



Information Systems Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

A Comparison of Methods for Treatment Assignment with an Application to Playlist Generation

Carlos Fernández-Loría, Foster Provost, Jesse Anderton, Benjamin Carterette, Praveen Chandar

To cite this article:

Carlos Fernández-Loría, Foster Provost, Jesse Anderton, Benjamin Carterette, Praveen Chandar (2023) A Comparison of Methods for Treatment Assignment with an Application to Playlist Generation. *Information Systems Research* 34(2):786-803. <https://doi.org/10.1287/isre.2022.1149>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Information Systems Research*. Copyright © 2022 The Author(s). <https://doi.org/10.1287/isre.2022.1149>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2022 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>


A Comparison of Methods for Treatment Assignment with an Application to Playlist Generation

Carlos Fernández-Loría,^{a,*} Foster Provost,^b Jesse Anderton,^c Benjamin Carterette,^c Praveen Chandar^c

^aHong Kong University of Science and Technology, Clear Water Bay, Hong Kong; ^bNew York University, New York, New York 10012;

^cSpotify, New York, New York 10007

*Corresponding author

Contact: imcarlos@ust.hk,  <https://orcid.org/0000-0003-4509-3768> (CF-L); fprovost@stern.nyu.edu (FP); janderton@spotify.com (JA); benjamin@spotify.com (BC); praveen@spotify.com (PC)

Received: July 31, 2020

Revised: April 21, 2021; January 28, 2022

Accepted: April 9, 2022


Published Online in Articles in Advance:
August 2, 2022

<https://doi.org/10.1287/isre.2022.1149>

Copyright: © 2022 The Author(s)

Abstract. This study presents a systematic comparison of methods for individual treatment assignment, a general problem that arises in many applications and that has received significant attention from economists, computer scientists, and social scientists. We group the various methods proposed in the literature into three general classes of algorithms (or metalearners): learning models to predict outcomes (the O-learner), learning models to predict causal effects (the E-learner), and learning models to predict optimal treatment assignments (the A-learner). We compare the metalearners in terms of (1) their level of generality and (2) the objective function they use to learn models from data; we then discuss the implications that these characteristics have for modeling and decision making. Notably, we demonstrate analytically and empirically that optimizing for the prediction of outcomes or causal effects is not the same as optimizing for treatment assignments, suggesting that, in general, the A-learner should lead to better treatment assignments than the other metalearners. We demonstrate the practical implications of our findings in the context of choosing, for each user, the best algorithm for playlist generation in order to optimize engagement. This is the first comparison of the three different metalearners on a real-world application at scale (based on more than half a billion individual treatment assignments). In addition to supporting our analytical findings, the results show how large A/B tests can provide substantial value for learning treatment-assignment policies, rather than simply for choosing the variant that performs best on average.

History: Olivia Liu Sheng, Senior Editor; Gautam Pant, Associate Editor.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Information Systems Research. Copyright © 2022 The Author(s). <https://doi.org/10.1287/isre.2022.1149>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/isre.2022.1149>.

Keywords: [treatment assignment](#) • [treatment effects](#) • [predictive modeling](#)

1. Introduction

Systems that make automated decisions are often deployed with the underlying goal of improving (rather than just predicting) outcomes. For example, when targeting online ads or retention incentives, the goal is often to encourage customers to make a purchase or stay with the company, rather than just predicting their behavior. This type of task is more generally known as a treatment-assignment problem (Manski 2004), where each possible course of action corresponds to a different “treatment” (e.g., “show ad” versus “do not show ad”), and ideally each individual is assigned to the treatment associated with the most beneficial outcome (e.g., the one with the highest profit).

Treatment-assignment policies may be estimated from data using statistical modeling, allowing decision makers to map individuals to the best treatment

according to their characteristics (e.g., preferences, behaviors, or history with the firm). However, there are different ways in which one could proceed, and various methods for the estimation of treatment-assignment policies from sample data have been proposed across different fields, including econometrics (Manski 2004), data mining (Lo 2002), and multiarmed bandits (Beygelzimer and Langford 2009).

This paper gathers these various methods into three general classes of algorithms (or metalearners) for individualized treatment assignment, all of which have been proposed in the literature. Table 1 describes these metalearners in terms of their estimand (i.e., the model that each metalearner intends to estimate) and their resulting treatment-assignment policy.¹ The first metalearner (*Outcome Learner (O-learner)*) learns a model that predicts the expected outcome (Y) given

Table 1. Metalearners for Estimating Treatment-Assignment Policies

Metalearners	Estimand in the learning procedure	Estimated treatment-assignment policy
Outcome Learner (O-learner)	$\mu(x, j) = \mathbb{E}[Y X = x, T = j]$	$\arg \max_j \hat{\mu}(x, j)$
Effect Learner (E-learner)	$\tau(x, j) = \mu(x, j) - \mu(x, 0)$	$\arg \max_j \hat{\tau}(x, j)$
Assignment Learner (A-learner)	$a^*(x) = \arg \max_j \tau(x, j)$	$\hat{a}^*(x)$

Note. Y is the outcome of interest, X are individual-level characteristics, and T is the treatment assignment.

individual-level characteristics ($X = x$) and a specific treatment assignment ($T = j$). Each individual is then assigned to the treatment with the best (e.g., largest) predicted outcome. The second metalearner (*Effect Learner* (*E-learner*)) learns a model for heterogeneous causal effect estimation, and each individual is then assigned to the treatment with the largest predicted causal effect. Finally, the third metalearner (*Assignment Learner* (*A-learner*)) directly learns the treatments that are expected to have the best outcome for each individual.

At a first glance, the policies estimated by the metalearners in Table 1 may look the same: They all seek to assign the treatment with the best outcome. As a main contribution, we reveal two key distinctions between these learning approaches and the implications that these distinctions have for personalized treatment assignment.

The first distinction between metalearners is the level of generality of the tasks that the machine-learned models can perform. For instance, models that predict outcomes ($\hat{\mu}$ in Table 1) may be used to estimate causal effects ($\hat{\tau}$ in Table 1), whereas models that predict causal effects generally cannot predict outcomes. Thus, a more general metalearner may be preferable when the machine-learned models are important for other reasons besides identifying the treatment assignment with the largest outcome, such as when there are decision-making constraints that depend on predicted outcomes or causal effects.

The second distinction is in the objective function that each metalearner optimizes when learning models. The first metalearner optimizes for outcome prediction, the second for causal effect prediction, and the third for optimal action (treatment assignment) prediction. It is clear from Table 1 that when models are highly accurate (i.e., $\hat{\mu} = \mu$, $\hat{\tau} = \tau$, and $\hat{a}^* = a^*$), then all the metalearners produce the same policy. In practice, however, machine-learned models are derived from sample data, so each metalearner may lead to a different treatment-assignment policy, depending on how its objective function frames the prediction problem. Importantly, and as we discuss in detail in this paper, optimizing models to predict outcomes or causal effects is not the same as optimizing models to predict treatment assignments.

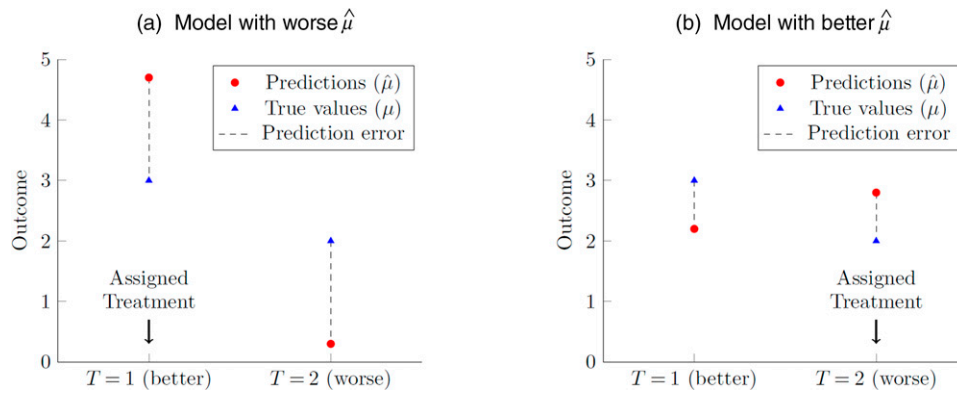
We illustrate this with an example in Figure 1, which compares the outcome predictions made by two different

models for a single individual. One model has high prediction errors (Figure 1(a)), and the other has low prediction errors (Figure 1(b)). The triangles correspond to the true conditional expectations (they are the same for both figures), whereas the dots correspond to the predictions. A larger distance between the triangles and the dots (represented by dashed lines) implies that the model makes worse outcome predictions.

In this example, the conditional expectation when $T = 1$ is larger than when $T = 2$ (as shown by the triangles), which implies that $T = 1$ is a better treatment assignment. Therefore, models make the optimal assignment when $\hat{\mu}(x, 1) > \hat{\mu}(x, 2)$. Figure 1(a) shows that the model with larger prediction errors actually makes the optimal treatment assignment because the rank ordering of the predicted outcomes across treatments is the same as the rank ordering of the true values. On the other hand, the second model depicted in Figure 1(b) makes a worse assignment, even though its prediction errors are smaller, because the ordering is inverted. Importantly, this can also occur when fitting models for causal effect prediction. Therefore, because better outcome or causal effect prediction can lead to worse treatment assignments, learning models that predict optimal assignments (i.e., using an A-learner) should, in principle, lead to better assignments than the other two metalearners.

As a second main contribution, we empirically assess the practical impact that this analytical finding can have on treatment-assignment performance by conducting a massive-scale, experimental comparison of the three metalearners in the context of content selection at Spotify. Our focal application is choosing, for each listener, which playlist-generation algorithm (treatment) to apply in order to maximize the number of songs streamed. To our knowledge, this is the first real-world, at-scale comparison of the three metalearners.

The experiment shows several things. First, it supports the analytical finding that models specifically trained to predict the best treatment are best for treatment assignment (i.e., the A-learner outperforms the O-learner and the E-learner for treatment assignment). This is the case even with training data consisting of more than half a billion observations, a surprising finding given that, in theory, all metalearners should converge to the same treatment-assignment policy

Figure 1. (Color online) Comparison of Outcome Prediction vs. Treatment Assignment for a Single Individual

Notes. The model depicted in (a) has larger outcome-prediction errors than the model depicted in (b), because the dashed lines in (a) are larger than the dashed lines in (b). However, the model in (a) makes a better treatment assignment than the model in (b), because the dots preserve the ranking of the triangles.

with large enough data. This finding also implies that we should reconsider some of the justifications made in prior research for using causal effect prediction methods for personalized treatment assignment.

The experiment also reveals interesting findings for the application to content selection in music streaming. Specifically, it shows that (1) using an algorithm-assignment policy can substantially improve total streaming compared with the typical approach of applying the same playlist-generation algorithm to everyone, and that (2) larger data sets lead to significantly better policies, illustrating the advantages of conducting massive-scale A/B tests for the purpose of learning treatment-assignment policies (rather than just for evaluation or for selecting the best variant).

2. Treatment-Assignment Problem

Treatment-assignment problems correspond to settings where a decision maker wants to maximize the overall causal effect of decisions on an outcome of interest (e.g., deciding what playlist-generation algorithm to use for each listener to maximize the number of streams). Each possible alternative corresponds to a different treatment, and the goal is to assign individuals to the treatment that maximizes their outcome. We formalize the problem in this section.

We consider settings in which decisions are independent and the treatment-assignment policy is learned from historical data on previous decisions made at random. This implies that each decision affects a single unit (or instance), and there is no selection bias in the data. In the causal inference literature, the first assumption is known as the *Stable Unit Treatment Value Assumption* (SUTVA) (Cox 1958). The second assumption is known as *unconfoundedness* and implies that there are no relevant systematic differences between users assigned to different treatments. Unconfoundedness also

goes by other names in different fields, including ignorability (Rosenbaum and Rubin 1983), the back-door criterion (Pearl 2009), and exogeneity (Wooldridge 2015). Practically speaking, these assumptions hold when we use a carefully designed randomized A/B test to gather data.

Let T be the treatment-assignment variable and Y be the observed outcome. We use potential outcomes to frame causality (Rubin 1974) and define $Y(j)$ as the outcome we would observe if we were to assign treatment j (out of k possible treatment alternatives), so that $Y = Y(j)$ if $T = j$. Then, the treatment assignment that would lead to the best outcome, on average, is:

$$a^* = \arg \max_j \mathbb{E}[Y(j)], \quad (1)$$

and can be estimated by using the sample mean ($\hat{\mathbb{E}}$) for each treatment:

$$\hat{a} = \arg \max_j \hat{\mathbb{E}}[Y|T = j]. \quad (2)$$

Equation (2) describes a standard A/B-test approach that compares multiple treatments across a predefined population. This approach, however, does not address individualized treatment assignments. In our setting, we *learn* optimal treatments for individuals or subpopulations. Suppose individuals vary with respect to a set of variables (features) X . We can then think of a feature vector x as a subpopulation where $X = x$ and formulate the optimal assignment (given x) as:

$$a^*(x) = \arg \max_j \mathbb{E}[Y(j)|X = x]. \quad (3)$$

Without the argmax, the right-hand side is essentially the formulation of a predictive model. Applying statistical modeling frees us from specifying in advance what are the particular subpopulations of interest. In Section 3, we present three metalearners that can be

used to estimate $a^*(x)$ from data, each producing a treatment-assignment policy $\hat{a}(x)$.

Treatment-assignment policies can be evaluated in terms of their ability to minimize the expected difference between the outcome when optimal assignments are made, $Y(a^*(X))$, and the outcome when the policy is deployed, $Y(\hat{a}(X))$. This evaluation measure is also known as “expected regret” (or just “regret”) in decision theory:

$$\text{Regret}(\hat{a}) = \mathbb{E}_{Y(a^*(X)), Y(\hat{a}(X))} [Y(a^*(X)) - Y(\hat{a}(X))], \quad (4)$$

and minimizing regret is the same as maximizing the expected outcome of deploying the policy:

$$\mathbb{E}_{Y(\hat{a}(X))} [Y(\hat{a}(X))]. \quad (5)$$

However, evaluating treatment-assignment policies using historical data (as is typical when building standard machine-learning models) is challenging because we do not observe all potential outcomes for any given individual; we only observe the single potential outcome for the treatment actually given. Therefore, if (for any given individual) the policy assigns a different treatment from the one that was assigned in the historical data, we do not know the corresponding potential outcome. Fortunately, given a data set of n individuals from a randomized A/B test, we can still obtain an unbiased estimate of Equation (5) (Li et al. 2010):

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}(\hat{a}(x_i) = t_i) \frac{y_i}{\mathbb{P}(T = t_i)}, \quad (6)$$

where, for each individual i , x_i is the feature vector, t_i is the assigned treatment in the data, y_i is the observed outcome, and $\mathbb{P}(T = t_i)$ is the probability of being assigned to treatment t_i in the data (a known quantity if the data were collected through a randomized A/B test). We present a simplified proof of this result in the online appendices (see Li et al. 2010 for a detailed proof).

3. Metalearners

The causal inference literature often focuses its attention on the estimation of aggregate causal effects, such as the so-called *average treatment (or causal) effect (ATE)*, which corresponds to the average effect of a treatment across the individuals in some well-defined population. However, estimating the ATE does not help us to target different individuals with different treatments, because it does not discriminate between the individuals in the population at all. A fundamental motivation behind this work is that the population exhibits *heterogeneous treatment effects (HTEs)*, which are defined in terms of the degree to which a treatment may have different effects on different individuals (Imai et al. 2013).

One can account for HTEs through the estimation of *conditional average treatment effects (CATEs)*, which

correspond to the average causal effect conditioned on a set of available features. Thus, to the extent that individuals in the population differ on their features (and those features are related to causal effects), we may estimate different causal effects for each individual. Of course, treatment effects may still vary among individuals that share the same features (because we may not be accounting for all aspects related to the causal effect). Nevertheless, the estimation of HTEs by using CATEs allows us to make different interventions for different individuals without knowing the relevant subpopulations in advance.

As we discuss in more detail in Section 6.1, the ideas behind CATE estimation have been fundamental to the development of methods for learning treatment-assignment policies from data. We group such methods into three classes of algorithms (or metalearners), described below and in Table 1. This grouping is meant to highlight the perspective that outcome prediction, causal effect prediction, and treatment assignment are different tasks, which has important implications for modeling and the use of predictive models for decision making.

1. **O-Learner:** Standard machine learning is used to learn a model that is optimized to predict the outcome under each treatment. The policy assigns individuals to the treatment with the largest predicted outcome.

2. **E-Learner:** A machine-learning method specifically designed to estimate CATEs is used to learn a model optimized to predict the *differences* between the outcome of each treatment and a baseline (or control). The policy assigns individuals to the treatment with the largest predicted difference (i.e., treatment effect).

3. **A-Learner:** The problem of estimating optimal treatment assignments can be transformed into a weighted classification problem, wherein each treatment corresponds to a class, and the optimal classifier corresponds to the optimal treatment-assignment policy (Zadrozny 2003, Zhang et al. 2012). Weights are often defined by using the observed outcome under each treatment condition, so that weights are larger for treatments with larger outcomes. Thus, standard machine learning can be used to learn a weighted classification model optimized to predict the treatment with the largest weight. The resulting classifier can then be used to assign individuals to the class (treatment) that is predicted to have the largest weight (outcome).

Section 6.1 discusses multiple studies that have used or recommended specific instances of these metalearners for treatment assignment. Importantly, the three metalearners converge to optimal treatment assignments with large enough samples, assuming that the machine-learning procedure that is used to learn the models is a consistent estimator of the estimands presented in Table 1. However, there are two

Table 2. Comparison of Metalearners for Estimating Treatment-Assignment Policies

Metalearners	Model may be used to predict:			Learning procedure optimized for
	Outcomes	Effects	Assignments	
Outcome learner (O-learner)	✓	✓	✓	Outcomes (MSE_{μ} in Equation (9))
Effect learner (E-learner)	×	✓	✓	Causal effects (MSE_{τ} in Equation (12))
Assignment learner (A-learner)	×	×	✓	Assignments (WMR in Equation (13))

key differences between the metalearners that are critical in practice, as summarized in Table 2 and discussed next.

3.1. Distinction 1: Level of Generality

The first key distinction between the three metalearners is their level of generality (the metalearners are listed above from the most general to the least general). O-learners are the most general of the metalearners because they produce models that predict outcomes, and such models may also be used to predict causal effects or optimal treatments. Specifically, causal effect predictions may be obtained by taking the difference between the predicted outcomes of two treatments under consideration, and optimal-treatment predictions may be obtained by selecting the treatment with the largest predicted outcome. Therefore, O-learners may be used for all three different purposes.

Models that predict causal effects cannot be used to predict outcomes. For such models, predictions estimate the expected marginal change in the outcome that results from assigning some specific treatment, but the predictions cannot be used to estimate expected outcomes under an arbitrary treatment condition. Therefore, although causal effect predictions may still be used to predict optimal treatments (by selecting the treatment with the largest predicted effect), E-learners are not as general as O-learners.

Finally, models trained to predict optimal treatment assignments (i.e., learned with an A-learner) can only be used for that purpose. These models, the least general, cannot predict the outcome or the effect that would result from making those assignments.

This distinction implies that more general metalearners (O-learners and E-learners) can be preferable over A-learners when outcome and causal effect predictions are important for other reasons besides their usefulness to determine the treatment with the most beneficial outcome. For example, McFowland et al. (2021) consider treatment-assignment settings where there are budget constraints and the decision maker faces costs that are unknown ex ante. In such settings, quantifying the benefit of each individual decision (e.g., via causal effect prediction) and the cost of each possible course of action (e.g., via outcome prediction) is important to allocate resources in the most profitable way.

Feasible treatment-assignment rules can also be constrained for ethical, legislative, or political reasons. For example, a public policy maker may want to prioritize the assignment of subsidies to individuals in some protected class, unless the predicted effect of the subsidy on annual income is below a certain threshold or the individual is predicted to already have an annual income above a certain threshold. Because assessing whether an individual meets these two conditions would require causal effect and outcome predictions, implementing an A-learner in this type of setting may be counterproductive or infeasible.

A-learner decision rules may also be more difficult to implement in settings where the models are intended to support (rather than automate) human decision making. For example, for economic policy and medical treatment assignment, decision makers may need to weigh the potential benefit of the treatment alternatives with respect to some other information not available to the model (e.g., how the individuals affected by the treatments feel about the treatment alternatives), so predicting the treatment with the “most beneficial outcome” may not suffice. Additionally, information about causal effects and outcomes can be important for other reasons beyond decision making (e.g., for users to trust the model, to debug the model, or to develop more effective treatments in the future).

Nonetheless, in settings where outcomes and effects are relevant only for the model to assign individuals to the most beneficial treatment, we should expect A-learners to make better treatment assignments because they are specifically designed for treatment assignment; we elaborate on this premise in detail in the rest of the paper.

3.2. Distinction 2: Learning Procedure

The second key distinction is that each metalearner uses a different learning objective (or loss function) for the machine learning. O-learners use a loss function designed to optimize outcome predictions; E-learners use a loss function designed to optimize causal effect predictions; and A-learners use a loss function designed to optimize treatment assignments. This implies that, although all metalearners share the same ultimate goal (optimizing treatment assignments, as specified by Equations (4), (5), and (6)), they

differ with respect to the procedures they use to learn from data.

This distinction is important because an improvement in the prediction of outcomes or causal effects does not imply an improvement in treatment assignment (as previously shown in Figure 1). In fact, the improvements may occur at the expense of worse treatment assignments! Thus, we should expect machine learning with loss functions specifically tailored to optimize treatment assignments to produce better models: The A-learner should outperform the other metalearners with finite training data when making treatment-assignment decisions.

Nevertheless, O-learners and E-learners are much more commonly used among scholars and practitioners in marketing and information systems (IS), even though those metalearners are optimized to minimize prediction errors in outcomes or causal effects, rather than decision-making errors. One goal of this study is to encourage a more widespread consideration, study, and use of A-learners among management researchers by showing how decisions can be substantially improved when machine-learning models are directly optimized for treatment assignment (decision making).

4. Choice of Objective Function

In this section, we compare the three metalearners analytically to illustrate how their choice of objective function may affect their performance in treatment assignment.

4.1. Outcome Prediction

As mentioned, the O-learner assigns treatments by learning one or more models that predict the expected outcome of each treatment ($\hat{\mu}$):

$$\hat{\mu}(x, j) = \hat{\mathbb{E}}[Y|X = x, T = j], \quad (7)$$

and then selecting the treatment with the best predicted outcome:²

$$\hat{a}_\mu(x) = \arg \max_j \hat{\mu}(x, j). \quad (8)$$

A standard approach to fit Equation (7) is to regress outcome Y on features X and treatment-assignment T using various machine-learning methods designed to minimize the mean squared error for the outcome (MSE_μ):

$$MSE_\mu(\hat{\mu}, j) = \mathbb{E}_{X,Y}[(Y - \hat{\mu}(X, j))^2 | T = j], \quad (9)$$

and then to choose the model(s) with the lowest MSE_μ .

The premise here is that minimizing MSE_μ implies better outcome predictions and, therefore, better treatment assignments. However, optimizing for outcome prediction (by minimizing MSE_μ or other measures,

such as mean absolute error or cross-entropy) does not necessarily optimize for treatment assignment. Going back to the earlier example in the introduction, Figure 1(a) shows that the model with larger prediction errors makes the optimal treatment assignment because the rank ordering of the predicted outcomes is the same as the rank ordering of the true values. The second model makes a worse assignment, even though its prediction errors are smaller, because the ordering is inverted. Therefore, choosing the model with the lower (and, thus, better) MSE_μ leads to a worse treatment assignment.

Multiple researchers have noted the potential of overfitting when multiple outcome models are used to estimate treatment effects instead of directly fitting a causal effect model (Nie and Wager 2017, Künzel et al. 2019). For example, suppose that features X_1 and X_2 are predictive of outcomes, but only feature X_1 is predictive of effects. This implies that, for the purposes of estimating effects and assigning treatments, segmenting individuals using exclusively X_1 is more statistically efficient than segmenting them according to X_1 and X_2 . Hence, by focusing statistical power on features that are predictive of effects, models optimized for treatment-effect estimation can achieve lower bias and lower variance than models optimized for outcome prediction. Subsequently, many researchers have proposed methods that directly model effects (rather than outcomes) to make treatment assignments; these are instances of the E-learner. However, as we discuss next, optimizing for causal effects is not the same as optimizing for treatment assignments either.

4.2. Causal Effect Prediction

The second metalearner, the E-learner, consists of learning one or more models to estimate the CATEs ($\hat{\tau}$) directly:

$$\hat{\tau}(x, j) = \hat{\mathbb{E}}[Y|X = x, T = j] - \hat{\mathbb{E}}[Y|X = x, T = 0], \quad (10)$$

where $T = 0$ corresponds to a baseline treatment (e.g., the control in an A/B-test setting). The optimal treatment may then be chosen as follows:

$$\hat{a}_\tau(x) = \arg \max_j \hat{\tau}(x, j). \quad (11)$$

The (sometimes unstated) goal of machine-learning methods designed for the estimation of CATEs is to minimize the mean squared error for treatment effects (MSE_τ):

$$MSE_\tau(\hat{\tau}, j) = \mathbb{E}_{X,Y(j),Y(0)}[(Y(j) - Y(0) - \hat{\tau}(X, j))^2]. \quad (12)$$

Therefore, these methods are not optimized to predict outcomes, but to predict causal effects, which are usually defined as the difference between potential outcomes (i.e., $Y(1) - Y(0)$). The main challenge is that we

only observe one potential outcome for any given individual, so we cannot calculate Equation (12) directly. However, we may use alternative formulations to estimate MSE_{τ} from data (Schuler et al. 2018), allowing us to compare (and optimize) models on the basis of how good they are at predicting causal effects.

Unfortunately for our application, and similarly to the previous section, optimizing causal effect predictions (by minimizing MSE_{τ}) is not the same as optimizing treatment assignments either. We illustrate this in Figure 2, which shows a similar example to the one illustrated in Figure 1, except that it compares the causal effect (rather than outcome) predictions made by two models.³ Therefore, the triangles in this example represent the causal effects of the treatments for a specific individual (they are the same in both graphs), and the dots represent the estimation of the effects by the models. As before, the first model has high prediction errors (Figure 2(a)), but makes a better assignment, whereas the second has lower prediction errors (Figure 2(b)), but makes a worse assignment. Thus, the model that makes a better causal effect prediction (i.e., that has lower MSE_{τ}) makes a worse treatment assignment.

Surprisingly, this implies that models that are (relatively) bad at causal effect prediction may be good at making treatment assignments. This result, although seemingly counterintuitive at first, may be attributed to the bias-variance decomposition of errors. In the machine-learning community, it is well known that models that have a good classification performance are not necessarily good at estimating class probabilities, and vice versa (Friedman 1997). A useful analogy in our context is to think about treatment assignment as a classification problem and to think about causal effect estimation as a probability-estimation problem; the two tasks are closely related, but not exactly the same. Importantly, the bias and variance components

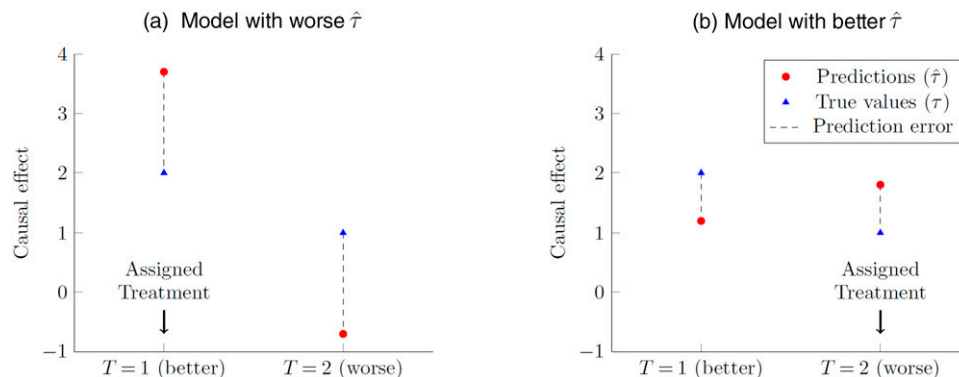
of the estimation error in causal effect predictions may combine to influence treatment-assignment errors in a very different way than with the squared error of the predictions themselves (Fernández-Loría and Provost 2022a).

Figure 3 illustrates this in more detail by depicting the sampling distribution of the causal effect estimates previously shown in Figure 2. Specifically, Figure 3(a) shows that the large prediction errors of the model in Figure 2(a) are the result of high bias because the sampling distributions are not centered on the causal effect estimands. However, this model works very well for treatment assignment because $\mathbb{P}(\hat{\tau}(x,2) < \hat{\tau}(x,1)) \approx 1$ and $\tau(x,2) < \tau(x,1)$. On the other hand, Figure 3(b) shows that the model in Figure 2(b) is an unbiased estimator of causal effects and has lower mean squared error. However, this model is more likely to make the incorrect assignment due to sampling errors (variance). Importantly, what matters in this case is not the accuracy of the causal effect estimates, but how good they are at discriminating between treatment alternatives. We elaborate more on this in Section 4.4, after discussing the objective function of the A-learner.

4.3. Treatment-Assignment Prediction

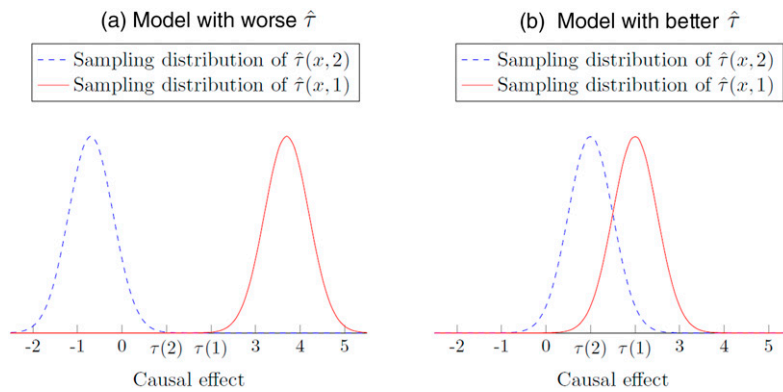
The third metalearner, the A-learner, estimates the treatment-assignment policy by directly learning the treatment assignments that lead to the best outcomes. As Zhang et al. (2012) describe in detail, the treatment-assignment problem can be transformed into a weighted classification problem. The idea is that each treatment alternative can be mapped to a class, and classes are associated with weights that correspond to the cost of not predicting the corresponding class. Weights are generally defined in terms of the potential outcomes associated with each treatment alternative (Beygelzimer and Langford 2009, Zhao et al. 2012, Kitagawa and Tetenov 2018),

Figure 2. (Color online) Comparison of Causal Effect Prediction vs. Treatment Assignment for a Single Individual



Notes. The model depicted in (a) has larger effect-prediction errors than the model depicted in (b), because the dashed lines in (a) are larger than the dashed lines in (b). However, the model in (a) makes a better treatment assignment than the model in (b), because the dots preserve the ranking of the triangles.

Figure 3. (Color online) Sampling Distributions of $\hat{\tau}$



Notes. The model depicted in (a) is a biased estimator of causal effects, whereas the model depicted in (b) is unbiased. However, the model in (b) is more likely to estimate that $\hat{\tau}(x, 1) < \hat{\tau}(x, 2)$ (and, thus, make the wrong assignment) because of sampling error.

but more broad definitions exist (Zhang et al. 2012). Thus, the goal is to “classify” individuals into the class (treatment) with the largest weight in order to minimize misclassification costs.

The challenge here is that we only observe a single weight for any given individual (the one associated with the treatment assigned in the data), so we do not observe correct classifications at the individual level. Fortunately, samples from treatment-assignment problems can be transformed into weighted classification samples, so that any importance-weighted classification algorithm can be used to learn treatment-assignment policies. For example, given a probability distribution $\mathbb{P}(T)$ over the treatment (e.g., the probability that an individual gets assigned to treatment T in the A/B-test data), each observation $(x, y, \text{ and } t)$ can be transformed into an importance-weighted example, where $y/\mathbb{P}(t)$ is the cost of not predicting treatment t given input x (Beygelzimer and Langford 2009, Zhao et al. 2012); if t is predicted, then the cost is zero. The weighted misclassification rate (*WMR*) of classifier \hat{a} under this setting is:

$$WMR(\hat{a}) = \mathbb{E}_{X,Y,T} \left[\mathbf{1}(\hat{a}(X) \neq T) \frac{Y}{\mathbb{P}(T)} \right], \quad (13)$$

which is directly tied to treatment-assignment performance because minimizing the *WMR* is equivalent to minimizing expected regret (as defined in Equation (4)). We present a simplified proof of this result in the appendices (see Beygelzimer and Langford 2009 for more details). As a result, we should expect *WMR* to be a better objective function than MSE_{μ} or MSE_{τ} when the goal is to make the best possible treatment assignments.

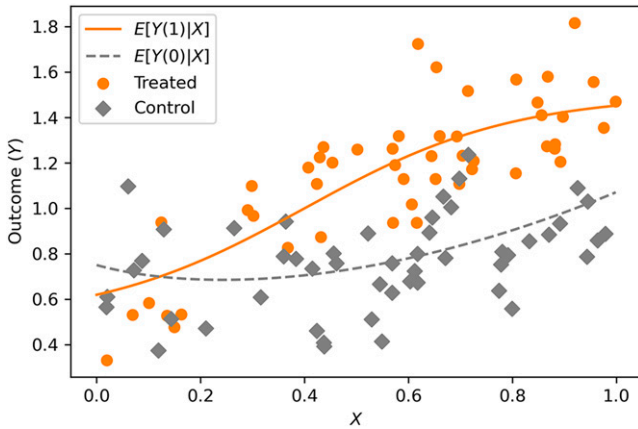
4.4. Bias-Variance Tradeoff

In this section, we use a simulated example to illustrate how the A-learner can exploit the bias-variance

tradeoff in the learning procedure to make better treatment assignments than the E-learner. The data from the simulated example are shown in Figure 4. There are two treatment alternatives, treat and not treat, and the goal is to use feature X to learn a treatment-assignment policy to discriminate individuals into those that would benefit from the treatment (the solid line is above the dashed line) and those that would not (the dashed line is above the solid line). The lines represent the expected outcomes under both treatment conditions, and the dots and diamonds represent the available training data. See the online appendices for a description of the data-generating process.

Similarly to previous studies that have considered models with constrained functional forms for ethical, legislative, or political reasons (Kitagawa and Tetenov 2018, Athey and Wager 2021), suppose that we are considering the class of treatment rules that split the population only once (e.g., decision trees with a single split). One alternative is to split individuals according to causal effect heterogeneity and then make treatment assignments according to their estimated causal effects, which corresponds to the E-learner.


Figure 5 shows the result of learning a single-split causal tree—a tree that splits individuals according to causal effects (Athey and Imbens 2016). The estimand (solid line) is the expected treatment effect given X (the estimand in causal effect estimation); the dotted line is the prediction from the tree that is learned from the data shown in Figure 4; and the best-in-class (dash-dotted line) is the prediction from a tree that is learned with unlimited data (i.e., the best possible single split that could be made to minimize MSE_{τ}). Finally, the dashed line corresponds to the actual decision boundary: It is optimal to treat when the expected treatment effect (solid line) is above the boundary. The treatment-assignment policy treats when

Figure 4. (Color online) Sample Data for Treated and Control Individuals

Notes. The lines represent the expected potential outcomes for treated and untreated, and the dots and diamonds represent the data points in the sample. Individuals should be treated when the solid line is above the dashed line and should not be treated otherwise.

the predicted effect (dash-dotted or dotted lines) is above the boundary.

Note that a treatment-assignment policy based on estimated causal effects (dotted line) would not treat individuals with $X \leq 0.6$, even though most individuals with such values for X would actually benefit from the treatment. This is the result of variance in the estimation procedure. With more data, errors due to variance eventually disappear, and the estimated causal effects converge to the best-in-class predictor.

Figure 5. (Color online) Learning a Single-Split Tree by Optimizing for Effect Prediction

Notes. The solid line is the expected causal effect; the dotted line is the estimated causal effect, according to a single-split tree learned from data; and the dash-dotted line is the best possible effect estimation with a single-split tree. Neither the estimated tree nor the best-in-class tree lead to optimal assignments.

Unfortunately, the best-in-class predictor does not lead to optimal assignments either, because it estimates that the treatment is beneficial for everyone, even though individuals with a small value for X do not benefit from the treatment. In this case, errors in treatment assignment occur due to bias in the estimation procedure: The causal tree is not complex enough to identify who does not benefit from the treatment. Nevertheless, this does not imply that the class of treatment rules that split the population only once is not complex enough to model treatment assignments. It just implies that learning a tree by splitting the population according to effect heterogeneity (an E-learner) does not lead to optimal assignments.

A second alternative is to use an A-learner to split the population according to *preferred treatment assignments*; in other words, to learn a classification model optimized to minimize Equation (13) (WMR) instead of Equation (12) (MSE_τ). An important challenge, however, is that Equation (13) can be viewed as a weighted version of 0-1 loss, and it is well known in the machine-learning literature that minimizing such loss is difficult due to its discontinuity and nonconvexity. A common approach to address this challenge is to use a surrogate loss to learn a scoring model, such as the negative log likelihood in logistic regression or the hinge loss in support vector machines (Zhao et al. 2012), and then classify individuals according to their scores. The predictions of the resulting scoring model ($\hat{\theta}$) would correspond to:

$$\hat{\theta}(x, j) = \hat{\mathbb{P}}(\tilde{T} = j | X = x), \quad (14)$$


where \tilde{T} corresponds to the weighted treatment class, and the scoring model may be used to choose the optimal treatment as follows:

$$\hat{a}_\theta(x) = \arg \max_j \hat{\theta}(x, j). \quad (15)$$

Nevertheless, in the case of tree-based algorithms, it is tractable to optimize according to 0-1 loss (and, hence, according to WMR). Figure 6 shows the result of learning such a tree. Similarly to before, the solid line is the target score in the weighted classification task (the estimand); the dotted line is the scoring model that is learned from the data shown in Figure 4; the dash-dotted line is the scoring model that is learned with unlimited data (i.e., the best possible single split that could be made to minimize WMR); and the dashed line corresponds to the decision boundary.

Note that, although the estimated scoring model (dotted line) also suffers from errors due to variance and bias, these errors do not affect decision making (treatment assignment) as much. For example, there is a substantial underestimation of the treatment score

Figure 6. (Color online) Learning a Single-Split Tree by Optimizing for Treatment Assignment



Notes. The solid line is the target treatment score; the dotted line is the estimated score, according to a single-split tree learned from data; and the dash-dotted line is the best possible score estimation with a single-split tree. The estimated tree does a good job assigning treatments, and the best-in-class tree leads to optimal assignments, even though it does not provide entirely accurate estimations of scores.

for individuals with a small value for X , but this does not affect decision making because the optimal decision for those individuals is not to intervene. Importantly, with more data, the model eventually converges to the best-in-class predictor, which leads to optimal assignments, even though the treatment scores are biased due to the simplicity of the treatment rules under consideration.

Essentially, this example shows that minimizing errors in causal effect (or outcome) predictions may not imply better treatment assignments because the direction of the errors is critical. For the purposes of decision making, overestimations do not hurt when the treatment is beneficial, and underestimations do not hurt when the treatment is detrimental. This point is important because E-learners may correct such errors at the expense of increasing errors that will hurt decision making (as shown in Figure 3). In contrast, A-learners are specifically designed to minimize errors that hurt decision making; errors that do not affect decisions are essentially ignored.

Table 3. Comparison of Splitting Criteria

Splitting criterion	Squared bias ^a	MSE_{τ}^b	Regret ^b
MSE_{τ} (E-learner split)	0.007	0.007	0.008
WMR (A-learner split)	0.019	0.019	0

^aSquared bias = $\mathbb{E}_X[(\tau(X) - \mathbb{E}[\hat{\tau}(X)])^2]$

^bThese measures exclude idiosyncratic noise and assume unlimited data, so they are exclusively driven by bias.

In our example, the E-learner splits the data to minimize the bias in the causal effect predictions, whereas the A-learner splits the data to minimize the bias that negatively affects decisions. Table 3 compares the two metalearners with unlimited training data (resulting in the best-in-class in Figures 5 and 6). The A-learner split leads to larger bias (and, hence, larger MSE_{τ}) than the E-learner split, but that bias does not have negative implications for decision making (regret).

Of course, if decisions are based on *true* causal effects (i.e., the estimand in Figure 5), then regret is also minimized. So, with more data, one could learn a more complex causal effect model (e.g., a tree with more splits), decrease the modeling bias, and eventually converge to optimal decision making. Nonetheless, as our subsequent empirical analysis shows, the A-learner can outperform the O-learner and the E-learner even when the training sample consists of hundreds of millions of observations.

5. Experiments and Results

We now present an empirical comparison of the three metalearners for choosing which playlist-generation algorithm to apply for each listener (see Liebman et al. 2019 for an overview of prior playlist-generation studies in IS).

5.1. Application Setting

In our playlist-generation setting, the treatment variants consist of different algorithmic playlist-generation systems tested in production by Spotify, a media services provider. Each system uses a different algorithm to select and rank songs in “algorithmic” playlists (playlists that are built dynamically according to user data). The company has multiple goals when deploying such systems (e.g., converting users from free to premium, reducing churn, and increasing engagement with the platform). However, the complexities of data collection, modeling, and deployment have historically made it prohibitively difficult for systems to be directly optimized in terms of these goals. Therefore, the models that underlie these systems are often heuristic (e.g., songs are ranked based on their similarity to other songs the user has played).

We focus specifically on the number of streamed songs in the playlist as a proxy for engagement and use it as our target outcome metric; this measure is significantly less noisy than other alternative engagement metrics and is available for all users.⁴ Thus, the goal is to assign users to the playlist-generation system with which they would listen to the most songs. Firms typically run A/B tests to compare new machine-learned systems with the existing production system (as a baseline) and decide whether to replace the production system with one of the new systems. This essentially

chooses the same treatment assignment for all users. However, as we have argued throughout this paper, different variants may work better for different users: If system A is best for new users and system B is best for more experienced users, then deploying the same system for all users would lead to suboptimal treatment assignments. Because the outcome of interest in this case is song streams, one could then learn a treatment-assignment policy that deploys different systems for different users in order to maximize the number of song streams.

In the online appendices, we provide a detailed discussion on how the deployment of content-selection systems differs from other treatment-assignment problems.

5.2. Data

We compare the three metalearners using data from a massive, production A/B test. The A/B test produced a data set in which four different playlist-generation systems were randomly assigned to users to build algorithmic playlists: three newly developed playlist-generation systems and the system that was currently in production. More specifically, each observation corresponded to a user who selected an algorithmic playlist, and each playlist was built by using one of the four systems (chosen at random) to select and rank songs.

There are 770 million observations in the data: 86.68% assigned to the production system and 4.44% for each of the new variants. For each observation, we have the following categorical features: country (19 values), playlist ID (this serves as an identifier of the pool of songs that can be used to build the algorithmic playlist; six values), platform (e.g., Android; three values), user tenure in days (transformed into a discrete variable with four values), and product (e.g., free or premium; eight values). For each categorical variable, the categories with fewer than 10,000,000 observations were grouped together in a category named “Other,” resulting in the number of values for each variable reported above. Descriptive statistics for the features are shown in Table 4, and balance tests with respect to these features confirmed an adequate randomization

Table 4. Descriptive Statistics of the Categorical Features Available for Treatment Assignment

Categorical feature	No. of values	Entropy ^a	Mode
Platform	3	0.471	iOS
Country	19	0.596	United States
Playlist ID	6	0.169	Other
Product	8	0.559	Free subscription
User tenure	4	0.309	>179 days

^aEntropy was normalized to range between zero and one. Larger values imply a more uniform distribution of observations among feature values.

of the systems. For each observation, we also have the number of total streams for the user (which is the target outcome).

Importantly, these massive data allow us to assess what would happen with data sets of many different sizes. Given a universal approximator (e.g., a tree-induction algorithm), all metalearners converge to the same (optimal) treatment-assignment policy when the training data are large enough. So, the interest here is not to compare the metalearners when there are “unlimited” (i.e., very large) training data. Instead, the interest is to compare them across data sets of different sizes.

5.3. Learning and Evaluating Policies

We compare treatment-assignment policies estimated with each metalearner: (1) the *O-learner policy*, which assigns treatments based on a model that predicts the total number of streams for each system (Equation (8)); (2) the *E-learner policy*, which assigns treatments based on a model that predicts the increase in the total number of streams (compared with control) for each system (Equation (11)); and (3) the *A-learner policy*, which uses a classification model to predict and assign the system that is estimated to produce the largest number of streams (Equation (15)).

We use tree-based algorithms to learn all models, so that differences in performance can be attributed to the loss functions used by each metalearner, rather than the machine-learning algorithm being used. We chose trees over other alternatives for multiple reasons. First, this choice allows us to demonstrate that the simulated example in Section 4.4 is not merely hypothetical: A-learner trees can, indeed, lead to substantially better assignments than E-learner trees. Second, tree models can be adapted to predict outcomes, effects, and optimal assignments, regardless of the metalearner; this will be important to compare the metalearners in Section 5.4.3. Finally, trees were the best-performing models out of the multiple alternatives we considered. In the online appendices, we provide an extended analysis that also considers random forests and linear models.

The O-learner uses a decision-tree regressor that minimizes MSE_{μ} (Equation (9)) to learn $\hat{\mu}$ (Equation (7)). The E-learner uses a decision-tree regressor on the transformed variable proposed by Athey and Imbens (2016) to learn $\hat{\tau}$ (Equation (10)) by minimizing MSE_{τ} (Equation (12))—that is, a “causal tree.” The E-learner policy uses three causal trees, one for each system, except control (the regressor does not support nonbinary treatments). Finally, the A-learner uses a weighted decision-tree classifier that minimizes a proxy⁵ of WMR (Equation (13)) to learn $\hat{\theta}$ (Equation (14)).

The models were learned, tuned, and evaluated by using 10-fold nested cross-validation, which separates the

cross-validation used for hyperparameter optimization from the test folds used for evaluation (Provost and Fawcett 2013). All hyperparameters were tuned to optimize their respective loss functions: The O-learner was tuned to optimize MSE_{μ_i} ; the E-learner was tuned to optimize MSE_{τ} ; and the A-learner was tuned to optimize WMR. We used the empirical measure described in Equation (6) to evaluate all policies, using $\mathbb{P}(t_i) = 86.68\%$ when $t_i = 0$ (control), and $\mathbb{P}(t_i) = 4.44\%$ otherwise. However, for clarity, the analysis that follows compares this quantity relative to assigning the incumbent (control) system to everyone, which is the percentage increase in streams.

5.4. Results

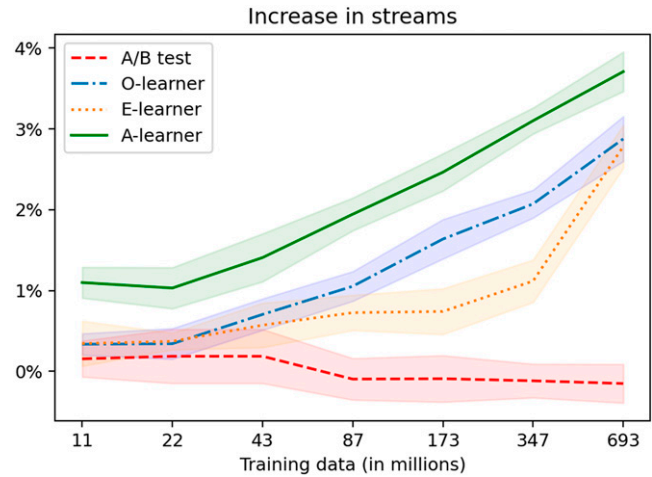
As mentioned, all metalearners eventually converge to the same treatment-assignment policy if the training data are large enough. However, most firms don't have access to unlimited experimental data, or even an experimental data set as large as the one presented in this study. Thus, this analysis assesses how the metalearners compare with various data sizes.

Figure 7 shows the performance of each metalearner (measured as the increase in streams relative to the baseline) as the size of the training data increases. The dashed line is the policy where the system that performs best, on average, is applied to everyone—this is what we would get from a standard A/B test. The other lines correspond to the treatment-assignment policies estimated by the O-learner, the E-learner, and the A-learner. The areas around the lines represent 95% confidence intervals calculated using the 10 results from the cross-validations.

5.4.1. The Importance of Individualized Treatment-Assignment Policies. The first interesting finding in Figure 7 is that choosing the system that performs best, on average, does not significantly increase the total number of song streams. This implies that no single “best system” for all users performs much better than the baseline (existing production system). Importantly, if we were to follow a traditional A/B-test approach, we might erroneously decide to use the system currently in production for everyone, because no other system produces an increase in streams that is statistically significant at the population level.

We show in Table 5 the percentage of users that would be assigned to each system when the entire data are used to estimate treatment-assignment policies. Recall that our analysis was conducted using nested cross-validation, so the table was built using out-of-sample treatment assignments for all users. The first row in the table shows that different systems may be selected as the best system (on average), depending on the cross-validation folds used to compare the systems: System $T = 1$ is selected for six out of 10 folds, whereas system $T = 2$ is selected for the other four

Figure 7. (Color online) Treatment-Assignment Performance by Data Size



Notes. All treatment-assignment metalearners improve substantially with more data and work better than assigning everyone to the system that works best according to the A/B test (the dashed line). Also, optimizing for assignment prediction (A-learner) works better than optimizing for outcome prediction (O-learner) or causal effect prediction (E-learner).

folds. Combined with Figure 7, this suggests that there is no real performance difference between the systems when applied to the entire population.

However, total streams can be increased substantially when different systems are applied to different users. Table 5 shows that the policies that were estimated by using machine learning (O-learner, E-learner, and A-learner) exhibit high heterogeneity in their treatment assignments, and these policies also perform substantially better (as shown in Figure 7) than assigning the single best playlist-generation system to everyone. To control for the fact that high heterogeneity may be the result of using different folds to estimate the models, we computed the entropy in treatment assignments for each test fold and then obtained the average entropy across all 10 folds (i.e., column (4) in Table 5). As we can see, all metalearners exhibit a large entropy (i.e., high heterogeneity) in treatment assignments within the folds, resulting in substantially more streams.

5.4.2. The Importance of Large A/B Tests. Another important result is that treatment-assignment policies become increasingly better with more training data, illustrating the importance of conducting large A/B tests to generate unconfounded training data to learn models for personalized treatment assignments. As we mentioned in Section 3, the estimation of causal effects by using CATEs instead of the ATE is a substantial improvement for the purposes of deciding on individual interventions. However, the estimation

Table 5. Percentage of Users Assigned to Each Treatment

Policy	$T = 0$ (%)	$T = 1$ (%)	$T = 2$ (%)	$T = 3$ (%)	Average entropy ^a
Best on average ^b	0.0	60.0	40.0	0.0	0
O-learner	11.1	25.3	33.5	30.1	1.896
E-learner	9.3	32.1	29.0	29.5	1.879
A-learner	9.3	32.0	29.4	29.3	1.879

^aAverage entropy of treatment assignments across folds. The minimum is zero, and the maximum is two.

^bDifferent systems perform best (on average) depending on the folds that are used to select the best system. Thus, not everyone is assigned to a single system when using cross-validation to evaluate and analyze the “best on average” policy.

of accurate CATEs requires much larger data sets; otherwise, the models are likely to overfit. Conducting large A/B tests alleviates this problem because the data can be partitioned into fine-grained subpopulations of users without losing as much statistical power. Correspondingly, we can fit more complex models with less overfitting.

5.4.3. The Importance of the Learning Objective.

Figure 7 also shows that learning policies by optimizing treatment assignments (A-learner) works better than learning policies by optimizing outcome and causal effect predictions (O-learner and E-learner, respectively), thus validating our analytical findings. Importantly, this is the case *even* when the training data contain more than half a billion observations. As discussed in detail, objective functions that optimize things other than treatment-assignment prediction (i.e., better outcome or causal effect predictions) do not necessarily favor better treatment assignments.

Going back to our analytical examples, we would expect each metalearner to perform best doing whatever it is optimized to do. For example, the predictive model estimated by the O-learner should perform better at predicting outcomes than the models estimated by the E-learner and the A-learner. However, as discussed in Section 3.1, we cannot (in general) use causal effect models or treatment-assignment models to estimate outcomes. Therefore, in order to compare the performance of the metalearners at different tasks, we adapt the models estimated by the E-learner and the A-learner in the following analysis.

Because the metalearners in our analysis estimate tree-based models, we can generalize the E-learner and the A-learner models by using different prediction functions to aggregate the training observations at each leaf, depending on the task at hand. For instance, if we want to make outcome predictions, the prediction function would consist of the average outcome of the observations in the leaf (rather than the average causal effect, in the case of the E-learner model, or the treatment with the largest average outcome, in the case of the A-learner model). Thus, the structure of each tree model remains the same, but the prediction

function at the leaf level can be changed to predict outcomes, causal effects, or best treatment assignments.

Table 6 shows the performance of each metalearner at the three different tasks, evaluated by using nested cross-validation on the entire data set. The tasks are predicting outcomes (a lower MSE_{μ} is better), predicting causal effects (a lower MSE_{τ} is better), and predicting treatment assignments (where the goal is to increase song streams). As expected, each metalearner is best at doing what it was optimized to do (the best result in each column is in bold). Thus, the models with the best performance in outcome prediction (O-learner) and causal effect prediction (E-learner) are not the best models at making treatment assignments. In fact, the improvement in streams produced by the A-learner is more than 28% larger than the improvement produced by either the O-learner or the E-learner.

6. Discussion

This paper groups treatment-assignment methods into three general metalearners: the outcome learner, the causal effect learner, and the treatment-assignment learner. The grouping allows us to compare treatment-assignment methods in terms of (1) their level of generality and (2) the objective function they use to learn models from data; both of these characteristics have important implications for modeling and decision making, as discussed in detail in Section 3.

One of the major implications is that optimizing for outcome or causal effect predictions (O-learner and E-learner, respectively) is not the same as optimizing for treatment assignments (A-learner), so the latter ought to perform better in practical (nonasymptotic)

Table 6. Policy Comparison at Different Tasks

Metalearners	MSE_{μ}	MSE_{τ} ^a	Increase in streams (%) ^b
O-learner	0.057	46.287	2.88
E-learner	0.111	46.276	2.78
A-learner	0.059	46.305	3.71

Note. Bold text correspond to the best result for each column.

^a MSE_{τ} corresponds to the MSE of the transformed outcome proposed by Athey and Imbens (2016) to estimate causal effects.

^bRelative to the system in production.

settings due to its ability to exploit the bias-variance tradeoff that results from framing the treatment-assignment problem as a classification task, instead of a numeric prediction task. We support this claim analytically and also empirically for the real-world application of choosing, for each listener, which playlist-generation algorithm to apply in order to maximize the number of song streams. Although, in theory, all metalearners should converge to the same (optimal) treatment-assignment policy with unlimited data, we find that the A-learner’s advantage over the other metalearners can persist, even when treatment-assignment policies are estimated with more than half a billion observations.

Our study also illustrates how large A/B tests can provide substantial value for learning treatment-assignment policies (rather than simply choosing the variant that performs best, on average). In our application, none of the individual treatments (the different playlist-generation algorithms) increases streams when applied over the whole population, but the best of the treatment-assignment policies we estimated can increase the number of song streams by 3.7%.

6.1. Methods for Treatment Assignment

Our metalearner categorization encompasses many methods proposed in the literature for learning treatment-assignment policies from data. We discuss below how the methods proposed in multiple fields of study fit into our categorization (see Table 7 for a summary).

6.1.1. Econometrics. The field of econometrics has long recognized that treatment assignment is a distinct problem from the point-estimation and hypothesis-testing problems usually considered in the treatment-effects literature (Manski 2004, Dehejia 2005, Hirano and Porter 2009, Bhattacharya and Dupas 2012). However, most of the methods proposed in the econometrics literature correspond to instances of the O-learner and the E-learner. As we discuss in Section 6.2, some recent studies recommend the A-learner, showing that it is statistically efficient under a wide range of theoretical

conditions (Kitagawa and Tetenov 2018, Athey and Wager 2021).

6.1.2. Uplift Modeling. Uplift modeling consists of assigning treatments based on the estimation of the incremental (causal) impact of a treatment on individuals’ behaviors (Lo 2002, Kane et al. 2014). It has been recommended by the data-mining community for targeting applications such as online advertising and customer retention (Radcliffe and Surry 2011). The uplift-modeling literature typically focuses on settings where treatment assignments and outcomes are binary, so methods are usually grouped into two main categories (Rzepakowski and Jaroszewicz 2012): the two-model approach and the single-model approach (which are specific instances of the O-learner and the E-learner, respectively). Most studies favor the use of the E-learner over the O-learner, but benchmark studies show that the O-learner can perform better depending on the data set (Jaskowski and Jaroszewicz 2012, Olaya et al. 2020). To our knowledge, none of these studies have considered the A-learner.

6.1.3. Causal Effect Estimation. As mentioned, most of the causal inference literature focuses on the estimation of causal effects, rather than treatment-assignment policies. However, the main motivation behind the use of machine-learning methods for CATE estimation is often treatment assignment (Athey and Imbens 2019; 2017), which corresponds to the E-learner. Popular machine-learning methods for CATE estimation include Bayesian additive regression trees (Hill 2011), causal random forests (Wager and Athey 2018), and regularized causal support vector machines (Imai et al. 2013). There is also a relatively large number of papers showing asymptotic properties of the E-learner for treatment assignment when an efficient or consistent estimator of CATEs is known (see Kitagawa and Tetenov 2018 and Athey and Wager 2021 for an overview), but most of them do not discuss the results of deploying such systems in practice. Dorie et al. (2019) provide an overview and a benchmark of several CATE estimation methods. Recent studies in the marketing and

Table 7. Summary of Literature for Treatment Assignment

Field of study	Discussion about approaches
Econometrics	Recommendations include the O-learner and the E-learner. Recent studies recommend the A-learner and show that it should be preferred under certain theoretical conditions.
Uplift modeling	Comparisons and recommendations include the O-learner and the E-learner. The E-learner is often the preferred approach, but the O-learner has been shown to be competitive in benchmark studies.
Causal effect estimation	The E-learner is often recommended for treatment assignment, including in IS and marketing.
Multiarmed bandits	The O-learner is the most commonly used approach, but the A-learner was first proposed and recommended in this field.
Treatment assignment as classification	Some studies have noted the close similarity between treatment assignment and classification. They recommend the A-learner.

IS literature recommend such methods for treatment assignment (see Section 6.3 for examples).

6.1.2. Multiarmed Bandits. Treatment-assignment policies are also at the core of contextual multiarmed bandits. Models for multiarmed-bandit problems may be used to learn how to make decisions in situations where the payoff of only one choice is observed (Beygelzimer and Langford 2009, Dudík et al. 2011). Such methods have been used to make automated decisions about online news recommendations to maximize clicks (Li et al. 2010), for example. It is precisely in this stream of research that it was first noted that the treatment-assignment problem (as defined in Section 2) is mathematically equivalent to a weighted classification problem (Zadrozny 2003, Beygelzimer and Langford 2009), leading to the suggestion of the A-learner. Nonetheless, the O-learner has also been recommended for multiarmed-bandit algorithms—LinUCB being a well-known example (Li et al. 2010)—and remains the most common approach in multiarmed-bandit problems.

An important distinction between our setting and the multiarmed-bandit problem is that the goal in bandit problems is to learn a treatment-assignment policy while actively making treatment-assignment decisions for incoming subjects. Therefore, there is an exploration/exploitation dilemma that plays an important role in the decision-making procedure, whereas, in our case, the decision maker cannot re-estimate the treatment-assignment policy after making each decision.⁶ Our setting is also referred to as “offline learning” in this literature (Beygelzimer and Langford 2009).

6.1.3. Treatment Assignment from a Classification Perspective. Zhang et al. (2012) proposed a generalized classification framework to show how several estimators of optimal treatment regimes can be represented as special cases of weighted classification within their framework; these estimators are essentially instances of the A-learner. This framework defines weights in terms of a contrast function that may represent outcomes (as in Zhang et al. 2012 and Kitagawa and Tetenov 2018), causal effects (as in Athey and Wager 2021), or some other business-oriented importance weight (as in Lemmens and Gupta 2020). In the personalized medicine literature, the A-learner is also referred as “outcome-weighted learning” (Zhao et al. 2012, Chakraborty and Moodie 2013).

6.2. Prior Comparisons of the Metalearners

To our knowledge, no prior study has compared the three metalearners as defined in our study, either analytically or on a real-world application at scale, but there are some partial exceptions.

Schuler et al. (2018) propose a framework to compare and select models based on their ability to predict

outcomes, causal effects, and optimal treatments, but they do not consider how learning models based on these three criteria may affect treatment-assignment performance. Interestingly, they show through simulations that selecting models based on their ability to predict causal effects generally leads to better treatment assignments than selecting models based on their ability to predict outcomes or optimal treatments (hence, suggesting that the E-learner is a good candidate for treatment assignment). However, their comparisons do not include the A-learner.

Beygelzimer and Langford (2009) provide a theoretical regret analysis for multiarmed-bandit problems showing that, for a given family of regressors (e.g., decision trees), the A-learner has a smaller lower-bound regret than the O-learner. These analytical results are supported by experiments on multiclass benchmark data sets that were repurposed to simulate potential outcomes, showing the A-learner as a superior alternative than the O-learner.

Other recent studies have also made theoretical developments that are in line with our study and argue in favor of the A-learner. Kitagawa and Tetenov (2018) investigate the statistical performance of the A-learner (to which they refer as Empirical Welfare Maximization) in terms of its uniform convergence rate of regret. They show that in settings where propensity scores are known (e.g., when data are acquired through A/B tests), the A-learner attains minimax optimal rates in finite samples over various classes of feasible data distributions. Athey and Wager (2021) extend these results to show the asymptotical efficiency of the A-learner when estimating policies from observational data, so they address cases where (1) propensity scores are not known or (2) there is endogeneity (so unconfoundedness is not met), but effects can be estimated using instrumental variables.

Our study builds on this past work in several important ways. First, we elaborate on the advantages and disadvantages of the A-learner with respect to other metalearners commonly proposed in the treatment-assignment literature. Second, we show that the A-learner can outperform other metalearners as a result of optimizing the bias-variance tradeoff with respect to decision-making errors (i.e., treatment assignment), rather than conventional prediction errors in outcomes or causal effects. Third, we empirically demonstrate the advantages of the A-learner for the deployment of content-selection systems in the context of music streaming.

Other studies in the uplift-modeling literature have also compared the O-learner and the E-learner (Jaskowski and Jaroszewicz 2012, Olaya et al. 2020). Notably, Olaya et al. (2020) conducted an extensive benchmark study comparing the performance of 13 uplift-modeling methods that are instances of the O-learner and the E-learner, as defined by our categorization. The comparisons included data sets across many domains of

interest, including marketing, political behavior, and clinical trials. They found that none of the evaluated techniques consistently outperform the other techniques. Although the uplift-modeling literature generally favors the E-learner over the O-learner, the consensus is that choosing between the O-learner and the E-learner should be an empirical undertaking because the O-learner sometimes beats the E-learner (Jaskowski and Jaroszewicz 2012, Olaya et al. 2020). As discussed next, this may also occur when choosing between the A-learner and the other metalearners.

6.3. Generalizability of the Results

Our results do not imply that the A-learner will necessarily outperform the other metalearners in other settings. Results may change, depending on the data-generating process, the available features, the size of the data, and the machine-learning algorithm used by the metalearners (see the online appendices for an extended empirical analysis that considers some of these factors). However, our argument is that, in general, the A-learner should be a strong (if not the strongest) contender.

This point is important because several recent papers in the marketing and the IS literature employ machine learning to estimate causal effects that are then used for treatment assignment (Ascarza 2018, Miller and Hosanagar 2020, Yang et al. 2020, McFowland et al. 2021); these methods would be considered O-learners or E-learners under our categorization. As an exception, Lemmens and Gupta (2020) propose a profit-based loss function that uses a weighting scheme that could be cast as an instance of an A-learner; they show that their approach performs better than the O-learner when managing churn to maximize profits. However, among practitioners and scholars who consider the causal impact of treatments when making treatment assignments,⁷ O-learners and E-learners are, by far, the most common approaches for treatment assignment in the marketing and IS literature and in many other fields of study.

6.4. Limitations and Future Research

One important assumption in our study is that the training data are generated from A/B tests. For cases where this assumption is not met, the A-learner can be used to leverage observational data where causal effects can be identified by using a variety of strategies, including selection on observable features and instrumental variables (Athey and Wager 2021). In a similar vein, Fernández-Loría and Provost (2019) assessed the impact of confounding bias on treatment assignments when causal effects cannot be identified from data, but it is unclear from their study how the three metalearners would compare with each other in settings where confounding is prevalent. This presents a promising direction for future research.

The second major assumption in this work is that all decisions are independent. However, this is unlikely to be the case when multiple treatments are assigned to the same individuals over time. Treatment-assignment policies for such settings are also known as dynamic treatment regimes (Murphy 2003, Chakraborty and Murphy 2014). We are unaware of any A-learner methods specifically designed to model dependent decisions, so this a natural next step to extend this line of work.

Another related setting that our study does not consider is that of constrained optimization, such as when treatment assignments are made under budget constraints. Under such circumstances, it is no longer necessarily the case that assigning individuals to the treatment with the most beneficial outcome is the optimal thing to do. McFowland et al. (2021) propose a framework to maximize the overall utility of treatment assignments under budget constraints by using machine learning for predicting costs and causal effects and then optimizing an integer linear program that converts predictions into treatment assignments. This is also known in the operations literature as a predict-then-optimize framework (Elmachtoub and Grigas 2022). That is, machine learning is used first to predict unknown input parameters of an optimization problem, and then decisions are made by solving the optimization problem using the predicted parameters.

An implication from our study is that the assignments should be better when the objective function in the machine learning measures errors in the assignments induced by the predicted input parameters, as opposed to errors in the prediction of the input parameters themselves. Such methods have been applied in the operations literature for a variety of optimization problems with linear objectives (Elmachtoub et al. 2020, Elmachtoub and Grigas 2022). Future research could also explore and apply these ideas in the context of treatment-assignment problems, as suggested in Fernández-Loría and Provost (2022b).

Finally, we hope that this study will encourage other researchers to further explore other parallels between policy estimation and “traditional” predictive modeling, besides the classification analogy discussed in this paper. For example, policy estimation (as defined in this paper) is somewhat related to learning to rank (LTR)—as discussed in the information-retrieval literature (e.g., Burges et al. 2005). More specifically, for each user, we could rank a set of actions according to their potential outcomes and choose the action at the top of the ranking. Of course, LTR is not the same type of task because, in our setting, we only care about the difference in outcomes between the action we choose and the action at the top of the ranking (not the entire ranking). Nonetheless, LTR methods could potentially be adapted to address treatment-assignment

problems, perhaps in a similar fashion to the weighted-classification transformation discussed in this paper.

Endnotes

¹ When the goal is to identify the treatments that lead to the lowest outcomes (e.g., when Y is the number of hospitalizations), the argmax should be replaced with an argmin in this and all subsequent equations.

² Note that this policy does not use confidence intervals to make treatment-assignment decisions; only the point estimates are used. This is because, from the perspective of a regret minimizer, the best possible choice (in expectation) is the treatment with the largest point estimate, regardless of the confidence intervals. For this reason, confidence intervals are not incorporated as part of the decision making in our study.

³ In this case, the treatments $T=1$ and $T=2$ are being compared with a baseline treatment $T=0$. Therefore, the y -axis in Figure 2 represents the difference in outcomes with respect to $T=0$.

⁴ Alternative engagement metrics, such as whether the user added the playlist to favorites, consist of actions that may not be available to all users and are generally harder to optimize due to their rare occurrence.

⁵ For the empirical analysis, we estimated the weighted decision-tree classifier by splitting based on the (weighted) Gini impurity because the machine-learning library we used (scikit-learn) did not offer the option of learning the tree by directly optimizing for the weighted misclassification rate.

⁶ Few firms have the ability to deploy full-blown online machine-learning systems that can manage the exploration/exploitation tradeoff dynamically. It is much more common to deploy the learned models/prediction systems than the machine-learning systems themselves.

⁷ As discussed in Ascarza (2018), it is not uncommon for practitioners to estimate treatment-assignment policies from data without any sort of causal modeling. Fernández-Loría and Provost (2022a) give a modeling justification for this.

Acknowledgments

The authors thank the participants at the Workshop on Information Systems and Economics (WISE) 2020 for their valuable feedback, in particular Ravi Bapna, who was the discussant of our paper at the workshop. Foster Provost thanks Ira Rennert for his support. They also thank the NYU/Stern Fubon Center for supporting research on AI and data analytics for business.

References

Ascarza E (2018) Retention futility: Targeting high-risk customers might be ineffective. *J. Marketing Res.* 55(1):80–98.

Athey S, Imbens G (2016) Recursive partitioning for heterogeneous causal effects. *Proc. Natl. Acad. Sci. USA* 113(27):7353–7360.

Athey S, Imbens GW (2017) The state of applied econometrics: Causality and policy evaluation. *J. Econom. Perspect.* 31(2):3–32.

Athey S, Imbens GW (2019) Machine learning methods that economists should know about. *Annu. Rev. Econ.* 11:685–725.

Athey S, Wager S (2021) Policy learning with observational data. *Econometrica* 89(1):133–161.

Beygelzimer A, Langford J (2009) The offset tree for learning with partial labels. *Proc. 15th ACM SIGKDD Internat. Conf. Knowledge Discovery Data Mining* (ACM, New York), 129–138.

Bhattacharya D, Dupas P (2012) Inferring welfare maximizing treatment assignment under budget constraints. *J. Econometrics* 167(1):168–196.

Burges C, Shaked T, Renshaw E, Lazier A, Deeds M, Hamilton N, Hullender G (2005) Learning to rank using gradient descent. *Proc. 22nd Internat. Conf. Machine Learning* (ACM, New York), 89–96.

Chakraborty B, Moodie EE (2013) *Statistical Methods for Dynamic Treatment Regimes, Statistics for Biology and Health*, vol. 2 (Springer, New York).

Chakraborty B, Murphy SA (2014) Dynamic treatment regimes. *Annu. Rev. Statist. Appl.* 1:447–464.

Cox DR (1958) *Planning of Experiments* (Wiley, New Jersey).

Dehejia RH (2005) Program evaluation as a decision problem. *J. Econometrics* 125(1-2):141–173.

Dorie V, Hill J, Shalit U, Scott M, Cervone D (2019) Automated vs. do-it-yourself methods for causal inference: Lessons learned from a data analysis competition. *Statist. Sci.* 34(1):43–68.

Dudík M, Langford J, Li L (2011) Doubly robust policy evaluation and learning. Getoor L, Scheffer T, eds. *Proc. 28th Internat. Conf. Machine Learning* (Omnipress, Madison, WI), 1097–1104.

Elmachtoub A, Jason CNL, McNellis R (2020) Decision trees for decision-making under the predict-then-optimize framework. Daumé III H, Singh A, eds. *Internat. Conf. Machine Learning* (PMLR), 2858–2867.

Elmachtoub AN, Grigas P (2022) Smart “predict, then optimize.” *Management Sci.* 68(1):9–26.

Fernández-Loría C, Provost F (2019) Observational vs experimental data when making automated decisions using machine learning. Preprint, submitted August 29, <https://dx.doi.org/10.2139/ssrn.3444678>.

Fernández-Loría C, Provost F (2022a) Causal classification: Treatment effect estimation vs. outcome prediction. *J. Machine Learning Res.* 23(59):1–35.

Fernández-Loría C, Provost F (2022b) Causal decision making and causal effect estimation are not the same. . . and why it matters. *INFORMS J. Data Sci.*, ePub ahead of print March 10, <https://doi.org/10.1287/ijds.2021.0006>.

Friedman JH (1997) On bias, variance, 0/1-loss, and the curse-of-dimensionality. *Data Mining Knowledge Discovery* 1(1):55–77.

Hill JL (2011) Bayesian nonparametric modeling for causal inference. *J. Comput. Graphical Statist.* 20(1):217–240.

Hirano K, Porter JR (2009) Asymptotics for statistical treatment rules. *Econometrica*. 77(5):1683–1701.

Imai K, Ratkovic M (2013) Estimating treatment effect heterogeneity in randomized program evaluation. *Ann. Appl. Statist.* 7(1):443–470.

Jaskowski M, Jaroszewicz S (2012) Uplift modeling for clinical trial data. Elhadad N, Hauskrecht M, eds. *ICML Workshop Clinical Data Anal.* (International Machine Learning Society, San Diego).

Kane K, Lo VSY, Zheng J (2014) Mining for the truly responsive customers and prospects using true-lift modeling: Comparison of new and existing methods. *J. Marketing Anal.* 2(4):218–238.

Kitagawa T, Tetenov A (2018) Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86(2):591–616.

Künzel SR, Sekhon JS, Bickel PJ, Yu B (2019) Metalearners for estimating heterogeneous treatment effects using machine learning. *Proc. Natl. Acad. Sci. USA* 116(10):4156–4165.

Lemmens A, Gupta S (2020) Managing churn to maximize profits. *Marketing Sci.* 39(5):956–973.

Li L, Chu W, Langford J, Schapire RE (2010) A contextual-bandit approach to personalized news article recommendation.

- Proc. 19th Internat. Conf. World Wide Web* (ACM, New York), 661–670.
- Liebman E, Saar-Tsechansky M, Stone P (2019) The right music at the right time: Adaptive personalized playlists based on sequence modeling. *MIS Quart.* 43(3):765–786.
- Lo VSY (2002) The true lift model: A novel data mining approach to response modeling in database marketing. *ACM SIGKDD Explorations Newslett.* 4(2):78–86.
- Manski CF (2004) Statistical treatment rules for heterogeneous populations. *Econometrica* 72(4):1221–1246.
- McFowland E, Gangarapu S, Bapna R, Sun T (2021) A prescriptive analytics framework for optimal policy deployment using heterogeneous treatment effects. *MIS Quart.* 45(4):1807–1832.
- Miller A, Hosanagar K (2020) Personalized discount targeting with causal machine learning. Karahanna E, Sarker S, Oestreicher-Singer G, eds. *Internat. Conf. Inform. Systems ICIS* (Association for Information Systems, Atlanta).
- Murphy SA (2003) Optimal dynamic treatment regimes. *J. Roy. Statist. Soc. Ser. B Statist. Methodol.* 65(2):331–355.
- Nie X, Wager S (2017) Quasi-oracle estimation of heterogeneous treatment effects. Preprint, submitted December 13, <https://arxiv.org/abs/1712.04912>.
- Olaya D, Coussemont K, Verbeke W (2020) A survey and benchmarking study of multitreatment uplift modeling. *Data Mining Knowledge Discovery* 34(2):273–308.
- Pearl J (2009) *Causality: Models, Reasoning and Inference* (Cambridge University Press, Cambridge, UK).
- Provost F, Fawcett T (2013) *Data Science for Business: What You Need to Know About Data Mining and Data-Analytic Thinking* (O’Reilly Media, Inc., Newton, MA).
- Radcliffe NJ, Surry PD (2011) Real-world uplift modelling with significance-based uplift trees. White Paper TR-2011-1, Stochastic Solutions, Edinburgh, UK.
- Rosenbaum PR, Rubin DB (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1):41–55.
- Rubin DB (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Ed. Psych.* 66(5):688–701.
- Rzepakowski P, Jaroszewicz S (2012) Decision trees for uplift modeling with single and multiple treatments. *Knowledge Inform. Systems* 32(2):303–327.
- Schuler A, Baiocchi M, Tibshirani R, Shah N (2018) A comparison of methods for model selection when estimating individual treatment effects. Preprint, submitted April 14, <https://arxiv.org/abs/1804.05146>.
- Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* 113(523):1228–1242.
- Wooldridge JM (2015) *Introductory Econometrics: A Modern Approach* (Nelson Education, Toronto).
- Yang J, Eckles D, Dhillon P, Aral S (2020) Targeting for long-term outcomes. Preprint, submitted October 29, <https://arxiv.org/abs/2010.15835>.
- Zadrozny B (2003) Policy mining: Learning decision policies from fixed sets of data. Unpublished PhD thesis, University of California, San Diego, La Jolla, CA.
- Zhang B, Tsiatis AA, Davidian M, Zhang M, Laber E (2012) Estimating optimal treatment regimes from a classification perspective. *Stat* 1(1):103–114.
- Zhao Y, Zeng D, Rush AJ, Kosorok MR (2012) Estimating individualized treatment rules using outcome weighted learning. *J. Amer. Statist. Assoc.* 107(499):1106–1118.