: A Field Experiment in Sierra Leone,” University of California, Berkeley Working Paper. [1124]
- KENNEDY, EDWARD H. (2020): “Optimal Doubly Robust Estimation of Heterogeneous Causal Effects,” arXiv preprint. arXiv:2004.14497. [1130]
- KISH, LESLIE, AND MARTIN RICHARD FRANKEL (1974): “Inference From Complex Samples,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1–37. [1130]
- KITAGAWA, TORU, AND ALEKSEY TETENOV (2018): “Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice,” *Econometrica*, 86 (2), 591–616. [1135]
- KUHN, MAX (2008): “Caret Package,” *Journal of Statistical Software*, 28 (5), 1–26. [1159]
- LEE, JASON D., DENNIS L. SUN, YUEKAI SUN, AND JONATHAN E. TAYLOR (2016): “Exact Post-Selection Inference, With Application to the Lasso,” *The Annals of Statistics*, 44 (3), 907–927. [1130]
- LIANG, TENGYUAN, ALEXANDER RAKHLIN, AND KARTHIK SRIDHARAN (2015): “Learning With Square Loss: Localization Through Offset Rademacher Complexity,” in *Conference on Learning Theory*. PMLR, 1260–1285. [1148]
- LOW, MARK G. et al. (1997): “On Nonparametric Confidence Intervals,” *The Annals of Statistics*, 25 (6), 2547–2554. [1127]
- MEINSHAUSEN, NICOLAI, LUKAS MEIER, AND PETER BÜHLMANN (2009): “P-Values for High-Dimensional Regression,” *Journal of the American Statistical Association*, 104 (488), 1671–1681. [1130,1140,1142]
- MOSTELLER, FREDERICK, AND JOHN WILDER TUKEY (1977): *Data Analysis and Regression: A Second Course in Statistics*. Addison–Wesley Series in Behavioral Science: Quantitative Methods. [1130]
- MURPHY, KEVIN P. (2012): *Machine Learning: A Probabilistic Perspective*. MIT press. [1122]
- NEKIPELOV, DENIS, VIRA SEMENOVA, AND VASILIS SYRGKANIS (2022): “Regularised Orthogonal Machine Learning for Nonlinear Semiparametric Models,” *The Econometrics Journal*, 25 (1), 233–255. [1130]
- NEYMAN, JERSEY (1923): “Sur les applications de la theorie des probabilites aux experiences agricoles: essai des principes (Masters Thesis); Justification of Applications of the Calculus of Probabilities to the Solutions of Certain Questions in Agricultural Experimentation. Excerpts English Translation (Reprinted),” *Stat Sci*, 5, 463–472. [1126]
- NIE, XINKUN, AND STEFAN WAGER (2020): “Quasi-Oracle Estimation of Heterogeneous Treatment Effects,” *Biometrika*, 108 (2), 299–319. [1130,1132,1145-1148,1158]
- OYO-ITA, ANGELA, CHARLES S. WIYSONGE, CHIOMA ORINGANJE, CHUKWUEMEKA E. NWACHUKWU, OLABISI ODUWOLE, AND MARTIN M. MEREMIKWU (2016): “Interventions for Improving Coverage of Childhood Immunisation in Low-and Middle-Income Countries,” *Cochrane Database of Systematic Reviews*, 7. [1124]
- R CORE TEAM (2022): “R: A Language and Environment for Statistical Computing,” Vienna, Austria: R Foundation for Statistical Computing, <https://www.R-project.org/>. [1140]
- REGAN, ANNETTE K., LAUREN BLOOMFIELD, IAN PETERS, AND PAUL V. EFFLER (2017): “Randomized Controlled Trial of Text Message Reminders for Increasing Influenza Vaccination,” *The Annals of Family Medicine*, 15 (6), 507–514. [1125]
- RINALDO, ALESSANDRO, LARRY WASSERMAN, MAX G’SSELL, JING LEI, AND RYAN TIBSHIRANI (2016): “Bootstrapping and Sample Splitting for High-Dimensional, Assumption-Free Inference,” arXiv preprint. arXiv:1611.05401. [1130]
- ROBINSON, PETER M. (1988): “Root- N -Consistent Semiparametric Regression,” *Econometrica*, 56 (4), 931–954. ISSN 0012-9682. [1147]
- ROSENBAUM, PAUL R., AND DONALD B. RUBIN (1983): “The Central Role of the Propensity Score in Observational Studies for Causal Effects,” *Biometrika*, 70 (1), 41–55. ISSN 0006-3444. [1126]
- RUBIN, DONALD B. (1974): “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies,” *Journal of educational Psychology*, 66 (5), 688. [1126]
- SEMENOVA, VIRA, AND VICTOR CHERNOZHUKOV (2021): “Debiased Machine Learning of Conditional Average Treatment Effects and Other Causal Functions,” *The Econometrics Journal*, 24 (2), 264–289. [1129,1130,1145,1160]

