

IRTG 1792 Discussion Paper 2019-026



Affordable Uplift: Supervised Randomization in Controlled Experiments

Johannes Haupt^{*}
Daniel Jacob^{*}
Robin M. Gubela^{*}
Stefan Lessmann^{*}



^{*} Humboldt-Universität zu Berlin, Germany

This research was supported by the Deutsche
Forschungsgesellschaft through the
International Research Training Group 1792
"High Dimensional Nonstationary Time Series".

<http://irtg1792.hu-berlin.de>
ISSN 2568-5619

International Research Training Group 1792

AFFORDABLE UPLIFT: SUPERVISED RANDOMIZATION IN CONTROLLED EXPERIMENTS

MANUSCRIPT

Johannes Haupt, Daniel Jacob, Robin M. Gubela, Stefan Lessmann
Humboldt-Universität zu Berlin
{johannes.haupt, daniel.jacob, robin.gubela, stefan.lessmann}@hu-berlin.de

ABSTRACT

Customer scoring models are the core of scalable direct marketing. Uplift models provide an estimate of the incremental benefit from a treatment that is used for operational decision-making. Training and monitoring of uplift models require experimental data. However, the collection of data under randomized treatment assignment is costly, since random targeting deviates from an established targeting policy. To increase the cost-efficiency of experimentation and facilitate frequent data collection and model training, we introduce *supervised randomization*. It is a novel approach that integrates existing scoring models into randomized trials to target relevant customers, while ensuring consistent estimates of treatment effects through correction for active sample selection. An empirical Monte Carlo study shows that data collection under supervised randomization is cost-efficient, while downstream uplift models perform competitively.

Keywords: *Uplift Modeling, Causal Inference, Experimental Design, Selection Bias*

1 Introduction

Direct marketing plays a key role in consumer markets. The continuous growth of e-commerce, accounting for 1.8 trillion Euros globally in 2019 [1], is accompanied by simultaneous growth in online and email advertising. Spending on traditional print advertising like catalog marketing has shown a similar growth [2]. At the core of scalable direct marketing, campaign analysts employ models to predict future customer behavior and target responsive clients [3].

For example, a decision tree could be trained to predict the probability for a customer to purchase in the next week based on known characteristics. The expected behavior of the customer could then be used to inform operational decision-making in that customers with a probability below average are targeted with an incentive. However, the predictive model is agnostic to the marketing policy, the overall effectiveness of the marketing action and the effect of the marketing action on individual customers. Outcome models provide an estimate of customer behavior, but fail to provide an estimate of the potential change in customer behavior, which is the goal of marketing intervention.

A growing research stream advocates that the decision which customers to target should be addressed directly through causal inference in the form of uplift models [4, 5]. Instead of predicting customer behavior, uplift models estimate the causal effect of a marketing action on an individual customer given their characteristics. In the above example, an uplift tree could be trained to predict the increase in probability for a customer to conduct a purchase in the next week if a catalog was sent. Uplift models thus provide an estimate of the incremental benefit from the marketing treatment, which can explicitly be used as a direct criterion for operational decisions by comparing it to the incremental cost. Conceptually, uplift models align with the actual decision problem of choosing the action with the highest incremental gain for each customer.

Uplift models are trained on experimental data and estimate the treatment effect by comparing the observed behavior of a group of individuals who have received the treatment, the treatment group, and a distinct group of individuals who have not received the treatment, the control group. Similarly, experimental data is required to evaluate the performance of uplift models [6]. In contrast, non-causal models of customer behavior are trained and evaluated on customers of which all or none have received the treatment. Collecting experimental data in randomized experiments is well established in practice in the form of A/B tests. Although used to evaluate the gross benefit of campaigns, A/B tests are not commonly used for uplift modeling to estimate individualized treatment effects [7].

During experiments, random assignment of individuals to either the treatment or control group is crucial to train unbiased uplift models. However, data collection through randomized experiments is costly, since random targeting withholds marketing spending on some customers that would be targeted under the established targeting policy and applies spending on customers that would otherwise not be targeted. The deployment of uplift models exacerbates data collection costs since decision support systems typically require continuous or frequent evaluation and occasional retraining on recent observations, which in turn require fresh experimental data.

We propose a novel approach for the collection of experimental data for uplift modeling based on the combination of cost-optimized randomization at the time of data collection and selection bias correction during model building, which we refer to as *supervised randomization*. In a nutshell, supervised randomization introduces a stochastic component to the existing targeting model and extends the standard experimental design with full randomization by considering customers that are rejected by the targeting policy.

Our contribution is two-fold. First, we show that supervised randomization can be used to integrate existing scoring models into randomized trials. The integration of existing scoring models into group assignment increases the cost-efficiency of experimentation and facilitates continuous data collection during regular business operation. Continuous data collection is critical for non-disruptive experimentation, monitoring the performance of uplift models under deployment and recurring model training. Facilitating model training and monitoring has the additional benefit to improve the acceptance of causal models by management and stakeholders.

Second, we introduce inverse probability weighting and doubly robust estimation as methods to control for biased treatment assignment to the uplift literature. Uplift models have so far relied on the assumption of data collected under full or imbalanced randomization in randomized controlled trials. We show that recent advances in the econometrics literature extend the applicability of uplift models to cases with non-standard treatment assignment. The bias-corrected uplift models are shown to perform competitively on simulated data.

2 Background

Consider a marketing action applied to an individual user i as a treatment intended to change an observed outcome Y_i . Let $D_i \in \{0, 1\}$ be an indicator if the individual has been treated and denote the outcome with and without treatment as Y_i^1 and Y_i^0 , respectively. Then the individual treatment effect is the incremental gain caused by a marketing action $Y_i^1 - Y_i^0$. Because a customer either does or does not receive the marketing action, the actual treatment effect is not observable. We can, however, estimate the treatment effect on the population or on the individual level. We denote the average campaign uplift as average treatment effect (ATE) and the customer-level uplift $\tau = E[Y_i^1 - Y_i^0 | x = X_i]$ as individualized treatment effect (ITE) [8, 9], sometimes alternatively denoted conditional average treatment effect (CATE) in the econometrics literature. Furthermore, we refer to a model used to estimate the outcome Y_i as outcome model and a model used to estimate the treatment effect τ as a causal model, as a more general alternative to the term uplift model. The operational decision problem posed in uplift modeling is to decide if an individual customer should receive the marketing treatment. The decision is automated through the targeting policy, a mapping from the estimated ITE to the binary decision of whether to treat the individual.

Three assumptions are needed for causal inference following the potential outcome theorem [10]. First, the *Stable Unit Treatment Value Assumption (SUTVA)* guarantees that the potential outcome of a customer is unaffected by changes in the treatment assignment of other customers. This assumption may be violated when treatment effects propagate through the social network of customers [11]. In settings of low value or low involvement products, research on treatment effects in marketing typically assumes that no interaction takes place [12] or that the indirect effect of treatment on other customers is at least substantially smaller than the direct effect of the treatment [13].

The second assumption is *conditional unconfoundedness*, i.e. the independence between the potential outcomes and the treatment assignment given the observed covariates (X) [10].

$$(Y_i^1, Y_i^0) \perp\!\!\!\perp D_i | X_i = x. \quad (1)$$

The third assumption, called *overlap*, guarantees that for all $x \in \text{supp}(X_i)$ the probability to receive treatment $e(x) = P(D = 1 | X_i = x)$ is bounded away from 0 and 1:

$$0 < e(x) < 1. \quad (2)$$

When the treatment assignment process is under the control of the experiments as in the customer targeting setting, conditional independence and overlap can be ensured by design through fully randomizing treatment assignment with treatment probability $e(x) = e \in (0, 1)$. Randomized experiments assign individuals at random to one of at least two conditions, where each condition entails a specific treatment. In controlled experiments, one condition is the

control condition in which individuals receive no treatment. In combination, randomized controlled trials (RCT) are the gold standard of data collection for causal inference. We refer to uniform treatment assignment as *full randomization*. Supervised randomization provides a framework that preserves the advantage of the randomized experimental design but allows some control over the probability of treatment assignment on the individual level.

3 Literature Review

The unbiased training of causal models and targeting policies requires data that fulfills the assumptions of the potential outcome theorem. In addition, the unbiased evaluation of causal models and policies also requires experimental data and metrics developed for counterfactual prediction [6, 12]. Violation of the unconfoundedness (Eq. 1) and overlap assumptions (Eq. 2) in observational studies cannot be substituted by collecting more data in the form of more covariates or more observations [14]. Randomized experiments are thus considered a prerequisite to uplift modeling. However, the design and costs of RCT are often not discussed in the literature. We aim to fill this gap by proposing a more efficient design for randomized experiments. We first summarize recent developments in causal machine learning, related research on efficient experimental design and methods to correct for treatment assignment in observational studies.

Causal machine learning methods can be divided into direct and indirect approaches. Direct estimation algorithms construct a feasible loss to estimate a model for the ITE. Indirect approaches model the expected customer response conditional on the treatment group and estimate the ITE as the difference between expected responses. This study employs a robust, indirect two-model logistic regression and a state-of-the-art, direct causal forest for the empirical comparison and provides a discussion of these models below. For an in-depth discussion and benchmark of recent methods for ITE estimation see [8, 9, 15].

Indirect approaches estimate the treatment effect via estimating the response with and without treatment using common statistical learners. The two-model approach [6], or K-model approach in settings with more than one treatment, estimates a separate model for the outcome in the treatment group and control group data and estimates the ITE as the difference between the predicted outcomes. The two-model approach is flexible with regard to the underlying outcome models. K-nearest neighbors learners [16] and deep neural networks [17] have demonstrated promising model performance in the two-model framework. While recent research advocates discretizing the outcome variable to use classification models in continuous settings [4], the two-model approach extends naturally to both categorical and continuous outcomes. This facilitates the use of classification and regression models to forecast, for example, purchase completion or customer spending, respectively.

A number of well-known machine learning algorithms have been extended to estimate the ITE directly without the need to model the customer response [18, 19]. Note that the average treatment effect within a subgroup provides a useful estimate of the treatment effect for individuals within that subgroup. Hence, algorithms that split the data into groups to calculate estimates on the subset are inherently applicable to causal modeling and modifications of the k-nearest neighbor estimator [12] and tree-based models [20, 21] have been applied to estimate individualized treatment effects. Causal tree models modify the Classification and Regression Tree by a splitting criterion maximizing the expected variance in treatment effects between leaves [21]. Within each terminal node conditional on the covariate splitting, the conditional average treatment effect can be estimated and provides an ITE for the observations falling into that node. Causal trees can be combined into ensembles through bagging or boosting. [8] propose a gradient-boosted ensemble of causal trees and an algorithm using multivariate adaptive regression splines. Causal forests are similarly flexible models and have been shown to be consistent and asymptotically normal for a fixed covariate space [22].

Both direct and indirect approaches to ITE estimation share the need for experimental data. The collection of experimental data has not been explicitly explored in the uplift literature. However, concerns over the organizational difficulty and the opportunity cost of running randomized controlled trials have led to research on the optimal use of available data and efficient experimental design in related fields.

A popular strategy for the evaluation of multiple targeting policies is to avoid experimentation for each candidate policy and instead to estimate each policy’s performance using one existing, fully randomized experiment. The cost-efficient evaluation is possible through extrapolation from observations where the policy recommendation matches the observed random treatment assignment, weighted to match the actual population [23, 24, 12]. This evaluation strategy requires existing experimental data, while our goal is to decrease the cost of collecting experimental data through efficient randomization. Approaches to efficient evaluation and efficient randomization are therefore complementary.

The experimental design of previous studies indicates awareness of data collection costs. Table 1 shows the marketing goal, data accessibility, number of observations and the imbalance between treatment and control group sizes of experimental campaigns in customer targeting applications. The large number of observations in recent studies is unsurprising since common technologies in e-commerce settings (e.g., web cookies) facilitate the collection of large

Table 1: Randomized treatment data in marketing

Application	Source	Access	Observations	Treatment / Control Ratio
Direct mail in office supplies	[25]	Closed	460,000	17:1
Mail promotion	[26]	Closed	550,000	10:1
Cross-selling mail in insurance	[27]	Closed	34,370	9:1
MSN subscription	[28]	Closed	110,000	9:1
Criteo online advertising campaign	[29]	Open	29,105,905	7:1
Direct mail in financial services	[25]	Closed	1,144,000	5:1
Simulation study	[18]	Closed	100,000	4:1
Customer retention mail in insurance	[30]	Closed	11,968	2:1
Catalog marketing	[12]	Closed	441,000	2:1
E-mail promotion in merchandising	[31]	Open	64,000	2:1
E-mail promotion in holiday marketing	[32]	Closed	282,277	1:1

data volumes of customer interactions in online shops. However, large-scale experimentation imposes substantial costs by randomly withholding profitable treatment for a sizeable control group. We reason that large experiment sizes indicate that companies perceive potential gains from causal modeling and are willing to collect data on a sufficient scale. Since the costs of experimentation are a result of the randomization of treatment assignment, we propose that supervised randomization can lead to cost reductions that are economically relevant in practice given the scale of experimentation. The savings potential increases with the targeting cost and will thus be most effective for catalog or telephone marketing, where resource-intensive treatments drive cost, and in customer churn management, where targeting customers may increase awareness of contract expiration and induce churn in otherwise passive customers. We further observe that 10 of 11 datasets show a substantial difference in size between the treatment and control group, which we denote *imbalanced full randomization*. The imbalance implies that companies assign customers to the treatment group with probabilities 2-17 times higher than assignment to the control group. The observed imbalanced experimental design is more efficient than equal assignment to treatment and control group when the marketing action is expected to be profitable on average and treatment is the dominant targeting strategy. Companies are thus conducting active cost management of random experiments based on an assessment of overall treatment effectiveness. Our approach follows the same motivation, but extends cost management to the individual level based on an assessment of the individual treatment effectiveness.

The design of randomization on the individual level is more thoroughly discussed in the medical literature [33]. On the one hand, administering a pharmaceutical to a random patient can induce severe health issues, so randomized trials pose a risk for patient health. On the other hand, new treatment may prove to be a substantial improvement over comparative options, so that withholding treatment can be seen as suboptimal care. The latter concern for optimal treatment of patients has motivated research on adaptive randomization procedures, where patients are more likely to be assigned to treatment for which positive outcomes have been previously observed over the course of the study [34, 35]. Response-adaptive randomization in medical trials differs from our approach in that we use a scoring model to adjust treatment probability conditional on customer characteristics rather than observing treatment outcomes during the trial.

The trade-off between collecting more data to improve the scoring model and applying an existing treatment policy deterministically corresponds to the exploration-exploitation problem in reinforcement learning and multi-armed bandit approaches. Supervised randomization is related to the ϵ -greedy algorithm [36], extended by heterogeneous exploration probabilities $\epsilon_i = 1 - e(x)$. In comparison to upper confidence bound sampling or Thompson sampling [36] which favor exploration of uncertain predictions, supervised randomization favors exploration close to the decision boundary of the policy and facilitates straightforward logging of the true treatment probability.

This study considers supervised randomization for continuous evaluation and periodical updating of treatment effect models. We do not adapt the scoring model and conditional treatment probabilities during the duration of the experiment as opposed to online learning of the treatment effect model under reinforcement learning. We leave a more in-depth comparison for future research.

Supervised randomization introduces dependency between the covariates and treatment assignment in the data as a side effect of adjusting treatment probabilities on the individual level. This violates the conditional independence assumption and without correction would lead to biased treatment estimates known as selection bias. An intuitive interpretation is that selection bias is due to the covariates being non-identically distributed between the treatment and control group because group assignment is itself based on the observed covariates. Statistical analysis of the average or individualized treatment effect on data that violates the unconfoundedness assumption thus requires correction for

the effect of the covariates on the individual probability to receive treatment. To the best of our knowledge, methods to systematically correct for selection bias have not yet been studied in the uplift community. Instead, research on uplift modeling assumes the feasibility of randomized control trials where there is no such selection bias by design, but ignores the associated costs of data collection in practice.

However, selection bias corrections are well-understood and common for research using observational data, where random treatment assignment is not ethical or feasible. The most common technique for selection bias correction are the inverse probability weighting estimator [37] (IPW) and its extension to the doubly robust (DR) estimator [38]. Selection bias correction has recently been integrated into popular ITE estimators, which are applied in the observational studies prevalent in economics research [22, 15].

Within the field of information systems, the IPW correction is used in observational studies [39] and research on recommender systems and reinforcement learning. In recommendation settings, explicit and some forms of implicit feedback can be understood as the outcome of a non-randomized experiment, where users evaluate a subset of items that they select based on their preferences. The customer feedback can be corrected by weighting the feedback by the probability of a customer to interact with an item before rating it [40].

Reinforcement learning research makes regular use of existing log data, which is cost-efficient to collect and available at scale, but not fully randomized. Under some conditions, unbiased training or evaluation of a learning algorithm is possible using IPW to correct for the treatment policy at the time of data collection [23]. Interestingly, if treatment under the existing policy is stochastic with a known probability and the treatment probability is logged, the resulting process can be seen as an online version of the supervised randomization process.

For observational data, the treatment probability used to correct for the selection bias is unknown and IPW thus follows a two-step process. In the first step, the treatment probabilities used for IPW correction are estimated from the observed data. In the second step, the treatment probability estimates are used to correct subsequent estimates of treatment effects. The approach proposed here is substantially different from the common applications of IPW in the first step. Under supervised randomization, the true treatment probability for each observation is actively controlled and thus known, eliminating the need for estimation and the potential for estimation error through unobserved variable bias or model misspecification.

4 Efficiently Randomized Experimental Design

Within this study, we take a holistic view of the causal modeling process emphasizing the interaction between data collection and model building. Data collection through a randomized controlled trial is a necessary part of the causal modeling process. To collect RCT data, the established targeting policy is temporarily replaced by randomized treatment assignment. However, the replacement of an efficient targeting strategy by a random assignment has negative side effects in practice, even when restricted to a subset of customers. First, randomized treatment assignment carries opportunity costs resulting from targeting the wrong customers. Compared to an existing effective targeting strategy, profitable customers are less likely to be targeted while less profitable customers are more likely to be targeted. Second, customers may misconceive data collection periods as a decrease in service or advertising quality. Since customers are not informed about the temporary replacement of the targeting model, they will attribute the random treatment assignment to the targeting efforts of the company.

Instead of replacing the established targeting policy with randomized treatment allocation, we propose to introduce a stochastic component to the existing targeting model as *supervised randomization*. Under supervised randomization, treatment assignment is largely driven by the effective targeting model but sufficiently randomized to allow the estimation and evaluation of causal models. Embedding the existing targeting model into the experimental design has three merits.

First, supervised randomization increases the return on treatment during experimentation when compared to full randomization. Supervised randomization allows us to actively decrease the cost of running randomization experiments by treating profitable customers identified by the targeting model with a higher probability than customers identified as less profitable by the scoring model.

Second, supervised randomization with conservative propensity mapping could facilitate continuous experimentation. Since both training and evaluation of causal models require experimental data, regular repetition of experimental data collection is necessary when causal models are deployed. Continuous experimentation could reduce variability in service from the perspective of the customer and streamline data collection for the company by avoiding interference with operation to run experiments and reducing the need to justify and approve the need of data collection at intervals. These goals are shared with reinforcement or bandit learning with the important difference that our approach facilitates model estimation through standard machine learning or uplift methodology.




Figure 1: Experimental design of full randomization (left) and supervised randomization (right). Note the heterogeneity in treatment probability for supervised randomization. *Rand.* indicates random assignment

4.1 Supervised randomization

Supervised randomization introduces heterogeneity in treatment probabilities into a randomized controlled experiment. We discuss the proposed approach as an extension to A/B testing through the introduction of a targeting model in several steps. We describe the process for one treatment and one control group in the online context where customers arrive in sequence, but note that the same process extends to more than one treatment group and other static settings. A/B testing for treatment evaluation is an instance of randomized controlled experiments with a single treatment. Each arriving customer is randomly assigned to the treatment or control group. The probability to receive treatment is identical for all customers $e(x) = e$ with the probability for assignment to the control group $1 - e$. The probability of treatment assignment can be equal $e = 1 - e = 0.5$ or imbalanced towards the preferred strategy for $e \in (0; 1)$. As discussed in the literature review, imbalanced probabilities are used to control the costs of the experiment in practice. For the case of multiple treatments, a different probability can be assigned to each treatment.

During regular business operation, the existing scoring model assigns a score $S(x)$ to each customer, where $S(x)$ could be an estimate of the conversion probability or the ITE. The model score $S(x)$ is compared to a threshold θ to classify customers into groups, where the group *high potentials* consists of the customers with the higher score, e.g., the highest probability to respond positively to the treatment. The *high potential* group would be targeted during regular operation, while the *low potential* group would receive no treatment. Figure 1a visualizes a scoring model during A/B testing. For the purpose of experimental data collection, the classification and deterministic targeting is replaced by random targeting. Independent of group assignment, each customer has an equal probability to receive treatment $e(x) = e$.

The proposed process of supervised randomization (Algorithm 1) integrates the scoring model into the randomized treatment assignment. As an intermediate step, let the treatment probability be dependent on the classification by the targeting strategy as depicted in Figure 1b. Different to the A/B test described in Figure 1a, where the targeting policy does not affect the treatment assignment, we now treat high potential customers with probability e_h and low potential customers with probability e_l , where $e_h \neq e_l$ and $e_l, e_h \in (0; 1)$. Note that e_h and e_l do not need to sum up to 1. We increase the treatment probability in the high potential group relative to the low potential group by choosing e_h and e_l so that $e_h > e_l$. Thereby, more high potential customers than low potential customers are treated, in accordance with the scoring model and approximating the regular targeting policy. Simultaneously, we preserve a degree of randomization in the treatment/control assignment, since each customer has some probability to be assigned to the treatment or control group, respectively. The randomization is required to fulfill the overlap assumption (Eq. 2) and should be large enough in practice to ensure coverage over the range of customer characteristics in both the treatment and control group. By

Algorithm 1: Supervised Randomization for a Controlled Experiment with K Treatments**Input:** Scoring model $S(\cdot)$; Treatment probability mapping $M(\cdot)$ **Output:** Treatment probability $e_{i,k}$; Treatment assignment $D_i \in \{0, 1, \dots, K\}$; Outcome Y_i

```

for  $i = 1, \dots, N$  do
  Observe customer  $X_i$ 
  Calculate customer score  $s_{i,k} = S(X_i)$ 
  Set treatment probability  $e_{i,k} = M(s_{i,k})$ 
  Draw treatment  $D_i \sim \text{Categorical}(e_{i,k})$ 
  if  $D_i == 0$  then
    Do not treat individual  $i$ 
    Observe outcome  $Y_i(0)$ 
  else
    for  $k$  in  $1, \dots, K$  do
      if  $D_i == k$  then
        Treat individual  $i$  with treatment  $k$ 
        Observe outcome  $Y_i(k)$ 
      end
    end
  end
end

```

violating the overlap assumption and setting $e_h = 1$ and $e_l = 0$, we recover the deterministic targeting policy of the classification model, where only customers in the *high potential* group are treated. Note that if we choose a constant treatment probability $e_h = e_l$, the process simplifies to an A/B test on the whole population as shown in Figure 1a.

We can further approximate individualized targeting by introducing more groups, each with a unique treatment probability e_k . Define a set of thresholds $[\theta_1, \theta_2, \dots, \theta_K]$ and corresponding treatment probabilities $[e_1, e_2, \dots, e_K]$ to target customer i with probability e_k for which $\theta_{k-1} < S(x_i) \leq \theta_k$. As above, we require $e_k \in (0; 1)$ and $\sum_k e_k = 1$. By increasing the number of thresholds K , we approximate a continuous mapping $M : S(x) \rightarrow e$, where each customer is assigned an individual treatment probability e_i proportional to the individual model score.

The specific mapping from model scores to treatment probability should follow the requirements of the application. We propose to determine the mapping by defining a set of k equal-sized intervals on the range of the model score in the training data $[\min(S(X_{train})), \max(S(X_{train}))]$ and assigning a linearly increasing treatment probability e_k to each interval, while setting the lowest treatment probability at $e_1 = 0.05$ and the highest at $e_K = 0.95$. Note that asymmetric mappings result in a controlled shift of average treatment probability. The design of the mapping thus allows the straightforward extension to imbalanced supervised randomization. We reiterate that supervised randomization (Algorithm 1) randomly assigns each customer to the treatment or control group, but adjusts the probability of this assignment based on the output of a scoring model, so that customers with higher score are treated with higher probability. The assigned individual treatment probabilities are logged and used in the subsequent analyses.

4.2 Inverse probability weighting

Using the targeting model to adjust the individual probability to receive treatment introduces a sampling bias into the experiment. The sampling bias is a direct result from the violation of independence between treatment probability and the individual characteristics via the scoring model. This type of selection bias commonly occurs in observational studies, where customers self-select into the treatment group, or in natural experiments. In both situations, the sample shows measurable distributional differences between the control and treatment group. Subsequent evaluation or model building need to correct for the selection bias to ensure unbiased estimates of the treatment effect. We will discuss IPW as a method that is easily integrated into model building and evaluation and discuss the doubly robust estimator as a recent extension. For a comprehensive overview of approaches including IPW see [9]. The idea underlying all approaches is to weight each observation in the treatment or control group by the inverse of its respective probability to be assigned to the observed group.

In contrast to observational studies where the treatment probability is estimated, the true probability at which customers receive the treatment is assigned actively based on a scoring model and a set of observed variables and is consequently known exactly under supervised randomization. Without the need to estimate the treatment probability from the data, we avoid confoundedness due to unobserved variables or misspecification of the propensity model by design.

IPW restores the hypothetical distribution as it would look like in a fully randomized experiments by weighting every customer with regard to the individual treatment probability. Intuitively, customers who were assigned by chance to the treatment group, even though their characteristics result in a low treatment probability, are underrepresented in the treatment group. IPW assigns these customers a higher weight. For example, if the probability of being in the treatment group for a customer is $e(x) = 0.2$ then the observed outcome if this customer received treatment is multiplied by $1/e(x) = 1/0.2 = 5$. Vice versa, if the same customer was assigned to the control group, which happened with a probability of $1 - e(x) = 0.8$, the customer's outcome in the control group is weighted by $1/0.8 = 1.25$. The IPW corrected ATE can then be estimated as:

$$\widehat{ATE}_{IPW} = \frac{1}{N} \left(\sum_{i=1}^N \frac{D_i Y_i}{e(X_i)} - \sum_{i=1}^N \frac{(1 - D_i) Y_i}{1 - e(X_i)} \right). \quad (3)$$

In observational studies, the propensity scores are unknown and need to be estimated from observed covariates. The *doubly robust* (DR) estimator is consistent and unbiased if only one of the models, the regression or the propensity score, is correctly specified [41]:

$$\widehat{ATE}_{DR} = \frac{1}{N} \sum_{i=1}^n \frac{D_i Y_i - (D_i - e(X_i)) g_1(X_i)}{e(X_i)} - \frac{1}{N} \sum_{i=1}^n \frac{(1 - D_i) Y_i + (D_i - e(X_i)) g_0(X_i)}{1 - e(X_i)}. \quad (4)$$

Here $g_D(X_i) = E(Y|D, X_i = x)$ are models of the outcome variable on x , estimated separately for $D \in \{0, 1\}$.

Adjusting for the propensity score under full randomization has no effect on the point estimate for the average treatment effect. However, there is some evidence that even in fully randomized experiments the large-sample variance of the estimate can be reduced by using estimated propensity scores to control for random imbalance in covariates as well as orthogonalization with the mean as in the doubly robust [42].

5 Empirical Evaluation

We evaluate the proposed randomization procedures through a simulation study designed to represent a direct marketing setting, which is a common application of uplift modeling in marketing (see [6, 5]).¹

Since IPW correction with the true propensity score is feasible, ATE and ITE estimates are consistent under supervised randomization. The goal of the empirical study is to compare the increased conversion rate and cost savings due to supervised randomization with the loss in efficiency due to a less balanced sample. The efficiency of each randomization procedure has two dimensions, that is, 1) monetary cost of the experiment and 2) the quality of models trained on the data collected during the experiment measured on downstream tasks.

First, the campaign profit during the experiment provides a metric on which to compare the opportunity cost of different experimental designs. We compare the campaign profit under supervised randomization to the baseline of full randomization, which provides optimal data quality, and expect opportunity costs to be lower under the proposed supervised randomization.

Second, we evaluate the data generated from the experiment by comparing the predictive performance of ITE estimators trained on data under supervised randomization to the same estimators trained on data under full randomization. Our metrics of model performance are the mean absolute error to the true treatment effect (MAE), which is known in this simulation study but unknown in real-world settings, and the Qini coefficient, which is a standard metric in the uplift literature. The Qini coefficient is a rank metric similar to model lift based on the group-wise difference in conversion rates for customers ranked by their estimated treatment effect [6].

5.1 Simulation design

We compare the ATE and ITE estimates on experimental data collected under full and supervised randomization. An online evaluation of randomization procedures is challenging since it requires running a randomized experiment for each experimental design. We therefore evaluate the supervised randomization design in an offline study and leave online testing for future research. Our empirical Monte Carlo study uses real data to the extent possible to ensure a realistic setting in which we simulate the treatment effect and have full control over the treatment assignment [43, 9]. The UCI Bank Marketing dataset [44] provides data on 45,211 customers of a Portuguese bank through 17 continuous or categorical variables covering individual socio-demographic and financial information, campaign details

¹The R code for the empirical evaluation is available at https://github.com/Humboldt-WI/supervised_randomization.

and macroeconomic indicators. All customers were subject to a phone marketing campaign promoting a term deposit and the target variable indicates if a customer has agreed to a deposit following the campaign.

Based on the available data, we simulate the individual treatment effect and hypothetical outcomes following the procedure of [43]. The treatment effect in real data can be a complex, non-linear function of a subset of observed variables and unobserved variables [17]. Therefore, we simulate the treatment effect as a combination of the twelve variables containing personal or macroeconomic information. The treatment effect as a non-linear combination of covariates is then modelled by a neural network of one hidden layer with the number of nodes equal to the number of input variables and sigmoid activation, initialized with random weights drawn from a standard Gaussian. To simulate the existence of unobserved covariates, e.g. due to privacy concerns, we remove variables with personal information on the customers’ age and marital status from the subsequent analysis.

In marketing settings, we further expect the ATE to be positive but small and the ITE to be mostly non-negative as marketing theory suggests a direct marketing campaign to increase overall conversion, with potentially zero but rarely negative impact on customers [12]. We center the simulated ITE distribution at an ATE of 5% and scale the standard deviation to 0.04 for 89% of simulated ITE to be positive. For our application, an ATE of 5% implies that the telephone campaign will convince an additional 5% of randomly targeted customer to register a term deposit. Because all customers in the observed data have received the marketing treatment, we simulate the potential outcome without treatment by flipping outcome labels for observations chosen randomly in proportion to their treatment effect as in [43].

Supervised randomization integrates an existing customer scoring model into the experimental design. A more accurate estimate from the existing model increases the extent to which potential cost savings are realized during experimentation. Noisier estimates of the scoring model lead to treatment assignment that is less profitable but has less influence on downstream tasks. In particular, supervised randomization with a noisy scoring model samples more evenly in the covariate space, thus mitigating the efficiency loss in downstream tasks, assuming a stochastic estimation error. We control the quality of the existing targeting model by simulating a noisy causal model with predictions $\hat{\tau}_i = \tau_i + \varepsilon$ and $\varepsilon \sim \mathcal{N}(0, \sigma)$. We report results for $\sigma=0.025$, such that customers with ITE equal to the ATE have a 95% chance to receive a prediction in the range $[0,0.1]$ between the true and predicted treatment effect to provide a conservative estimate of the cost savings from supervised randomization. We split the data into four folds for cross validation and randomly assign treatment to each observation in the training data according to full or supervised randomization. For ITE estimation, we then estimate the ITE model on the training data and evaluate its prediction on the holdout fold. Since the random treatment assignment introduces additional randomness into the evaluation, we repeat the treatment assignment 50 times for each holdout fold and report the average over a total of 200 repetitions.

5.2 Statistical model performance analysis

We first establish the effectiveness of supervised randomization independent of any application-specific cost setting. We evaluate the cost efficiency during experimentation through a comparison of conversion rates, ATE estimates based on their variance and ITE estimates based on uplift-specific performance metrics.

For cost efficiency, Table 2 reports the mean target fraction and conversion rate for full randomization at equal probability, full but imbalanced randomization with treatment probability 0.66 and the proposed supervised randomization procedure. We provide statistics on targeting no or all customers for context. However, targeting no or all customers and other non-randomized targeting strategies do not allow experimental data collection. In other words, settings *None* and *All* are inapplicable in practice for targeting policy evaluation or treatment effect estimation.

Table 2: Ratio of targeted customers and corresponding conversion rate under each randomization procedure. *None/All* denote targeting no/all customers for reference

	None	Full	<i>Supervised</i>	Full (Imb.)	All
Targeted Fraction of Customers	0.000	0.500	<i>0.500</i>	0.666	1.000
Conversion Rate	0.109	0.135	<i>0.143</i>	0.143	0.160

The target fractions for full randomization is 0.5 by definition and for imbalanced full randomization 0.66 by design. Small deviations from the target fraction are possible since treatment assignment is randomized on the individual level. The direction and ratio of the imbalance between the size of treatment to control group are in practice set by the experimenter to match the expected average treatment effect or marketing requirements, e.g. campaign budget. We chose a ratio of 2:1 in favor of targeting a larger group of customers following the most common design observed for customer targeting data in related research (see Table 1).

The conversion rate under each randomization provides an indirect measure of the campaign success with a higher conversion rate as an indicator of monetary returns. The increase in conversion rate from targeting no customers at 10.9% to targeting all customers at 16% reflects the simulated positive average treatment effect, specifically that customers are on average 5 percentage points more likely to convert after receiving the marketing treatment. During a fully randomized experiment, we observe an increase in the overall conversion rate by 2.6 percentage points to 13.5%. At the same fraction of customers targeted, the proposed supervised randomization increases the conversion rate by another 0.8 points to 14.3%. The improvement due to supervised randomization is the direct result of adjusting each customer’s probability to be treated based on the targeting model and targeting customers with a high predicted treatment effect.

The benchmark strategy, imbalanced full randomization, increases treatment probability for all customers indiscriminately. The increase in individual treatment probability results in a conversion rate increase by 0.8 percentage points, identical to the increase under full randomization, but at a higher fraction of customers targeted. The managerial implication is that supervised randomization achieves the same conversion rate as current best practice, while reducing the targeting rate with its associated costs by 24%.

The higher conversion rate from targeting randomization towards relevant customers comes at the downside of collecting less data for customer groups with very high or very low treatment probability. While we can use the logged treatment probabilities to optimally correct for the sampling bias that is introduced by supervised randomization, estimates of the treatment effect will exhibit higher uncertainty through higher variance. Figure 2 shows the estimated ATE under each randomization procedure. We see that 1) deviations from full randomization in the form of imbalanced full randomization and supervised randomization return unbiased estimates and 2) the overall variance from the true value and the number of extreme deviations increases when moving from full randomization to supervised randomization.

A Kruskal-Wallis test verifies that there is no significant difference in the mean point estimate among the four settings ($df=3$, $\chi^2 = 1.00$). We are thus able to verify the theoretical exposition and to show that the selection bias introduced by supervised randomization can be corrected for by applying either IPW or DR as described above. The additional uncertainty due to supervised randomization is less pronounced when using DR to correct for heterogeneous treatment probabilities instead of IPW. DR estimates exhibit a significantly lower variance when compared to IPW estimates, based on a Levene-test for homogeneity of variance ($df=1$, $F=10.29$).




Figure 2: Estimated average treatment effect averaged over 200 iterations for each randomization procedure. The dashed horizontal line denotes the simulated true average treatment effect, dots within each boxplot denote the mean estimated ATE

As uplift applications are concerned with the estimation of individualized treatment effects for customer scoring, we proceed to evaluate the model performance of two causal models on the data collected under each randomization procedure. We select the two-model approach using logistic regression and the causal forest and report the performance of using the ATE as a constant prediction for reference. Since our focus is on the comparison of the randomization procedures rather than a comparison of ITE estimators, we manually set the parameters for the causal forest as follows: number of variables tried at each split ($mtry$) = 7, number of trees = 500, minimum node size = 20 and sample fraction for honest tree building = 0.5. Model predictions are evaluated using the mean absolute error to the true ITE and the Qini score on holdout data.

Table 3: Average profit-agnostic performance of causal models for each randomization procedure. We evaluate the models using the MAE to the (simulated) true treatment effect (lower is better) and the Qini coefficient (higher is better)

	ATE		Two-Model (LR)		Causal Forest	
	MAE	Qini	MAE	Qini	MAE	Qini
Full	0.0324	-	0.0353	0.0045	0.0276	0.0056
Full (imb.)	0.0324	-	0.0357	0.0045	0.0275	0.0057
<i>Supervised</i>	<i>0.0325</i>	-	<i>0.0383</i>	<i>0.0041</i>	<i>0.0295</i>	<i>0.0047</i>

We identify two takeaways in Table 3. First, the causal forest outperforms the two-learner approach on both MAE and Qini. The performance difference is consistent over all randomization procedures with the causal forest resulting in a MAE lower by about 0.008 points and a Qini higher by 0.001 points. Second, we observe that deviating from full randomization to supervised randomization leads to the expected decrease in model performance. Under supervised randomization, the difference to full randomization for the two-model approach is 0.003 points MAE and 0.0004 points Qini and for the causal forest 0.002 points MAE and 0.001 points Qini. Imbalanced full randomization at $e = 0.66$ shows no substantial performance decrease compared to balanced full randomization, although additional experiments indicate lower performance at higher levels of imbalance. The subsequent profit-based analysis aims to provide a comprehensible evaluation of the observed differences in a business context.

5.3 Profit analysis

We proceed to empirically show the extent to which supervised randomization can reduce the cost of running a randomized experiment and the size of the expected trade-off measured by the performance of models trained on the collected data. The profit setting for telephone marketing is described by the gross profit resulting from a conversion and the variable contact cost of making a call to the customer. If we assume a constant interest margin for the bank, the gross profit from a one-year term deposit Ω is equivalent to the net interest margin m and the deposit amount A , $\Omega_i = mA_i$.

The incremental gross profit due to a marketing campaign is defined as change in the conversion probability, the treatment effect, to earn the gross profit on conversion minus the contact cost c , i.e. $\Delta\Omega_i = \tau_i mA_i - c$.

Given an accurate estimate of the treatment effect τ_i , the decision to target a specific customer is profitable when the predicted incremental gross profit for the customer is positive, i.e. $\hat{\tau}_i mA_i - c > 0$.

To simplify interpretation, we consider cost ratios in the range of $[5, 10, \dots, 50]$ to 1. Evaluation over a range of cost settings ensures the robustness of our results and allows generalization to a variety of profit and cost scenarios that may arise across banks or industries, e.g. for catalog marketing. We can empirically confirm the plausibility of the range of cost ratios by analyzing the ratio of customers which are targeted under each cost setting. For cost ratios below 10:1 and above 50:1, individual targeting policies are dominated by indiscriminate targeting of no or all customers, respectively. The cost ratio corresponds to different values of the interest margin m and deposit amount A at standardized contact cost. Assuming a constant amount of the term deposit \bar{A} for each customer, the cost ratio can be interpreted as the ratio between the gross profit over a range of interest margins m standardized to contact costs of $c = 1$ per contact.

We evaluate the cost-saving potential of using supervised randomization during experimentation based on the campaign profit resulting from a randomized experiment for each randomization procedure. We report the campaign profit per prospective customer and the difference in campaign profit relative to full randomization in Table 4. As above, we include targeting no customers and targeting all customers for reference, but stress that non-randomized targeting strategies do not allow experimental data collection, making them inapplicable for causal modeling in practice.

The empirical results in Table 4 support the proposition that supervised randomization increases the campaign profit during experimentation relative to full randomization for the full range of conversion values we consider in this study. In relative terms, supervised randomization increases the experimental campaign profit by 7.1-9.4% compared to full randomization and by 2.9-8.2% compared to imbalanced randomization.

For a conversion value of €10, we observe a marginal profit of €0.85 per customer under full randomization and a marginal profit of €0.93 under supervised randomization. The absolute increase in campaign profit is more pronounced when the cost ratio is higher. A value of €50 corresponds to a marginal profit per customer of €6.23 under full randomization compared to €6.67 under the proposed supervised randomization. Cost savings per customer compared to full randomization amount to €0.08 and €0.44, respectively.

We translate the per customer savings to an experimental campaign of 40,000 prospective customers, who are randomly

Table 4: Campaign profit (per customer) for randomized experiments under each randomization procedure and across purchase margins. *None/All* denote targeting no/all customers for reference. *Full (Imb.)* denotes full randomization with a treatment probability of 66%

Conversion Value (€)	Campaign profit per customer (€)				
	None	Full	<i>Supervised</i>	Full (Imb.)	All
10	1.09	0.85	<i>0.93</i>	0.76	0.60
15	1.64	1.52	<i>1.65</i>	1.48	1.40
20	2.18	2.19	<i>2.37</i>	2.19	2.20
25	2.73	2.86	<i>3.09</i>	2.91	3.00
30	3.27	3.54	<i>3.80</i>	3.62	3.80
35	3.82	4.21	<i>4.52</i>	4.34	4.60
40	4.36	4.88	<i>5.24</i>	5.05	5.40
45	4.91	5.56	<i>5.96</i>	5.77	6.20
50	5.45	6.23	<i>6.67</i>	6.48	7.00

targeted. This is the size of the observed telephone marketing campaign and, with less observations than 9 of the 11 experimental marketing campaigns summarized in Table 1, may provide a conservative estimate. The total cost savings per experiment when replacing full randomization with supervised randomization translate to € 3,200 for a marginal profit of € 10, € 10,400 for a marginal profit of € 30 and € 17,600 for a marginal profit of € 50. Experiment costs and the related savings arise whenever data is collected for policy evaluation or (re-)estimation of the customer scoring model.

For conversion values greater or equal € 20, targeting all customers is more profitable than not targeting any customer. The imbalanced full randomization, which we identify as standard in practice, is more profitable than full randomization only at values above € 20. At these values, imbalanced randomization achieves savings of 0 to € 0.25 per customer compared to full randomization for conversion values between € 20 and € 50, respectively. Compared to imbalanced full randomization, the proposed supervised randomization generates additional cost savings per customer of about € 0.18 for all values between € 20 and € 50. Again translated to an experiment campaign of 40,000 prospective customers, the total cost savings per experiment of supervised randomization when compared to the industry-standard range from € 7,200 for a marginal profit of € 20 to € 7,600 for a marginal profit of € 50. Note that it is possible to combine supervised randomization with imbalanced targeting. Increasing the average treatment probability through a custom treatment probability mapping may further increase campaign profit in settings where treatment is highly profitable.

Having discussed the expected cost savings during experimentation, we next discuss the opportunity costs on downstream tasks associated with the increase in model uncertainty under supervised randomization. We first report the campaign profit per customer when customers are targeted by the two-model approach or causal forest and each model is trained on experimental data collected under the different randomization procedures.

Table 5 shows that the expected decrease in profit for scoring models trained on data collected under supervised randomization is small but observable in the order of 1% of the absolute campaign profit per customer. For a basket margin of € 30, the two-model logistic regressions achieve a campaign profit of € 3.83 per customer under full randomization and a campaign profit of € 3.80 under supervised randomization, a decrease of 0.8%. The causal forest achieves a campaign profit of € 3.89 when trained on data from experiments under full randomization with a decrease by 1.3% to € 3.84 under supervised randomization. Compared over all values, supervised randomization induces a decrease in per customer profit between € 0 and € 0.04 for the two-model approach and € 0 and € 0.05 for the causal forest compared to full randomization.

6 Conclusion

Customer targeting is a continuously growing and widely studied application of scoring models. While research has focused on the prediction of future customer behavior to inform decision-making, a growing research stream has established uplift models to estimate the causal effect of a marketing action on each customer based on observed customer characteristics. The training and evaluation of causal models require data collected through experiments, in

Table 5: Campaign profit using targeting models trained on data collected under each randomization procedure. We evaluate the campaign profit per customer over a range of cost ratios

Conversion Value (€)	Two-Model (Logit)			Causal Forest		
	Simple	Simple (Imb.)	<i>Supervised</i>	Simple	Simple (Imb.)	<i>Supervised</i>
10	1.06	1.06	<i>1.06</i>	1.09	1.09	<i>1.09</i>
15	1.65	1.65	<i>1.65</i>	1.66	1.67	<i>1.65</i>
20	2.33	2.33	<i>2.32</i>	2.36	2.36	<i>2.33</i>
25	3.07	3.06	<i>3.05</i>	3.11	3.12	<i>3.08</i>
30	3.83	3.82	<i>3.80</i>	3.89	3.89	<i>3.84</i>
35	4.60	4.60	<i>4.57</i>	4.67	4.67	<i>4.62</i>
40	5.38	5.38	<i>5.35</i>	5.45	5.45	<i>5.41</i>
45	6.16	6.16	<i>6.13</i>	6.24	6.24	<i>6.20</i>
50	6.95	6.95	<i>6.91</i>	7.03	7.03	<i>7.00</i>

which customers are randomly assigned to treatments. However, experimental data collection incurs high costs by temporarily replacing an established targeting policy with random targeting.

We propose supervised randomization as a solution to reduce the cost of experimentation by integrating an existing scoring model into the experimental design. By mapping model scores to individual treatment propensities, we are able to target more profitable customers while maintaining stochastic treatment assignment. An empirical Monte Carlo study on telemarketing shows that supervised targeting can reduce the cost of an experimental campaign on 40,000 prospective customers by 7.1-9.4% compared to full randomization and 2.9-8.2% compared to imbalanced randomization, depending on the specific profit-cost ratio.

Active management of treatment assignment during experimentation leads to an overrepresentation of profitable customers in the treatment group, which causes selection bias when standard estimators are applied to estimate treatment effects. We consequently summarize inverse probability weighting and doubly robust estimation as well-studied methods to correct for selection bias when estimating average and individualized treatment effects. We show that the estimated treatment effects are unbiased and provide indicators of the increase in uncertainty related to supervised randomization. Empirical evaluation indicates that higher uncertainty of the scoring model may lead to a decrease in campaign profit by 0.8-1.3% depending on the specific profit-cost ratio. Further evaluation in real-world experiments is necessary to establish net cost savings in practice.

Overall, we argue that the methodology developed in the medical and econometric literature has not yet been fully studied and applied in the uplift setting. Doubly robust estimation serves as one example of a wider set of tools to correct for selection issues in the data. We further identify experimental data collection as a fundamental part of causal modeling. We expect that supervised randomization provides a first step towards a wider analysis of practical experimental design.

References

- [1] Statista. ecommerce. <https://www.statista.com/outlook/243/ecommerce>, 2019.
- [2] Statista. Advertising spending in the catalog, mail-order houses industry in the united states. <https://www.statista.com/statistics/470620/catalog-mail-order-houses-industry-ad-spend-usa/>, 2017.
- [3] David L. Olson and Bongsug Chae. Direct marketing decision support through predictive customer response modeling. *Decision Support Systems*, 54(1):443–451, 2012.
- [4] Robin Marco Gubela, Stefan Lessmann, Johannes Haupt, Annika Baumann, Tillmann Radmer, and Fabian Gebert. Revenue uplift modeling. In *Proceedings of the 38th International Conference on Information Systems (ICIS'17)*, Seoul, South Korea, 2017.
- [5] Floris Devriendt, Darie Moldovan, and Wouter Verbeke. A literature survey and experimental evaluation of the state-of-the-art in uplift modeling: A stepping stone toward the development of prescriptive analytics. *Big Data*, 6(1):13–41, 2018.

- [6] Nicholas J Radcliffe. Using control groups to target on predicted lift: Building and assessing uplift models. *Direct Marketing Journal, Direct Marketing Association Analytics Council*, 1:14–21, 2007.
- [7] Eva Ascarza. Retention futility: targeting high risk customers might be ineffective. *Journal of Marketing Research*, 55(1), 2018.
- [8] Scott Powers, Junyang Qian, Kenneth Jung, Alejandro Schuler, Nigam H Shah, Trevor Hastie, and Robert Tibshirani. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in Medicine*, 37(11):1767–1787, 2018.
- [9] Michael C. Knaus, Michael Lechner, and Anthony Strittmatter. Machine Learning Estimation of Heterogeneous Causal Effects: Empirical Monte Carlo Evidence. *IZA Discussion Paper*, 12039, 2019-01.
- [10] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- [11] Eva Ascarza, Peter Ebbes, Oded Netzer, and Matthew Danielson. Beyond the target customer: social effects of customer relationship management campaigns. *Journal of Marketing Research*, 54(3):347–363, 2017.
- [12] Günter J Hitsch and Sanjog Misra. Heterogeneous treatment effects and optimal targeting policy evaluation, 2018.
- [13] Guido W Imbens and Jeffrey M Wooldridge. Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86, 2009-03.
- [14] Brett R Gordon, Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky. A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook. *Marketing Science*, 38(2):193–364, 2019.
- [15] Sören R Künzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019.
- [16] Robin Marco Gubela, Artem Bequé, Fabian Gebert, and Stefan Lessmann. Conversion uplift in e-commerce: a systematic benchmark of modeling strategies. *International Journal of Information Technology & Decision Making*, 18(3):747–791, 2019.
- [17] Max H. Farrell, Tengyuan Liang, and Sanjog Misra. Deep Neural Networks for Estimation and Inference: Application to Causal Effects and Other Semiparametric Estimands. *arXiv preprint arXiv:1809.09953*, 2018.
- [18] Victor SY Lo. The true lift model: a novel data mining approach to response modeling in database marketing. *ACM SIGKDD Explorations Newsletter*, 4(2):78–86, 2002.
- [19] Lukasz Zaniewicz and Szymon Jaroszewicz. Support Vector Machines for Uplift Modeling. In *Proceedings of the 13th International Conference on Data Mining Workshops*, pages 131–138. IEEE, 2013.
- [20] Piotr Rzepakowski and Szymon Jaroszewicz. Decision trees for uplift modeling with single and multiple treatments. *Knowledge and Information Systems*, 32(2):303–327, 2012.
- [21] Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.
- [22] Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- [23] Adith Swaminathan and Thorsten Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research*, 16:1731-1755, 2015.
- [24] Susan Athey and Stefan Wager. Efficient policy learning. *arXiv preprint arXiv:1702.02896v2*, 2017.
- [25] Kathleen Kane, Victor SY Lo, and Jane Zheng. Mining for the truly responsive customers and prospects using true-lift modeling: Comparison of new and existing methods. *Journal of Marketing Analytics*, 2(4):218–238, 2014.
- [26] Behram J Hansotia and Bradley Rukstales. Direct marketing for multichannel retailers: Issues, challenges and solutions. *Journal of Database Marketing & Customer Strategy Management*, 9(3):259–266, 2002.
- [27] Leo Guelman, Montserrat Guillén, and Ana María Pérez Marín. Optimal personalized treatment rules for marketing interventions: A review of methods, a new proposal, and an insurance case study. *UB Riskcenter Working Paper Series*, 2014/06, 2014.
- [28] David Maxwell Chickering and David Heckerman. A decision theoretic approach to targeted advertising. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pages 82–88. Morgan Kaufmann Publishers Inc., 2000.

- [29] Eustache Diemert, Artem Betlei, Christophe Renaudin, and Massih-Reza Amini. A large scale benchmark for uplift modeling. In *Proceedings of the AdKDD and TargetAd Workshop, KDD*, London, United Kingdom, 2018. ACM.
- [30] Leo Guelman, Montserrat Guillén, and Ana M Pérez-Marín. Uplift random forests. *Cybernetics and Systems*, 46(3-4):230–248, 2015.
- [31] Kevin Hillstrom. The minethatdata e-mail analytics and data mining challenge. MineThatData Blog. <https://blog.minethatdata.com/2008/03/minethatdata-e-mail-analytics-and-data.html>, 2008.
- [32] Behram Hansotia and Brad Rukstales. Incremental value modeling. *Journal of Interactive Marketing*, 16(3):35–46, 2002.
- [33] Kenneth F. Schulz and David A. Grimes. Generation of allocation sequences in randomised trials: chance, not choice. *The Lancet*, 359(9305):515–519, 2002.
- [34] John M. Lachin, John P. Matts, and L. J. Wei. Randomization in clinical trials: Conclusions and recommendations. *Controlled Clinical Trials*, 9(4):365–374, 1988.
- [35] William F. Rosenberger and John M. Lachin. The use of response-adaptive designs in clinical trials. *Controlled Clinical Trials*, 14(6):471–484, 1993.
- [36] Eric M. Schwartz, Eric T. Bradlow, and Peter S. Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017-07.
- [37] Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- [38] James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.
- [39] Marco Caliendo, Michel Clement, Dominik Papies, and Sabine Scheel-Kopeinig. The cost impact of spam filters: Measuring the effect of information system technologies in organizations. *Information Systems Research*, 23(3):1068–1080, 2012.
- [40] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48, pages 1670–1679, 2016.
- [41] Jared K Lunceford and Marie Davidian. Stratification and weighting via the propensity score in estimation of causal treatment effects: A comparative study. *Statistics in Medicine*, 23(19):2937–2960, 2004.
- [42] Elizabeth J Williamson, Andrew Forbes, and Ian R White. Variance reduction in randomised trials by inverse probability weighting using the propensity score. *Statistics in Medicine*, 33(5):721–737, 2014.
- [43] Xinkun Nie and Stefan Wager. Quasi-Oracle Estimation of Heterogeneous Treatment Effects. *arXiv preprint arXiv:1712.04912*, 2017.
- [44] Sérgio Moro, Paulo Cortez, and Paulo Rita. A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems*, 62:22–31, 2014.

IRTG 1792 Discussion Paper Series 2019



For a complete list of Discussion Papers published, please visit
<http://irtg1792.hu-berlin.de>.

- 001 "Cooling Measures and Housing Wealth: Evidence from Singapore" by Wolfgang Karl Härdle, Rainer Schulz, Taojun Xie, January 2019.
- 002 "Information Arrival, News Sentiment, Volatilities and Jumps of Intraday Returns" by Ya Qian, Jun Tu, Wolfgang Karl Härdle, January 2019.
- 003 "Estimating low sampling frequency risk measure by high-frequency data" by Niels Wesselhöfft, Wolfgang K. Härdle, January 2019.
- 004 "Constrained Kelly portfolios under alpha-stable laws" by Niels Wesselhöfft, Wolfgang K. Härdle, January 2019.
- 005 "Usage Continuance in Software-as-a-Service" by Elias Baumann, Jana Kern, Stefan Lessmann, February 2019.
- 006 "Adaptive Nonparametric Community Detection" by Larisa Adamyan, Kirill Efimov, Vladimir Spokoiny, February 2019.
- 007 "Localizing Multivariate CAViaR" by Yegor Klochkov, Wolfgang K. Härdle, Xiu Xu, March 2019.
- 008 "Forex Exchange Rate Forecasting Using Deep Recurrent Neural Networks" by Alexander J. Dautel, Wolfgang K. Härdle, Stefan Lessmann, Hsin-Vonn Seow, March 2019.
- 009 "Dynamic Network Perspective of Cryptocurrencies" by Li Guo, Yubo Tao, Wolfgang K. Härdle, April 2019.
- 010 "Understanding the Role of Housing in Inequality and Social Mobility" by Yang Tang, Xinwen Ni, April 2019.
- 011 "The role of medical expenses in the saving decision of elderly: a life cycle model" by Xinwen Ni, April 2019.
- 012 "Voting for Health Insurance Policy: the U.S. versus Europe" by Xinwen Ni, April 2019.
- 013 "Inference of Break-Points in High-Dimensional Time Series" by Likai Chen, Weining Wang, Wei Biao Wu, May 2019.
- 014 "Forecasting in Blockchain-based Local Energy Markets" by Michael Kostmann, Wolfgang K. Härdle, June 2019.
- 015 "Media-expressed tone, Option Characteristics, and Stock Return Predictability" by Cathy Yi-Hsuan Chen, Matthias R. Fengler, Wolfgang K. Härdle, Yanchu Liu, June 2019.
- 016 "What makes cryptocurrencies special? Investor sentiment and return predictability during the bubble" by Cathy Yi-Hsuan Chen, Roméo Després, Li Guo, Thomas Renault, June 2019.
- 017 "Portmanteau Test and Simultaneous Inference for Serial Covariances" by Han Xiao, Wei Biao Wu, July 2019.
- 018 "Phenotypic convergence of cryptocurrencies" by Daniel Traian Pele, Niels Wesselhöfft, Wolfgang K. Härdle, Michalis Kolossiatis, Yannis Yatracos, July 2019.
- 019 "Modelling Systemic Risk Using Neural Network Quantile Regression" by Georg Keilbar, Weining Wang, July 2019.

IRTG 1792, Spandauer Strasse 1, D-10178 Berlin
<http://irtg1792.hu-berlin.de>

This research was supported by the Deutsche
Forschungsgemeinschaft through the IRTG 1792.

IRTG 1792 Discussion Paper Series 2019



For a complete list of Discussion Papers published, please visit
<http://irtg1792.hu-berlin.de>.

- 020 "Rise of the Machines? Intraday High-Frequency Trading Patterns of Cryptocurrencies" by Alla A. Petukhina, Raphael C. G. Reule, Wolfgang Karl Härdle, July 2019.
- 021 "FRM Financial Risk Meter" by Andrija Mihoci, Michael Althof, Cathy Yi-Hsuan Chen, Wolfgang Karl Härdle, July 2019.
- 022 "A Machine Learning Approach Towards Startup Success Prediction" by Cemre Ünal, Ioana Ceasu, September 2019.
- 023 "Can Deep Learning Predict Risky Retail Investors? A Case Study in Financial Risk Behavior Forecasting" by A. Kolesnikova, Y. Yang, S. Lessmann, T. Ma, M.-C. Sung, J.E.V. Johnson, September 2019.
- 024 "Risk of Bitcoin Market: Volatility, Jumps, and Forecasts" by Junjie Hu, Weiyu Kuo, Wolfgang Karl Härdle, October 2019.
- 025 "SONIC: SOcial Network with Influencers and Communities" by Cathy Yi-Hsuan Chen, Wolfgang Karl Härdle, Yegor Klochkov, October 2019.
- 026 "Affordable Uplift: Supervised Randomization in Controlled Experiments" by Johannes Haupt, Daniel Jacob, Robin M. Gubela, Stefan Lessmann, October 2019.

IRTG 1792, Spandauer Strasse 1, D-10178 Berlin
<http://irtg1792.hu-berlin.de>

This research was supported by the Deutsche
Forschungsgemeinschaft through the IRTG 1792.