- SEMEANOVA, VIRA, MATT GOLDMAN, VICTOR CHERNOZHUKOV, AND MATT TADDY (2017): “Estimation and Inference on Heterogeneous Treatment Effects in High-Dimensional Dynamic Panels,” arXiv preprint. [arXiv:1712.09988](https://arxiv.org/abs/1712.09988). [1130,1132,1145–1148]
- STONE, CHARLES J. (1982): “Optimal Global Rates of Convergence for Nonparametric Regression,” *Ann. Statist.*, 10 (4), 1040–1053. ISSN 0090-5364. [1127]
- UDDIN, MD JASIM, MD SHAMSUZZAMAN, LILY HORNG, ALAIN LABRIQUE, LAVANYA VASUDEVAN, KELSEY ZELLER, MRIDUL CHOWDHURY, CHARLES P. LARSON, DAVID BISHAI, AND NURUL ALAM (2016): “Use of Mobile Phones for Improving Vaccination Coverage Among Children Living in Rural Hard-to-Reach Areas and Urban Streets of Bangladesh,” *Vaccine*, 34 (2), 276–283. [1125]
- UNICEF (2019): “20 Million Children Miss out on Lifesaving Measles, Diphtheria and Tetanus Vaccines in 2018,” <https://www.unicef.org/eca/press-releases/20-million-children-miss-out-lifesaving-measles-diphtheria-and-tetanus-vaccines-2018>. [1124]
- VOVK, VLADIMIR, ALEXANDER GAMMERMAN, AND GLENN SHAFER (2005): *Algorithmic Learning in a Random World*. Springer Science & Business Media. [1143]
- WAGER, STEFAN, AND SUSAN ATHEY (2018): “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests,” *Journal of the American Statistical Association*, 113 (523), 1228–1242. [1129,1130]
- WAGER, STEFAN, WENFEI DU, JONATHAN TAYLOR, AND ROBERT J. TIBSHIRANI (2016): “High-Dimensional Regression Adjustments in Randomized Experiments,” *Proceedings of the National Academy of Sciences*, 113 (45), 12673–12678. [1144]
- WAINWRIGHT, MARTIN J. (2019): *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, Vol. 48. Cambridge University Press. [1148]
- WAKADHA, HOTENZIA, SUBHASH CHANDIR, ELIJAH VICTOR WERE, ALAN RUBIN, DAVID OBOR, ORIN S. LEVINE, DUSTIN G. GIBSON, FRANK ODHIAMBO, KAYLA F. LASERSON, AND DANIEL R. FEIKIN (2013): “The Feasibility of Using Mobile-Phone Based sms Reminders and Conditional Cash Transfers to Improve Timely Immunization in Rural Kenya,” *Vaccine*, 31 (6), 987–993. [1124]
- WASSERMAN, LARRY (2016): “Machine Learning Overview,” in *Becker-Friedman Institute, Conference on ML in Economics*. [1122]
- WASSERMAN, LARRY, AND KATHRYN ROEDER (2009): “High Dimensional Variable Selection,” *Annals of statistics*, 37 (5A), 2178. [1130]
- WRIGHT, MARVIN N., AND ANDREAS ZIEGLER (2017): “Ranger: A Fast Implementation of Random Forests for High Dimensional Data in C++ and R,” *Journal of Statistical Software*, 77 (1), 1–17. [1133]
- ZHAO, QINGYUAN, DYLAN S. SMALL, AND ASHKAN ERTEFAIE (2017): “Selective Inference for Effect Modification via the Lasso,” arXiv preprint. [arXiv:1705.08020](https://arxiv.org/abs/1705.08020). [1129]
- ZIMMERT, MICHAEL, AND MICHAEL LECHNER (2019): “Nonparametric Estimation of Causal Heterogeneity Under High-Dimensional Confounding,” arXiv preprint. [arXiv:1908.08779](https://arxiv.org/abs/1908.08779). [1130]

Co-editor Guido W. Imbens handled this manuscript as an invited Fisher–Schultz lecture. The invitation to deliver the Fisher–Schultz lecture is also an invitation to publish a suitable version of the lecture in Econometrica.

Manuscript received 27 December, 2020; final version accepted 19 July, 2024; available online 7 April, 2025.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.10871241>. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices.