



Partial counterfactual identification and uplift modeling: theoretical results and real-world assessment

Théo Verhelst¹ · Denis Mercier² · Jeevan Shrestha² · Gianluca Bontempi¹

Received: 1 March 2022 / Revised: 8 November 2022 / Accepted: 7 February 2023 /
Published online: 14 March 2023

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2023

Abstract

Counterfactuals are central in causal human reasoning and the scientific discovery process. The uplift, also called conditional average treatment effect, measures the causal effect of some action, or treatment, on the outcome of an individual. This paper discusses how it is possible to derive bounds on the probability of counterfactual statements based on uplift terms. First, we derive some original bounds on the probability of counterfactuals and we show that tightness of such bounds depends on the information of the feature set on the uplift term. Then, we propose a point estimator based on the assumption of conditional independence between the counterfactual outcomes. The quality of the bounds and the point estimators are assessed on synthetic data and a large real-world customer data set provided by a telecom company, showing significant improvement over the state of the art.

Keywords Causal inference · Counterfactuals · Uplift modeling · Machine learning

Denis Mercier, Jeevan Shrestha and Gianluca Bontempi have contributed equally to this work.

Editors: Longbing Cao, Joao Gama, Nitesh Chawla, Joshua Huang.

✉ Théo Verhelst
theo.verhelst@ulb.be

Denis Mercier
denis1.mercier@orange.com

Jeevan Shrestha
jeevan.shrestha@orange.com

Gianluca Bontempi
gianluca.bontempi@ulb.be

¹ Machine Learning Group, Université Libre de Bruxelles, Campus de la Plaine ULB CP212, Boulevard du Triomphe, 1050 Brussels, Belgium

² Data Science Team, Orange Belgium, Avenue du Bourget 3, 1140 Brussels, Belgium

1 Introduction

Counterfactual statements (or counterfactuals for short) concern the potential of events in situations different from the actual state of the world. An example of counterfactual statement is “I got no effect since I made no action but something would have happened had I acted”. Counterfactuals are used in many fields, ranging from algorithmic recourse (Karimi et al., 2021) to online advertisement and customer relationship management (Li & Pearl, 2019).

Counterfactuals have been formally defined in terms of structural causal models by Pearl (2009). Nevertheless, since a counterfactual statement cannot be directly observed, the research focuses on estimating or bounding their probability (e.g. the probability that we have an effect given a treatment and no effect else). The probability of some specific counterfactual expressions have been studied in the literature (Tian & Pearl, 2000) because of their relevance in causal decision-making. The *probability of necessity* (PN) is the probability that an event y would not have occurred in the absence of an action or treatment t , given that y and t in fact occurred. Conversely, the *probability of sufficiency* (PS) is the probability that event y would have occurred in the presence of an action t , given that both y and t in fact did not occur. Lastly, the *probability of necessity and sufficiency* (PNS) is the probability that the event y occurs if and only if the event t occurs.

In the case of incomplete knowledge about the causal model, identification procedures indicate when and how the probability of counterfactuals can be computed from a combination of observational data, experimental data (i.e. data with randomized treatment), and causal assumptions (Correa et al., 2021). In situations where the exact probability of counterfactuals cannot be directly computed, an alternative consists in bounding this quantity. This problem, called *partial counterfactual identification*, has first been addressed by Tian and Pearl (2000), and more recently by Mueller et al. (2021) and Zhang et al. (2022).

Counterfactual reasoning has practical applications in business, notably churn modeling: consider a company wishing to use direct marketing actions to prevent customers from churning (i.e. stop using their service). The behavior of the customers in reaction to the two possible actions (contact or not) could be described in terms of counterfactual statements (Devriendt et al., 2019):

- *Sure thing*: customer not churning regardless of the action.
- *Persuadable*: customer churning only if not contacted.
- *Do-not-disturb*: customer churning only if contacted.
- *Lost cause*: customer churning regardless of the action.

Note that the probability of a customer being a *do-not-disturb* is an example of PNS (Tian & Pearl, 2000) while, to the best of our knowledge, the other three probabilities have not been labeled in the causal inference literature. Though not observable, those quantities are relevant for adequate decision-making, and partial counterfactual identification can help in reducing the uncertainty about the possible customer behaviors.

Uplift modeling, where uplift stands for the conditional average treatment effect (CATE), or heterogeneous treatment effect (Zhang et al., 2021), is another well-known approach for estimating causal effects. It returns an estimate at the individual level of the impact of some action on the probability of the outcome. In the example of churn prevention, uplift modeling estimates the impact of a promotional offer on the probability of churn for each customer. Most recent and powerful uplift models are based on machine

learning (Curth & van der Schaar, 2021). Some uplift models expect experimental data and are based on conventional classification models (Jaskowski & Jaroszewicz, 2012; Athey & Imbens, 2016). Other models accept observational (possibly confounded) data and estimate the uplift through some sort of adjustment, for example with propensity scores in (Künzel et al., 2019) and (Curth & van der Schaar, 2021).

Counterfactuals and uplift are closely related, yet formally distinct notions. The counterfactual distribution describes the probability of each possible combination of realized and hypothetical outcomes, while the uplift describes the change in outcome probability due to the treatment. While the counterfactual distribution is more informative, it is also more difficult to estimate than the uplift. In (Li & Pearl, 2019) it is mentioned that the similarity between these two notions can lead to confusion, especially since they collapse under the assumption of monotonicity (the absence of negative causal effects).

Existing works on partial counterfactual identification (Mueller et al., 2021; Zhang et al., 2022) make structural assumptions on the causal model to derive bounds whose estimate requires a combination of experimental and observational data. In this paper, we propose some original bounds on the probability of counterfactuals based on the uplift terms. The originality of our approach consists in defining bounds that depend on terms (like uplift) for which nowadays a lot of reliable estimators exist in literature. This is of particular interest in big data applications, where structural assumptions are hard to validate but a large number of observations about individual descriptors (covariates) and past behavior are available.

The main contributions of this paper are as follows:

- A set of original bounds on the probability of counterfactuals, expressed in terms of the uplift quantity.
- A formal derivation of the relationship between our original bounds and the state-of-the-art Fréchet bounds derived by Tian and Pearl (2000).
- A point estimator of the counterfactual probabilities based on the conditional independence assumption.
- A hierarchical Bayesian model for simulating counterfactual settings and assessing the accuracy of the sample version of the derived bounds.
- A real-world assessment of the proposed bounds with a large data set of customer churn campaigns and a discussion of the potential benefits.

The rest of this paper is organized as follows. In Sect. 2, we present related work in the literature on partial counterfactual identification. In Sect. 3, we present the formalism used throughout this paper. In Sects. 4 and 5, we derive bounds and point estimates on the probability of counterfactuals. We analyze the behavior of these estimators under various conditions with simulated examples in Sect. 6. We apply our estimator to a real-world data set from our industrial partner and estimate the suggested potential benefits in Sect. 7. Conclusions and limitations are given in Sect. 8.

2 Related work

The *probability of necessity and sufficiency* (PNS) as presented by Pearl (2009 p. 286) is one of the four counterfactual probabilities that we consider in this paper. Seminal works on partial counterfactual identification include (Balke & Pearl, 1994) and (Tian &

Table 1 Mathematical notation

$Y \in \{0, 1\}$	Outcome
$T \in \{0, 1\}$	Treatment indicator
X	Set of features
$\text{do}(T = t)$	Intervention $T = t$
$P(Y_t = y)$	Probability of the outcome $Y = y$ under $\text{do}(T = t)$
S_0, S_1	$P(Y_0 = 1), P(Y_1 = 1)$
$\alpha, \beta, \gamma, \delta$	$P(Y_0 = 0, Y_1 = 0), \dots, P(Y_0 = 1, Y_1 = 1)$ (see Equations (4) to (7))
$Q(x)$	Quantity Q conditioned on $X = x$, e.g. $S_0(x) = P(Y_0 = 1 X = x)$

Pearl, 2000). They show that in the exogenous case (e.g. when the treatment is randomly assigned), the bounds on the PNS reduce to the Fréchet bounds (Fréchet, 1935). We will use these bounds as a baseline in the remaining of this paper.

The PNS conditioned on a set of covariates x is called x -specific PNS in (Li & Pearl, 2019). The main focus of Li and Pearl (2019) is the estimation of the benefit generated by a customer retention campaign when the different types of customers have different values. For example, keeping a persuadable customer (a customer who does not churn only when targeted) might be more beneficial than keeping a customer who would never leave, besides the cost of the targeted action. In (Li & Pearl, 2022), the authors further refine the bounds on the campaign benefit based on causal assumptions derived from causal diagrams.

Mueller et al. (2021) derived tighter bounds on the PNS for a variety of causal diagrams, such as with sufficient covariates or with a mediator variable. In particular, the bounds in Theorem 5 in (Mueller et al., 2021) are formally very close to the bounds we develop in this paper, although they consider a set of discrete covariates, whereas we use uplift modeling which allows for arbitrary high-dimensional covariate sets. Zhang et al. (2022) express the problem of bounding the probability of counterfactuals into polynomial programming, providing tight bounds for any causal graph and combination of experimental and observational data.

Our approach in this paper differs from Mueller et al. (2021) and Zhang et al. (2022) in that we make very few causal assumptions (only that the treatment is randomized), but we suggest uplift modeling as a powerful way to estimate conditional probabilities, and we analyze the impact of mutual information between the conditioning set and the potential outcomes.

3 Notation

In this section, we present the mathematical notation used throughout this paper. A summary is given in Table 1.

We use Pearl's causal framework, which is based on the notion of *structural causal models* (SCM). A formal definition of SCMs is given by Pearl (2009 Def. 7.1.1). In this framework, T denotes the action or *treatment*, Y the causal effect (or *outcome*), X a set of features (or *covariates*) describing the unit/individual under treatment and the $\text{do}(T = t)$ operator denotes a causal intervention in the system. In this paper, we will limit ourselves to consider binary treatments and outcomes. For example, let T be the binary variable representing a medical treatment: the notation $\text{do}(T = 1)$ indicates that the treatment is forced

on an individual regardless of whether they would have received it without explicit intervention. The conditional probability of $Y = y$ given $X = x$ under the intervention $\text{do}(T = t)$ is written $P(Y = y \mid \text{do}(T = t), X = x)$. An alternative notation consists in indicating the intervened variable as a subscript to the other variables, such as¹ $P(Y_t = y \mid X = x)$. In our application about customer churn prevention, $Y = 1$ indicates that the customer churned, X a set of descriptive features of the customer and the treatment T denotes the exposure of the customer to a targeted marketing action in the form of an e-mail or a phone call ($T = 1$ when targeted, $T = 0$ otherwise).

We note the probability of the outcome Y under intervention $\text{do}(T = 0)$ given some feature $X = x$ as

$$S_0(x) = P(Y_0 = 1 \mid X = x). \quad (1)$$

Similarly, under the intervention $\text{do}(T = 1)$ we have

$$S_1(x) = P(Y_1 = 1 \mid X = x). \quad (2)$$

The *uplift* is defined to be the difference between these probabilities:

$$U(x) = S_0(x) - S_1(x). \quad (3)$$

Note that the uplift is also sometimes defined as $U(x) = S_1(x) - S_0(x)$, depending on the context and the meaning of the outcome Y . Throughout this paper, the argument x in quantities such as $S_0(x)$ indicates the conditioning on $X = x$. If omitted, the quantity is supposed to be no longer conditioned on x (e.g. $S_0 = P(Y_0 = 1)$). Equivalently, we can consider $S_0(x)$ as a function from the domain of X to $[0, 1]$, therefore we can define S_0 as $S_0 = \mathbb{E}_X[S_0(X)]$, and similarly for S_1 and U .

The probabilities $S_0(x)$ and $S_1(x)$ cannot be estimated without further assumptions. In this paper, we make the assumption of *unconfoundedness*² (Pearl, 2009, Def. 9.2.9):

Definition 1 (Unconfoundedness) A variable Y is *unconfounded* with respect to T given X if, for any values y , t and x ,

$$P(Y = y \mid \text{do}(T = t), X = x) = P(Y = y \mid T = t, X = x)$$

Or, alternatively, if for any value t ,

$$Y_t \perp T \mid X.$$

Note that in (Pearl, 2009), the unconfoundedness is called *exogeneity* and is defined without conditioning on X . The distinction is made between *weak* exogeneity and *strong* exogeneity: Definition 1 corresponds to weak exogeneity, while strong exogeneity assumes $\{Y_0, Y_1\} \perp T \mid X$. This distinction has no impact on the results presented in this paper.

Unconfoundedness allows the estimation of the scores $S_0(x)$ and $S_1(x)$ from data, since

¹ The features X should also receive the subscript t under the intervention $\text{do}(T = t)$. In our case, the treatment is supposed to occur after the measurement of X , and has thus no effect on X . This implies that $X_t = X$.

² Also called *ignorability* by Rosenbaum and Rubin (1983), *exogeneity* by Pearl (2009), or *conditional independence assumption* by Gutierrez and Gérardy (2016).

Table 2 The four categories of customers for churn prevention in terms of counterfactual outcomes

	$Y_0 = 0$	$Y_0 = 1$
$Y_1 = 0$	Sure thing	Persuadable
$Y_1 = 1$	Do-not-disturb	Lost cause

$$S_0(x) = P(Y = 1 \mid \text{do}(T = 0), X = x) = P(Y = 1 \mid T = 0, X = x)$$

and similarly for $S_1(x)$. Unconfoundedness is guaranteed when the treatment T is randomized. In absence of randomization by using a suitable adjustment set (i.e. satisfying the *back-door criterion* (Pearl, 2009)) an estimation method could still permit the unbiased estimation of $S_0(x)$ and $S_1(x)$. Such assumption is typically made in uplift approaches integrating propensity scores, notably the X-learner (Künzel et al., 2019) and, more recently, double machine learning estimators (Jung et al., 2021).

Let us suppose that $Y_0 = 1$, i.e. we observe $Y = 1$ after having assigned the treatment $T = 0$ to a given individual. Though we cannot observe the counterfactual outcome Y_1 , we can reason about the value it would have. If $Y_1 = 0$, the treatment would have a causal impact on the outcome, since the outcome Y changes by intervening on T . Otherwise, if $Y_1 = 1$, the treatment would have no causal influence on the outcome of this individual. More generally, the joint values of Y_0 and Y_1 define four different counterfactual expressions. In this paper their probability is noted

$$\alpha = P(Y_0 = 0, Y_1 = 0) \tag{4}$$

$$\beta = P(Y_0 = 1, Y_1 = 0) \tag{5}$$

$$\gamma = P(Y_0 = 0, Y_1 = 1) \tag{6}$$

$$\delta = P(Y_0 = 1, Y_1 = 1) \tag{7}$$

From which we can derive

$$S_0 = P(Y_0 = 1) = P(Y_0 = 1, Y_1 = 0) + P(Y_0 = 1, Y_1 = 1) = \beta + \delta \tag{8}$$

$$S_1 = P(Y_1 = 1) = P(Y_0 = 0, Y_1 = 1) + P(Y_0 = 1, Y_1 = 1) = \gamma + \delta. \tag{9}$$

Note that the *probability of necessity and sufficiency* (PNS) in (Pearl, 2009) is the γ term in (6).

In customer churn prevention, the four counterfactuals may be mapped to the four categories of customers presented in the introduction (Table 2). An effective campaign should then only reach out to *persuadable* customers (whose proportion in the population is β), since *sure-thing* and *lost cause* customers would not change their minds in reaction to the marketing action, and the *do-not-disturb* would react negatively to it.

The next sections will discuss the paper’s contributions on the estimation of the probabilities α, β, γ and δ .

4 Bounds on the probability of counterfactuals

Bounds on the probability of counterfactuals have first been derived in Tian and Pearl (2000), where the authors focus on $P(Y_0 = 0 \mid T = 1, Y = 1)$, $P(Y_1 = 1 \mid T = 0, Y = 0)$, and $P(Y_0 = 0, Y_1 = 1)$ (denoted γ in (6)) under various assumptions. They showed that the quantity γ can be bounded as

$$\max\{0, P(Y_1 = 1) - P(Y_0 = 1)\} \leq \gamma \leq \min\{P(Y_0 = 0), P(Y_1 = 1)\}. \quad (10)$$

The bounds derive from the classical Fréchet bounds (Fréchet, 1935) stating that for any pair of events A and B

$$\max\{0, P(A) + P(B) - 1\} \leq P(A, B) \leq \min\{P(A), P(B)\}. \quad (11)$$

For instance, by replacing A with $Y_0 = 0$ and B with $Y_1 = 1$, it is easy to derive the inequalities (10). Tighter bounds on counterfactual probabilities are derived in (Mueller et al., 2021; Zhang et al., 2022) by making structural assumptions on the causal directed acyclic graph (DAG).

In this paper, we focus on a setting where (i) no structural assumptions may be made (besides unconfoundedness) and (ii) an estimation of the uplift is possible on the basis of historical data. For this reason, we derive a set of original bounds that depend on the conditional probabilities terms $S_0(x) = P(Y_0 = 1 \mid X = x)$ and $S_1(x) = P(Y_1 = 1 \mid X = x)$.

Our derivation consists in first generalizing the Fréchet bounds to all four counterfactual probabilities, by substituting A with $Y_0 = 0$ or $Y_0 = 1$, and B with $Y_1 = 0$ or $Y_1 = 1$:

$$\max\{0, P(Y_0 = 0) - P(Y_1 = 1)\} \leq \alpha \leq \min\{P(Y_0 = 0), P(Y_1 = 0)\} \quad (12)$$

$$\max\{0, P(Y_0 = 1) - P(Y_1 = 1)\} \leq \beta \leq \min\{P(Y_0 = 1), P(Y_1 = 0)\} \quad (13)$$

$$\max\{0, P(Y_1 = 1) - P(Y_0 = 1)\} \leq \gamma \leq \min\{P(Y_0 = 0), P(Y_1 = 1)\} \quad (14)$$

$$\max\{0, P(Y_0 = 1) - P(Y_1 = 0)\} \leq \delta \leq \min\{P(Y_0 = 1), P(Y_1 = 1)\}. \quad (15)$$

Then, we assume that a reliable estimate (e.g. by uplift modeling) of the conditional scores $S_0(x)$ and $S_1(x)$ is available. Such scores can be used to refine the bounds on α, \dots, δ by leveraging Jensen's inequality.³ We apply first Jensen's inequality to the lower bounds of Equations (12)-(15) by taking f as the $\max(0, \cdot)$ function and then to the upper bounds with f as the $\min(\cdot, \cdot)$ function. We detail here the derivation for the lower bound on β , but the same reasoning can be easily extended to the other bounds as well.

$$\max\{0, P(Y_0 = 1) - P(Y_1 = 1)\} = \max\{0, S_0 - S_1\} \quad (16)$$

$$= \max\{0, \mathbb{E}[S_0(X) - S_1(X)]\} \quad (17)$$

³ Jensen's inequality, in its probabilistic form, states that for a convex function f and a random variable X , we have

$$f(\mathbb{E}[X]) \leq \mathbb{E}[f(X)].$$

$$\leq \mathbb{E}[\max\{0, S_0(X) - S_1(X)\}] \tag{18}$$

$$\leq \mathbb{E}[\beta(X)] = \beta. \tag{19}$$

Note that the quantity in (16) is the Fréchet bound, which by Jensen’s inequality is lower than (18). It follows that our derivation returns a tighter upper bound than the Fréchet upper bound. The inequality (19), derived from (13) conditioned on $X = x$, guarantees that this is a lower bound on β . To summarize, we propose to bound α, \dots, δ as follows

$$\mathbb{E}[\max\{0, 1 - S_0(X) - S_1(X)\}] \leq \alpha \leq \mathbb{E}[\min\{1 - S_0(X), 1 - S_1(X)\}] \tag{20}$$

$$\mathbb{E}[\max\{0, S_0(X) - S_1(X)\}] \leq \beta \leq \mathbb{E}[\min\{S_0(X), 1 - S_1(X)\}] \tag{21}$$

$$\mathbb{E}[\max\{0, S_1(X) - S_0(X)\}] \leq \gamma \leq \mathbb{E}[\min\{1 - S_0(X), S_1(X)\}] \tag{22}$$

$$\mathbb{E}[\max\{0, S_0(X) + S_1(X) - 1\}] \leq \delta \leq \mathbb{E}[\min\{S_0(X), S_1(X)\}]. \tag{23}$$

Hereafter we will refer to those bounds as the *uplift bounds* (UB) since they are defined in terms of the uplift terms. To assess whether these bounds improve the state-of-the-art Fréchet bounds, we consider their respective spans (i.e. the difference between the upper and the lower bound). It can be shown that the uplift bounds span Span_{UB} is the same for all the counterfactual probabilities:

$$\text{Span}_{\text{UB}} = \mathbb{E}[\min\{S_0(X), 1 - S_1(X)\}] - \mathbb{E}[\max\{0, S_0(X) - S_1(X)\}] \tag{24}$$

$$= \mathbb{E}[\min\{S_0(X), 1 - S_1(X)\} - \max\{0, S_0(X) - S_1(X)\}] \tag{25}$$

$$= \mathbb{E}[\min\{S_0(X), 1 - S_1(X)\} + \min\{0, S_1(X) - S_0(X)\}] \tag{26}$$

$$= \mathbb{E}[\min\{S_0(X), S_1(X), 1 - S_0(X), 1 - S_1(X)\}] \tag{27}$$

Where in (26) we used the equality $-\max\{a, b\} = \min\{-a, -b\}$, and in (27) the equality $\min\{a, b\} + \min\{c, d\} = \min\{a + c, a + d, b + c, b + d\}$.

The span of the Fréchet bounds, denoted by Span_{Fr} , is equal to

$$\text{Span}_{\text{Fr}} = \min\{S_0, S_1, 1 - S_0, 1 - S_1\}$$

For all four counterfactual probabilities. Note that Span_{Fr} depends solely on the marginal terms S_0 and S_1 (i.e. the average probability of the outcome in the control and target groups) whereas Span_{UB} is a function of the descriptive features (or covariates) X . This means that in the case of informative features (i.e. when the conditional entropy of Y_0 and Y_1 is smaller than the marginal entropy), the uplift bounds are tighter than the Fréchet ones. In the case of perfect knowledge (i.e. when Y_0 and Y_1 are deterministic functions of X), $S_0(x)$ and $S_1(x)$ are either 0 or 1, the span of the uplift bounds collapses to zero and the counterfactual distribution is fully determined. In the case of noninformative features (i.e. when the conditional entropy of Y_0 and Y_1 is equal to the marginal entropy) the uplift bounds reduce to the Fréchet bounds.

Such considerations can be formalized in terms of conditional entropy by the following Theorem:

Theorem 1 *As the conditional entropy $H(Y_0, Y_1 | X)$ approaches zero, the uplift bounds on the probability $P(Y_0 = y_0, Y_1 = y_1)$ collapse to the exact value of that probability. Conversely, as the conditional entropy $H(Y_0, Y_1 | X)$ approaches the entropy $H(Y_0, Y_1)$, the uplift bounds reduce to the Fréchet bounds.*

Proof We have

$$H(Y_0, Y_1 | X) = - \int \sum_{y_0, y_1} P(y_0, y_1 | x) \log P(y_0, y_1 | x) f_X(x) dx$$

Where the sum runs over the four possible values of Y_0, Y_1 , and $f_X(x)$ is the probability density function of X . This expression can also be noted

$$H(Y_0, Y_1 | X) = - \int (\alpha(x) \log(\alpha(x)) + \beta(x) \log(\beta(x)) + \gamma(x) \log(\gamma(x)) + \delta(x) \log(\delta(x))) f_X(x) dx.$$

It is minimized (in fact, equal to zero) when one of $\alpha(x), \dots, \delta(x)$ is equal to one and the three other ones are equal to zero for all x . Also, the span of the uplift bounds is

$$\begin{aligned} \text{Span}_{\text{UB}} &= \mathbb{E}[\min\{S_0(X), S_1(X), 1 - S_0(X), 1 - S_1(X)\}] \\ &= \int \min\{\beta(x) + \delta(x), \gamma(x) + \delta(x), \alpha(x) + \gamma(x), \alpha(x) + \beta(x)\} f_X(x) dx \end{aligned}$$

When one of $\alpha(x), \dots, \delta(x)$ is equal to one and the three other values are equal to zero for all x , this expression collapses to zero, since two of the four terms in the minimum will be equal to zero. In this case, the bounds collapse to the true value of the counterfactual probability. This proves the first part of the theorem.

For the second part of the theorem, let's assume that X brings no information about Y_0, Y_1 , which we formalize as $H(Y_0, Y_1 | X) = H(Y_0, Y_1)$, or also in terms of statistical independence as $(Y_0, Y_1) \perp X$. By definition of statistical independence, we know that $P(y_0 | x) = P(y_0)$ and $P(y_1 | x) = P(y_1)$ for all values y_0, y_1 and x . Hence, as an example for β , the uplift bounds simplify to

$$\mathbb{E}[\max\{0, P(Y_0 = 1) - P(Y_1 = 1)\}] \leq \beta \leq \mathbb{E}[\min\{P(Y_0 = 1), P(Y_1 = 1)\}].$$

The expected value is on the distribution of X , but since the terms in the expected value do not depend on X , the bounds reduce to

$$\max\{0, S_0 - S_1\} \leq \beta \leq \min\{S_0, S_1\}$$

Which are the Fréchet bounds on β . The same reasoning applies to the bounds on α, γ and δ . \square

4.1 Probability bounds and uplift estimation

The main motivation underlying the derivation of the uplift bounds is that in real-world settings characterized by large historical data sets (like churn modeling), it is possible to derive sample-based estimates of the terms bounding the counterfactual probabilities. In particular, we advocate the adoption of a plug-in estimator from an uplift model $\hat{S}_0(x), \hat{S}_1(x)$

on a data set $\mathcal{D} = \{x^{(1)}, \dots, x^{(N)}\}$. In this case, a sample-based version of the lower bound on β is

$$\mathbb{E}[\max\{0, S_0(X) - S_1(X)\}] \approx \frac{1}{N} \sum_{i=1}^N \max\{0, \widehat{S}_0(x^{(i)}) - \widehat{S}_1(x^{(i)})\} \tag{28}$$

And similarly for the other bounds on α, \dots, δ .

It is sometimes desirable to obtain a point estimate on the probability of counterfactuals, for example when a unique number is expected as the result of the counterfactual analysis. Though a naive estimator could be derived by taking the midpoint of the bounds, in the next section we will introduce a more theoretically founded estimator.

5 Point estimate of counterfactual probabilities

Counterfactual probabilities are latent yet very important quantities to be taken into consideration for decision-making. In the previous section, we proposed an original approach to bound their values. However, it is sometimes desirable to compute a point estimate of those probabilities, even if this requires stronger assumptions. Here we present a point estimator of the probabilities α, \dots, δ (Equations (4) to (7)) based on the conditional independence between Y_0 and Y_1 . The introduction of specific assumptions is required since those probabilities, e.g.

$$\alpha = P(Y_0 = 0, Y_1 = 0) = \mathbb{E}_X[P(Y_0 = 0, Y_1 = 0 | X)] = \mathbb{E}_X[\alpha(X)] \tag{29}$$

Cannot be estimated from observational or experimental data, given that one of the two outcomes will be necessarily unobserved. The conditional independence between Y_0 and Y_1 given $X = x$, which is formally expressed as $Y_0 \perp Y_1 | X = x$, allows developing the term $\alpha(x)$ in Equation (29) as

$$\alpha(x) \approx P(Y_0 = 0 | X = x)P(Y_1 = 0 | X = x). \tag{30}$$

In order to study the impact of the conditional independence assumption, we define the difference between $P(Y_0 = 0, Y_1 = 0 | X = x)$ and the approximation $P(Y_0 = 0 | X = x)P(Y_1 = 0 | X = x)$ as $\phi(x)$. This quantity appears in the other conditional probabilities too:

$$\alpha(x) = P(Y_0 = 0 | X = x)P(Y_1 = 0 | X = x) + \phi(x) \tag{31}$$

$$\beta(x) = P(Y_0 = 1 | X = x)P(Y_1 = 0 | X = x) - \phi(x) \tag{32}$$

$$\gamma(x) = P(Y_0 = 0 | X = x)P(Y_1 = 1 | X = x) - \phi(x) \tag{33}$$

$$\delta(x) = P(Y_0 = 1 | X = x)P(Y_1 = 1 | X = x) + \phi(x). \tag{34}$$

Note that the quantity $\phi(x)$ can be interpreted as a conditional measure of dependency between Y_0 and Y_1 and is similar to classical binary dependency measures, like the odd ratio, Yule’s Q coefficient, or the difference coefficient (Edwards, 1957). We will see in Theorem 2 that

$$\phi = \alpha\delta - \beta\gamma - \text{cov}_X(S_0(X), S_1(X)). \quad (35)$$

Meaning that ϕ depends both on the distribution of counterfactuals (α, β, γ and δ) and the dependency between the scores $S_0(x)$ and $S_1(x)$. From (29) we obtain

$$\alpha = \mathbb{E}[\alpha(X)] \quad (36)$$

$$= \mathbb{E}[P(Y_0 = 0 | X)P(Y_1 = 0 | X) + \phi(X)] \quad (37)$$

$$= \mathbb{E}[(1 - S_0(X))(1 - S_1(X))] + \phi \quad (38)$$

Where $\phi = \mathbb{E}[\phi(X)]$. If we assume $Y_0 \perp Y_1 | X$, then $\phi = 0$ and

$$\alpha \approx \mathbb{E}[(1 - S_0(X))(1 - S_1(X))]. \quad (39)$$

The question of the dependency between Y_0 and Y_1 has already been discussed in the causal inference literature (Imbens & Rubin, 2015, Sec.8.6). A possible approach could be to assume the maximum possible dependency between the potential outcomes. Alternatively, one could make no a priori preference between a positive and negative association between Y_0 and Y_1 (i.e. Y_0 and Y_1 taking similar or opposite values), thus assuming no association. Since, in absence of some preexisting knowledge, there is no a priori good answer, it is more interesting to reason about the dependency between Y_0 and Y_1 as follows:

- A positive correlation⁴ between Y_0 and Y_1 means that they are often equal, indicating that the treatment has little effect on the outcome. When the correlation is maximum, the upper bounds on α and δ in Equations (20) and (23) are met.
- A negative correlation between Y_0 and Y_1 indicates that the treatment has either a strongly positive or negative impact on the outcome. When the correlation is maximally negative, the upper bounds on β and γ in Equations (21) and (22) are met.
- The absence of dependency indicates an even mix of the two previous cases. This corresponds to the point estimator presented in this section.

5.1 Point estimate and uplift estimation

Given estimators $\hat{S}_0(x), \hat{S}_1(x)$ of the uplift terms, and an evaluation data set $\{x^{(i)}\}_{i=1, \dots, N}$, we propose to estimate α, \dots, δ as

$$\hat{\alpha} = \frac{1}{N} \sum_i (1 - \hat{S}_0(x^{(i)}))(1 - \hat{S}_1(x^{(i)})) \quad (40)$$

$$\hat{\beta} = \frac{1}{N} \sum_i \hat{S}_0(x^{(i)})(1 - \hat{S}_1(x^{(i)})) \quad (41)$$

$$\hat{\gamma} = \frac{1}{N} \sum_i (1 - \hat{S}_0(x^{(i)}))\hat{S}_1(x^{(i)}) \quad (42)$$

⁴ The correlation between Y_0 and Y_1 refers to the tendency of Y_0 and Y_1 to take identical or opposite values.

$$\hat{\delta} = \frac{1}{N} \sum_i \hat{S}_0(x^{(i)})\hat{S}_1(x^{(i)}). \tag{43}$$

The bias of these estimators is expressed in Theorem 2.

Theorem 2 *Given that $\hat{S}_0(x)$ and $\hat{S}_1(x)$ are unconfounded and unbiased estimators of $S_0(x)$ and $S_1(x)$ trained on a training set with distribution D , in the large sample limit the bias of $\hat{\alpha}, \dots, \hat{\delta}$ is*

$$\begin{aligned} \text{Bias}[\hat{\beta}] &= \text{Bias}[\hat{\gamma}] = -\text{Bias}[\hat{\alpha}] = -\text{Bias}[\hat{\delta}] \\ &= \alpha\delta - \beta\gamma - \text{cov}_X(S_0(X), S_1(X)) - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))] \end{aligned} \tag{44}$$

$$= \phi - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))]. \tag{45}$$

Proof We will derive the bias of $\hat{\beta}$, and the bias of the three other estimators can be derived in a similar way. The expected value of $\hat{\beta}$ over the distribution of training sets D is

$$\begin{aligned} \mathbb{E}_D[\hat{\beta}] &= \mathbb{E}_D \left[\frac{1}{N} \sum_{i=1}^N \hat{S}_0(x^{(i)})(1 - \hat{S}_1(x^{(i)})) \right] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}_D \left[\hat{S}_0(x^{(i)})(1 - \hat{S}_1(x^{(i)})) \right] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}_D[\hat{S}_0(x^{(i)})]\mathbb{E}_D[1 - \hat{S}_1(x^{(i)})] + \text{cov}_D(\hat{S}_0(x^{(i)}), 1 - \hat{S}_1(x^{(i)})) \\ &= \frac{1}{N} \sum_{i=1}^N S_0(x^{(i)})(1 - S_1(x^{(i)})) - \text{cov}_D(\hat{S}_0(x^{(i)}), \hat{S}_1(x^{(i)})). \end{aligned}$$

In the large sample limit ($N \rightarrow +\infty$), we can assume that this sum converges to

$$\mathbb{E}_D[\hat{\beta}] = \mathbb{E}_X[S_0(X)(1 - S_1(X))] - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))].$$

The first term can be expanded as

$$\begin{aligned} \mathbb{E}[S_0(X)(1 - S_1(X))] &= \mathbb{E}[S_0(X)]\mathbb{E}[1 - S_1(X)] + \text{cov}_X(S_0(X), 1 - S_1(X)) \\ &= S_0(1 - S_1) - \text{cov}_X(S_0(X), S_1(X)) \\ &= (\beta + \delta)(\beta + \alpha) - \text{cov}_X(S_0(X), S_1(X)) \\ &= \beta(\beta + \delta + \alpha) + \alpha\delta - \text{cov}_X(S_0(X), S_1(X)) \\ &= \beta(1 - \gamma) + \alpha\delta - \text{cov}_X(S_0(X), S_1(X)) \\ &= \alpha\delta - \beta\gamma + \beta - \text{cov}_X(S_0(X), S_1(X)). \end{aligned}$$

And thus

$$\mathbb{E}_D[\hat{\beta}] = \alpha\delta - \beta\gamma + \beta - \text{cov}_X(S_0(X), S_1(X)) - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))].$$

Finally, the bias of $\hat{\beta}$ is

$$\begin{aligned} \text{Bias}[\hat{\beta}] &= \mathbb{E}_D[\hat{\beta}] - \beta \\ &= \alpha\delta - \beta\gamma - \text{cov}_X(S_0(X), S_1(X)) - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))] \end{aligned}$$

Which proves Equation (44). Equation (45) is derived from

$$\mathbb{E}[S_0(X)(1 - S_1(X))] = \mathbb{E}[\beta(x) + \phi(x)] = \beta + \phi$$

And then

$$\begin{aligned} \text{Bias}[\hat{\beta}] &= \mathbb{E}_D[\hat{\beta}] - \beta \\ &= \mathbb{E}_X[S_0(X)(1 - S_1(X))] - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))] - \beta \\ &= \phi - \mathbb{E}_X[\text{cov}_D(\hat{S}_0(X), \hat{S}_1(X))]. \end{aligned}$$

□

While the three first terms in Equation (44) are inherent to the customer population, the last term depends also on the estimators $\hat{S}_0(x)$ and $\hat{S}_1(x)$, and the data distribution D . Without assumptions about these processes, the last term cannot be further reduced.

The proposed procedure to compute $\hat{\beta}$ as well as the two uplift bounds on β presented in Sect. 4 is described in Algorithm 1, where we assume we have two unbiased estimators of the scores $S_0(x)$ and $S_1(x)$.

Algorithm 1 Estimating the counterfactual probability β

Input: Data set $\mathcal{D} = \{(x^{(i)}, y^{(i)}, t^{(i)})\}_{i=1, \dots, N}$

Output: Point estimate $\hat{\beta}$, and bounds \widehat{LB}_β and \widehat{UB}_β such that $\widehat{LB}_\beta \leq \beta \leq \widehat{UB}_\beta$

Split \mathcal{D} into training set \mathcal{D}_{tr} and test set \mathcal{D}_{te}

Train uplift model on \mathcal{D}_{tr} to obtain estimators $(\hat{S}_0(x), \hat{S}_1(x))$

Compute $\hat{\beta} = \frac{1}{|\mathcal{D}_{te}|} \sum_i \hat{S}_0(x^{(i)})(1 - \hat{S}_1(x^{(i)}))$ on \mathcal{D}_{te}

Compute $\widehat{LB}_\beta = \frac{1}{|\mathcal{D}_{te}|} \sum_i \max\{0, \hat{S}_0(x^{(i)}) - \hat{S}_1(x^{(i)})\}$ on \mathcal{D}_{te}

Compute $\widehat{UB}_\beta = \frac{1}{|\mathcal{D}_{te}|} \sum_i \min\{\hat{S}_0(x^{(i)}), 1 - \hat{S}_1(x^{(i)})\}$ on \mathcal{D}_{te}

6 Bounds assessment by simulation

In this section, we assess the bounds and estimators presented in Sects. 4 and 5 by setting up a specific simulation environment. The simulated nature of the experiment allows us to compare the estimated bounds to the ground truth.

6.1 Methodology

Let α, β, γ and δ be the terms introduced in (4), (5), (6), (7). The aim of the simulation is to generate samples from a distribution where the scores $S_0(x)$ and $S_1(x)$ are conditional on a set of features X .

One possible approach is to model the X covariate distribution, the stochastic functional dependency between Y, X and T , train an uplift model $\widehat{S}_0(x), \widehat{S}_1(x)$ on a generated data set $\mathcal{D} = \{(x^{(i)}, y^{(i)}, t^{(i)})_{i=1, \dots, N}\}$, and finally apply the estimators presented in the previous sections. We did not consider this approach since the results would heavily depend on the model choices (e.g. the distribution of X and the class of functions for Y) and the learning algorithm.

Our simulation setting consists in directly sampling the distributions of the estimators \widehat{S}_0 and \widehat{S}_1 obtained as a noisy version of S_0 and S_1 which, according to (8) and (9), are functions of the terms α, \dots, δ . Since we sample the distribution of scores \widehat{S}_0 and \widehat{S}_1 but we do not sample X directly, we will denote individual scores with superscript i rather than as functions of x . The sampling process of our simulation is detailed in Equations (46) to (49):

$$(\alpha^{(i)}, \dots, \delta^{(i)}) \sim \text{Dir}(a, b, c, d) \tag{46}$$

$$S_0^{(i)} = \beta^{(i)} + \delta^{(i)} \quad S_1^{(i)} = \gamma^{(i)} + \delta^{(i)} \tag{47}$$

$$\widehat{S}_0^{(i)} \sim \frac{1}{v} \text{B}(v, S_0^{(i)}) \quad \widehat{S}_1^{(i)} \sim \frac{1}{v} \text{B}(v, S_1^{(i)}) \tag{48}$$

$$(Y_0^{(i)}, Y_1^{(i)}) \sim \text{Cat}(\alpha^{(i)}, \dots, \delta^{(i)}) \tag{49}$$

First, we generate N independent samples $(\alpha^{(i)}, \beta^{(i)}, \gamma^{(i)}, \delta^{(i)})_{i=1, \dots, N}$ according to a Dirichlet distribution $\text{Dir}(a, b, c, d)$. They represent the probabilities of counterfactuals at the individual level. The Dirichlet distribution is a natural candidate to sample numbers in a probability simplex (i.e. such that $\alpha^{(i)}, \beta^{(i)}, \gamma^{(i)}$ and $\delta^{(i)}$ are all positive and sum up to 1), since it is the conjugate prior of the multinomial distribution (Lin, 2016). Then, we derive the value of the scores $S_0^{(i)}$ and $S_1^{(i)}$ with the identities $S_0^{(i)} = \beta^{(i)} + \delta^{(i)}$ (Equation (8)) and $S_1^{(i)} = \gamma^{(i)} + \delta^{(i)}$ (Equation (9)). To emulate imperfect estimators $\widehat{S}_0^{(i)}$ and $\widehat{S}_1^{(i)}$, we draw $\widehat{S}_t^{(i)}$ (for $t = 0, 1$) according to a normalized binomial distribution $\frac{1}{v} \text{B}(v, S_t^{(i)})$, where v is the parameter controlling the variance of $\widehat{S}_t^{(i)}$. Such estimator distribution guarantees that $\widehat{S}_t^{(i)}$ takes values inside $[0, 1]$ and models the variability of $\widehat{S}_t^{(i)}$ due to a limited number of training examples of a binary outcome Y_t . Finally, the counterfactual outcomes $Y_0^{(i)}$ and $Y_1^{(i)}$ are sampled according to a categorical distribution $\text{Cat}(\alpha^{(i)}, \dots, \delta^{(i)})$ such that $P(Y_0^{(i)} = 0, Y_1^{(i)} = 0) = \alpha^{(i)}$, and similarly for $\beta^{(i)}, \gamma^{(i)}$ and $\delta^{(i)}$, reflecting Equations (4) to (7). Once the sampling process is executed, the bounds and estimators from Sects. 4 and 5 can be evaluated from the set of scores $\{(\widehat{S}_0^{(i)}, \widehat{S}_1^{(i)})\}_{i=1, \dots, N}$.

6.2 Simulation parameters

The simulation setting is defined by six main parameters: N, v, a, b, c and d .

- The parameter N represents the size of the data set on which the bounds and estimators are evaluated.
- The parameter ν emulates the variance of the simulated uplift model. Higher values of ν induce a lower variance since we can show⁵ that $\text{Var}(\widehat{S}_t^{(i)}) = S_t^{(i)}(1 - S_t^{(i)})/\nu$.
- The parameters a, b, c and d are proportional to the distribution of counterfactuals $P(Y_0^{(i)} = 0, Y_1^{(i)} = 1), \dots, P(Y_0^{(i)} = 1, Y_1^{(i)} = 1)$. For example, using the moments of the Dirichlet distribution, we have

$$P(Y_0^{(i)} = 1, Y_1^{(i)} = 0) = \mathbb{E}[\beta^{(i)}] = \frac{b}{A} \tag{50}$$

Where $A = a + b + c + d$.

- The value of A influences the distribution of $\alpha^{(i)}, \dots, \delta^{(i)}$. High values of A lead to samples $\alpha^{(i)}, \dots, \delta^{(i)}$ concentrated around their expected values (which can be computed from Equation (50)), while low values of A lead to samples where one of $\alpha^{(i)}, \dots, \delta^{(i)}$ is close to one while the three other values are close to zero. This has an impact on the scores $S_0^{(i)}, S_1^{(i)}$ as well: they are close to their expected values when A is large, and close to either zero or one when A is low. In loose terms, the quantity A represents the amount of information that the covariates X brings about the outcomes Y_0 and Y_1 : when the features are uninformative, the scores $S_0(x), S_1(x)$ are close to their prior probabilities $P(Y_0 = 1)$ and $P(Y_1 = 1)$, while when the features are highly informative, the scores are close to either zero or one.

Theorem 2 indicates that the bias of the points estimators $\widehat{\alpha}, \dots, \widehat{\delta}$ (transposed to the notation of this section) is

$$\mathbb{E}[\phi^{(i)}] - \mathbb{E}_{\alpha^{(i)}, \dots, \delta^{(i)}}[\text{cov}(\widehat{S}_0^{(i)}, \widehat{S}_1^{(i)})].$$

The second term is null because we sample $\widehat{S}_0^{(i)}$ and $\widehat{S}_1^{(i)}$ independently, but we can show using the product moments of the Dirichlet distribution (Lin, 2016) that the first term $\mathbb{E}[\phi^{(i)}]$ is

$$\mathbb{E}[\phi^{(i)}] = \mathbb{E}[\alpha^{(i)}\delta^{(i)} - \beta^{(i)}\gamma^{(i)}] = \frac{ad - bc}{A(A + 1)}. \tag{51}$$

Since the parameters a, b, c, d are sampled uniformly, the expression in Equation (51) will be different from zero. Therefore, the distribution of the bias of the point estimators $\widehat{\alpha}, \dots, \widehat{\delta}$ has a large variance. This is desirable to assess how violations of the hypothesis underlying our estimators affect the quality of the estimation.

6.3 Assessment of the theoretical results

In this section, we assess the quality of the uplift bounds and the point estimator, discussed in Sects. 4 and 5 respectively, for different values of the simulation parameters. The simulation process is repeated 5000 times with randomly chosen parameters. The size N of the evaluation set varies between 10 and 10000 and the variance

⁵ The variable $\widehat{S}_t^{(i)}$ is based on a Binomial distribution $B(\nu, S_t^{(i)})$, which has a variance $\nu S_t^{(i)}(1 - S_t^{(i)})$. We can develop the variance $\text{Var}(\widehat{S}_t^{(i)}) = \nu S_t^{(i)}(1 - S_t^{(i)})/\nu^2 = S_t^{(i)}(1 - S_t^{(i)})/\nu$.

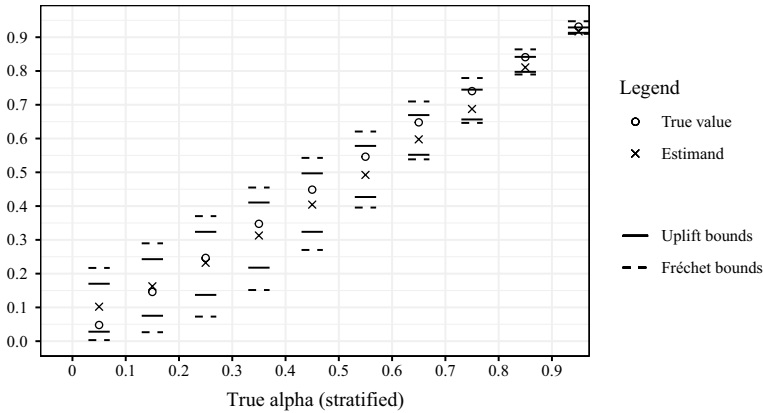


Fig. 1 The estimator $\hat{\alpha}$, the true value of α , and the bounds on α , for different values of α . We take the average over all experiments where α falls into the relevant range. The graph is quite similar for β, γ and δ

Table 3 Identification error for β . We compare the Fréchet bounds and the uplift bounds, and we also compare the point estimators with the bounds mid-point. We observe that the uplift bounds provide a clear improvement over the Fréchet bounds

Bounds	Mean width
Uplift bounds	12.6%
Fréchet bounds	24.9%
Estimator	RMSE
Point estimator $\hat{\beta}$	6.4%
Uplift bounds mid-point	5.9%
Fréchet bounds mid-point	8.0%

parameter ν varies between 5 and 50. The Dirichlet parameters a, b, c, d are fixed as $(a, b, c, d) = A(\alpha, \beta, \gamma, \delta)$ where A varies between 0.1 and 15, and the vector $(\alpha, \beta, \gamma, \delta)$ is sampled uniformly over the probability simplex (i.e. such that the four terms are positive and sum up to one).

Figure 1 plots the estimator $\hat{\alpha}$ (cross), the uplift bounds (continuous line) and the Fréchet bounds (dashed lines) with respect to the true α (circle). Since the plots for β, γ and δ are quite similar, they are omitted for the sake of conciseness. The values for the 5000 simulation runs are stratified according to the true value α in order to simplify the plot. For each stratum, the point reports the average of the estimated value and the horizontal bars report the average upper and lower bounds. The main conclusions of the simulation study are:

- The uplift bounds are significantly tighter than the Fréchet bounds, as shown in Fig. 1. The bounds span is typically reduced by half, as reported in Table 3.
- The point estimate provides a good approximation of the true counterfactual probability, with a root mean squared error (RMSE) of 6.4% (Table 3). In order to have a baseline for comparison, we also compute the RMSE of the bounds mid-point if those were taken as point estimators of the true counterfactual probability. We see that the Fréchet bounds mid-point has a larger error while the uplift bounds mid-point has an error comparable to the point estimate.




Fig. 2 Distribution of the point estimator bias, $\mathbb{E}[\phi^{(i)}]$, over 4000 simulation runs. Note that this is different from the distribution of $\phi^{(i)}$ in a given simulation run. Although the maximum is around zero, it is never exactly zero, indicating that the estimators are biased in our simulations. This is desirable to reflect violations of the hypotheses underlying our estimators in practical scenarios




Fig. 3 The bounds span as a function of the conditional entropy of $Y_0^{(i)}, Y_1^{(i)}$, which is directly influenced by the parameter A . We fixed $(\alpha, \beta, \gamma, \delta) = (0.947, 0.020, 0.017, 0.017)$, and $v = 50$ and $N = 2000$

- The distribution of the bias $\mathbb{E}[\phi(x)]$ of the point estimator, which is defined by Equation (51), is shown in Fig. 2. The fact that most of the bias realizations are different from zero is an indication of the realism of the simulation setting and a positive sign about the robustness of the theoretical results.

6.4 Sensitivity analysis of the simulation

In this section, we assess the influence of the training data (in terms of the number of samples or the information of the features) on the precision of the estimation. We plot the span of the uplift bounds and the error of the point estimator while varying one of the parameters A, N and v and keeping the other parameters fixed. The values of the fixed parameters are selected to clearly show the influence of the varying parameters. In particular, we set $(\alpha, \beta, \gamma, \delta) = (0.947, 0.020, 0.017, 0.017)$ based on the results of Sect. 7,

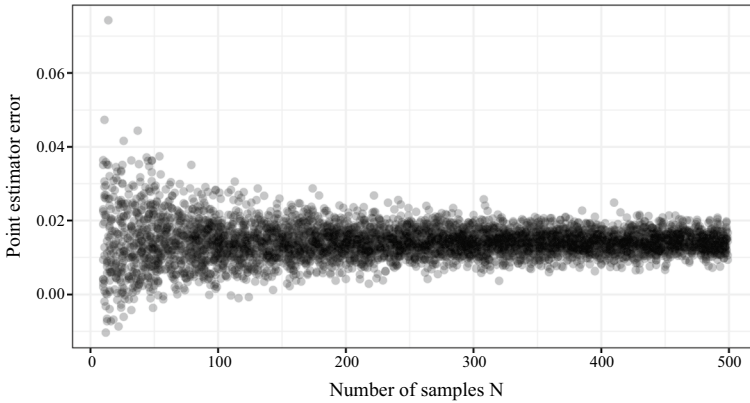


Fig. 4 The error of the point estimator as a function of the number of samples in the evaluation data set. We fixed $(\alpha, \beta, \gamma, \delta) = (0.947, 0.020, 0.017, 0.017)$, and $v = 20$ and $A = 1$




Fig. 5 The error of the point estimator as a function of model variance $\text{Var}(\hat{S}_t^{(i)})$. We fixed $(\alpha, \beta, \gamma, \delta) = (0.947, 0.020, 0.017, 0.017)$, and $N = 1000$ and $A = 10$. As the variance decreases, the estimator bias converges towards to its theoretical value

which represents the distribution of counterfactuals in a typical scenario of customer churn prevention in telecom. The main conclusions of this sensitivity analysis are:

- The uplift bounds span decreases as the conditional entropy of Y_0, Y_1 decreases (Fig. 3). This is an empirical illustration of Theorem 1. In the context of this simulation, since we do not model the features X , we instead note the conditional entropy as $H(Y_0^{(i)}, Y_1^{(i)} | \alpha^{(i)}, \dots, \delta^{(i)})$, and we compute its value as

$$\begin{aligned}
 & H(Y_0^{(i)}, Y_1^{(i)} | \alpha^{(i)}, \dots, \delta^{(i)}) \\
 &= \frac{1}{N} \sum_{i=1}^N (-\alpha^{(i)} \log \alpha^{(i)} - \beta^{(i)} \log \beta^{(i)} - \gamma^{(i)} \log \gamma^{(i)} - \delta^{(i)} \log \delta^{(i)}).
 \end{aligned}$$




Fig. 6 The uplift bounds span as a function of the model variance $\text{Var}(\widehat{S}_t^{(i)})$. We fixed $(\alpha, \beta, \gamma, \delta) = (0.947, 0.020, 0.017, 0.017)$, and $N = 1000$ and $A = 10$. A lower model variance is shown here to be associated with larger bounds. In fact, as the variance goes to zero (left side of the plot), the bounds span converges towards its theoretical value. A model with a high variance predicts more often low values, which artificially reduces the bounds span

We see that as the conditional entropy approaches zero (which emulates very informative features), the bounds span converges towards zero as well.

- The variance of the point estimator decreases as the number of samples N increases (Fig. 4) or the model variance $\text{Var}(\widehat{S}_t^{(i)})$ decreases (Fig. 5). In fact, the error converges towards the bias derived in Theorem 2. This demonstrates the convergence of our estimator in the large sample scenario.
- The uplift bounds span increases with the decrease of the model variance $\text{Var}(\widehat{S}_t^{(i)})$ (for $t = 0, 1$) (Fig. 6). This is because a model with a high variance predicts often lower or higher scores than the expected score. Since the bounds span is $\mathbb{E}[\min\{S_0^{(i)}, S_1^{(i)}, 1 - S_0^{(i)}, 1 - S_1^{(i)}\}]$ (see Equation (27)), this artificially reduces the bounds span.

7 Evaluation with real data

This section applies the theoretical results discussed so far to a real-world data set provided by our industrial partner Orange Belgium that includes 6 churn prevention campaigns.

7.1 Data set description

Churn prevention campaigns are used to mitigate customer churn by contacting customers at risk of leaving the company. They are offered an incentive to stay, such as a promotional offer or a suggestion for a better tariff plan. The retention campaigns were performed over 6 months in 2019 and 2020. Before each campaign, a churn prediction model (independent of the models evaluated in this section) was trained on the whole customer base to predict the churn risk. The riskiest customers were randomly split into target and

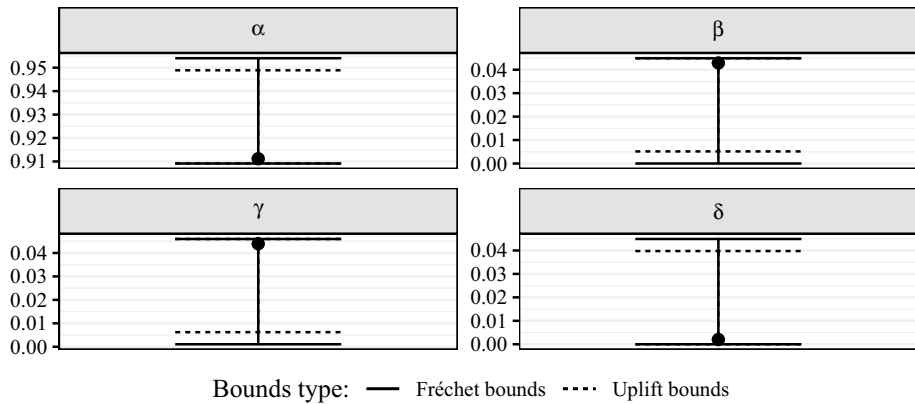


Fig. 7 Point estimate and bounds on α, \dots, δ . Note the different vertical axis for α


control groups. Customers in the target group were contacted by phone and were proposed a tariff plan adapted to the apparent root cause of potential churn. For example, if a large amount of mobile data was used, a tariff plan with a larger provision of mobile data was then suggested. The final data set used in this section comprises only customers selected in the target and control groups, all other customers that are not part of the campaign are discarded. The data set contains 11268 samples, for 145 features. Examples of features include the tariff plan of the customer, the number of calls over the last month, some socio-demographic information, the number of calls to customer service, and so on. The churn rate in the control group is 4.85%, while in the target group it is 4.03%. The control group amounts to 33% of the data set. Note that the treatment indicator in this data set indicates whether a call attempt to the customer was made, and does not indicate whether the customer answered the call or accepted the offer.

7.2 Methodology

We train an uplift random forest model (Guelman et al., 2015) on the O. data set using the R package *uplift* (Guelman, 2014). Other uplift models have been shown to be superior in accuracy (e.g. the X-learner (Künzel et al., 2019)) but we need here separate estimators for $S_0(x)$ and $S_1(x)$ to compute the uplift bounds and the point estimator. This condition is satisfied by the uplift random forest, as well as the T-learner approach (Künzel et al., 2019). The uplift random forest model is trained with 100 trees. Given the high imbalance of the data sets, we rely on the EasyEnsemble strategy (Liu et al., 2009) for class balancing. It consists in training k base learners ($k = 10$ in our case) on the whole set of positive instances (churners) and an equally sized random set of negative instances. This choice is based on previous literature on similar tasks with high imbalance and large class overlap (Zhu et al., 2017; Dal Pozzolo et al., 2014). The predictions of all the base learners are averaged to obtain the final prediction. When a resampling strategy such as EasyEnsemble is used to obtain a balanced data set, the prior probability of churn is modified (Batista et al., 2004), and the scores predicted by the trained model are biased. This bias is corrected with the calibration formula presented by Dal Pozzolo et al. (2015). To avoid overfitting on a specific train-test split, we repeat the experiment using a k -fold cross-validation scheme with $k = 5$.

Table 4 Numerical values of the estimated counterfactual distribution α, \dots, δ on the O. data set. The uplift bounds and the Fréchet bounds show similar results

	α	β	γ	δ
Point estimate	91.12	4.29	4.39	0.20
Uplift bounds	[90.91, 94.89]	[0.52, 4.49]	[0.62, 4.60]	[0.00, 3.98]
Fréchet bounds	[90.91, 95.40]	[0.00, 4.49]	[0.11, 4.60]	[0.00, 4.49]

Fig. 8 Point estimate and uplift bounds on β , for each month of the campaign

7.3 Results

The estimated distribution of counterfactuals is reported in Fig. 7 and Table 4. In Fig. 7, each of $\hat{\alpha}$, $\hat{\beta}$, $\hat{\gamma}$ and $\hat{\delta}$ is reported in a different sub-plot, together with the uplift and Fréchet bounds. We observe that the uplift bounds are consistently tighter than the Fréchet bounds, although not by a large margin. The value of $\hat{\beta}$ and $\hat{\gamma}$ are very close, with point estimates at respectively 4.29% and 4.39%. The value of $\hat{\alpha}$ is high, around 91.12%, as expected since most customers do not churn.

The proportion of persuadable customers is estimated as $\hat{\beta} = 4.29\%$, with a lower bound of 0.52% and an upper bound of 4.49%. This amounts to 483 customers, bounded between 58 and 505. This indicates that a maximum of 505 customers should have been called during the 6-months campaign, while in practice 7500 customers have been called. We applied the same methodology separately for each month instead of on the whole campaign data, and the results are reported in Fig. 8. We observe that, although the value of $\hat{\beta}$ seems to fluctuate from one month to the next, it tends to be close to the upper bound. This is because both $\hat{S}_0(x)$ and $\hat{S}_1(x)$ tend to be close to zero, and $\hat{\beta}(x)$ is estimated as $\hat{S}_0(x)(1 - \hat{S}_1(x))$ in Equation (41). Therefore $\hat{\beta}(x)$ is typically close to $\hat{S}_0(x)$, and the upper bound $\min\{\hat{S}_0(x), 1 - \hat{S}_1(x)\}$ from Equation (21) is almost always equal to $\hat{S}_0(x)$ as well.

7.4 Profit analysis

To give some intuition about these results, we now conduct a simplistic profit analysis. Let us suppose that each call has a cost $C = 1\text{€}$, and that the average customer lifetime value is $V = 120\text{€}$ (a customer pays on average 20€ per month and stays 6 months). The benefit due to the campaign as it actually happened can be computed as

$$\text{Profit} = NUV - NC \quad (52)$$

Where N is the number of contacted customers and $U = S_0 - S_1$ is the campaign uplift (approximately 0.8% in our case). The term NUV is the benefit generated by converting customers. The benefit of calling do-not-disturb customers cancels out the benefit of calling persuadable customers, since $U = \beta - \gamma$.⁶ The term NC in Equation (52) is the cost of calling N customers. By evaluating this expression on the O. data set, we obtain that the campaign incurred a net loss of 130€. However, if we suppose that we were able to call only the 483 persuadable customers, the campaign could generate a profit of up to 57477€. Note that this is a simplistic way to evaluate the profit generated by a campaign. For more detailed estimations of the profit, we refer the reader to (Li & Pearl, 2019; Verbraken et al., 2013; Verbeke et al., 2012; Gubela & Lessmann, 2021).

7.5 Discussion

The improvement of the uplift bounds with respect to the Fréchet bounds is directly related to the quantity of information between the features and the outcome (see Theorem 1). The small improvement observed in practice, as shown in Fig. 7, indicates that the uplift terms, and in turn counterfactual probabilities, are difficult to estimate in real-world settings such as customer churn prediction. A possible solution would be to add more informative features or design a more powerful uplift model. The bounds can also be further refined when observational data is available (i.e. data where the treatment assignment is not randomized), as demonstrated in (Mueller & Pearl, 2022). The results of this section provide nonetheless very valuable insights for our industrial partner Orange Belgium on the potential value of past retention campaigns and on the distribution of the different customer categories.

The results of this section do not indicate which customers should be targeted in order to maximize the profit from the retention campaign. This is the objective of uplift modeling. There is some debate on whether uplift modeling is always the best approach for causal decision-making. Fernández-Loria and Provost (2022a, 2022b) show that uplift models are sub-optimal under some circumstances, and that proxy targets such as the probability of the outcome are sometimes more effective for accurate causal decision-making. This is in line with the abundant literature on churn management that use predictive models instead of uplift models, e.g. (Amin et al., 2019; Coussemont et al., 2017; Óskarsdóttir et al., 2018) to cite a few. Li and Pearl (2019) consider the case where each of the four categories of customers (*persuadable*, *sure thing*, *lost cause* and *do-not-disturb*, see Table 2) have arbitrary associated costs. In this case, counterfactual identification is essential for accurate causal decision-making.

8 Conclusion

We have derived and empirically assessed new bounds and a point estimator on the probability of counterfactuals for binary outcomes under the assumption of unconfoundedness. Counterfactuals are essential for accurate decision-making for example in churn prevention in the telecom industry.

The proposed uplift bounds improve upon the classical Fréchet bounds by leveraging the scores estimated by an uplift model. We have demonstrated theoretically that

⁶ This can be shown by decomposing $U = P(Y_0 = 1) - P(Y_1 = 1)$ in terms of β , γ and δ .

the bounds improve as the quality of the uplift estimation increases. Simulated examples indicate that the uplift bounds typically provide a significant improvement over the Fréchet bounds. We have also derived a point estimator by assuming the conditional independence between the potential outcomes Y_0 and Y_1 . Simulated examples demonstrate that the estimator is still close to the true value even when this condition is not respected.

Our estimators are limited by several factors. The most important is the choice of the underlying uplift model. The uplift model should be unbiased, and the quality of the estimator depends on the quality of the uplift model. Since the two uplift terms $S_0(x)$ and $S_1(x)$ are used independently in our estimators, we are also limited to uplift estimators that can provide an estimation of these two terms separately.

Counterfactuals model individual behavior and as such can provide significant business insights about customers. In future work, we intend to explore the relationship between counterfactuals and customer features. This will allow describing the persuadable customers in terms of concrete characteristics, which very desirable from a business standpoint.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10994-023-06317-w>.

Author's contribution Conceptualization and methodology: all authors. Formal analysis, investigation and writing: TV. Funding acquisition: all authors. Resources and supervision at Orange Belgium: DM and JS. Academic supervision: GB.

Funding This work was supported by the Applied PhD Machu-Pichu project funded by Innoviris (2019-PHD-16). We thank this agency for allowing us to conduct both fundamental and applied research.

Availability of code, data and material An anonymized version of the data set and the software written for this article are available at <https://github.com/TheoVerhelst/Counterfactual-uplift-bounds>.

Declarations

Conflict of interest The authors have no competing interests to declare that are relevant to the content of this article.

Ethics approval, consent to participate, consent for publication Not applicable.

References

- Amin, A., Al-Obeidat, F., Shah, B., Adnan, A., Loo, J., & Anwar, S. (2019). Customer churn prediction in telecommunication industry using data certainty. *Journal of Business Research*, 94, 290–301. (Publisher: Elsevier) <https://doi.org/10.1016/j.jbusres.2018.03.003>
- Athey, S., & Imbens, G. (2016, July). Recursive partitioning for heterogeneous causal effects. In: *Proceedings of the National Academy of Sciences*, 113(27), 7353–7360. Retrieved 2020-05-06, from <https://www.pnas.org/content/113/27/7353> (Publisher: National Academy of Sciences) <https://doi.org/10.1073/PNAS.1510489113>
- Balke, A., & Pearl, J. (1994). Counterfactual probabilities: Computational methods, bounds and applications. *Uncertainty Proceedings 1994* (pp. 46-54). Elsevier.
- Batista, G.E., Prati, R.C., & Monard, M.C. (2004). A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 6(1), 20–29. (Publisher: ACM) <https://doi.org/10.1145/1007730.1007735>
- Correa, J., Lee, S., & Bareinboim, E. (2021). Nested counterfactual identification from arbitrary surrogate experiments. *Advances in Neural Information Processing Systems*, 34, 6856–6867.

- Coussement, K., Lessmann, S., & Verstraeten, G. (2017). A comparative analysis of data preparation algorithms for customer churn prediction: a case study in the telecommunication industry. *Decision Support Systems*, 95, 27–36. (Publisher: Elsevier) <https://doi.org/10.1016/j.dss.2016.11.007>
- Curth, A., & van der Schaar, M. (2021). Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In: *International Conference on Artificial Intelligence and Statistics* (pp. 1810–1818). (arXiv: Curth2021)
- Dal Pozzolo, A., Caelen, O., Johnson, R.A., & Bontempi, G. (2015). Calibrating probability with under-sampling for unbalanced classification. In: *2015 IEEE Symposium Series on Computational Intelligence* (pp. 159–166). IEEE.
- Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). *Learned lessons in credit card fraud detection from a practitioner perspective*. *Expert Systems with Applications*, 41(10), 4915–4928. Publisher: Elsevier.
- Devriendt, F., Berrevoets, J., & Verbeke, W. (2019). *Why you should stop predicting customer churn and start using uplift models*. *Information Sciences*. Publisher: Elsevier.
- Edwards, J.H. (1957). A note on the practical interpretation of 2x2 tables. *British Journal of Preventive and Social Medicine*, 11(2), 73. (Publisher: BMJ Publishing Group)
- Fernández-Loria, C., & Provost, F. (2022a). Causal classification: treatment effect estimation vs. outcome prediction. *Journal of Machine Learning Research*, 23(59), 1–35.
- Fernández-Loria, C., & Provost, F. (2022b). Causal decision making and causal effect estimation are not the same...and why it matters. *INFORMS Journal on Data Science*. Publisher: INFORMS.
- Fréchet, M. (1935). Généralisation du théorème des probabilités totales. *Fundamenta mathematicae*, 1(25), 379–387.
- Gubela, R. M., & Lessmann, S. (2021). Uplift modeling with value-driven evaluation metrics. *Decision Support Systems*, p. 113648. Publisher: Elsevier.
- Guelman, L. (2014). R Package ‘uplift’. CRAN
- Guelman, L., Guillén, M., & Préz-Marín, A.M. (2015). Uplift random forests. *Cybernetics and Systems*, 46(3–4), 230–248. (Publisher: Taylor & Francis) <https://doi.org/10.1080/01969722.2015.1012892>
- Gutiérrez, P., & Gérardy, J.-Y. (2016, January). Causal Inference and Uplift Modelling: A Review of the Literature. In: Hardgrove, C., Dorard, L., Thompson, K., & Douetteau F.(Eds.) *Proceedings of The 3rd International Conference on Predictive Applications and APIs* (Vol. 67, pp. 1–13). Microsoft NERD, Boston, USA: PMLR. Retrieved from <http://proceedings.mlr.press/v67/gutierrez17a.html> (Series Title: Proceedings of Machine Learning Research)
- Imbens, G. W., & Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Jaskowski, M., & Jaroszewicz, S. (2012). Uplift modeling for clinical trial data. In: *ICML Workshop on Clinical Data Analysis*.
- Jung, Y., Tian, J., & Bareinboim, E. (2021). Estimating Identifiable Causal Effects through Double Machine Learning. In: *Proceedings of the 35th AAAI Conference on Artificial Intelligence*.
- Karimi, A.-H., Schölkopf, B., & Valera, I. (2021). Algorithmic recourse: from counterfactual explanations to interventions. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 353–362).
- Künzel, S.R., Sekhon, J.S., Bickel, P.J., & Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. In: *Proceedings of the National Academy of Sciences of the United States of America*, 116(10), 4156–4165. (arXiv: 1706.03461) Publisher: National Acad Sciences) <https://doi.org/10.1073/pnas.1804597116>
- Li, A., & Pearl, J. (2019). Unit Selection Based on Counterfactual Logic. *IJCAI* (pp. 1793–1799). In: *International Joint Conferences on Artificial Intelligence Organization*. Retrieved from <https://doi.org/10.24963/ijcai.2019/248>
- Li, A., & Pearl, J. (2022). Unit selection with causal diagram. In: *Proceedings of the AAAI conference on artificial intelligence* (Vol. 36, pp. 5765–5772). (Number: 5)
- Lin, J. (2016). *On the dirichlet distribution*. *Master's Report*. Publisher: Queen's University Kingston Ontario, Canada.
- Liu, X.-Y., Wu, J., & Zhou, Z.-H. (2009). Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2), 539–550. (Publisher: IEEE) <https://doi.org/10.1109/tsmcb.2008.2007853>
- Mueller, S., Li, A., & Pearl, J. (2021). Causes of Effects: Learning individual responses from population data. arXiv preprint [arXiv:2104.13730](https://arxiv.org/abs/2104.13730) .
- Mueller, S., & Pearl, J. (2022). Personalized decision making—a conceptual introduction. arXiv preprint [arXiv:2208.09558](https://arxiv.org/abs/2208.09558)

- Óskarsdóttir, M., Van Calster, T., Baesens, B., Lemahieu, W., & Vanthienen, J. (2018). Time series for early churn detection: Using similarity based classification for dynamic networks. *Expert Systems with Applications*, 106, 55–65. (Publisher: Elsevier) <https://doi.org/10.1016/j.eswa.2018.04.003>
- Pearl, J. (2009). *Causality: models, reasoning, and inference*. Cambridge University Press.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55. Publisher: Oxford University Press.
- Tian, J., & Pearl, J. (2000). Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1), 287–313. Publisher: Springer.
- Verbeke, W., Dejaeger, K., Martens, D., Hur, J., & Baesens, B. (2012). New insights into churn prediction in the telecommunication sector: A profit driven data mining approach. *European Journal of Operational Research*, 218 (1), 211–229. Retrieved from (Publisher: Elsevier B.V.) <https://doi.org/10.1016/j.ejor.2011.09.031>
- Verbraken, T., Verbeke, W., & Baesens, B. (2013). A novel profit maximizing metric for measuring classification performance of customer churn prediction models. *IEEE Transactions on Knowledge and Data Engineering*. (Publisher: IEEE) <https://doi.org/10.1109/tkde.2012.50>
- Zhang, J., Tian, J., & Bareinboim, E. (2022). Partial counterfactual identification from observational and experimental data. In: *international conference on machine learning* (pp. 26548-26558). (tex.organization: PMLR)
- Zhang, W., Li, J., & Liu, L. (2021). A unified survey of treatment effect heterogeneity modelling and uplift modelling. *ACM Computing Surveys (CSUR)*, 54(8), 1–36. Publisher: ACM New York, NY.
- Zhu, B., Baesens, B., & vanden Broucke, S.K.L.M. (2017). An empirical comparison of techniques for the class imbalance problem in churn prediction. *Information Sciences*, 408, 84–99. (Publisher: Elsevier) <https://doi.org/10.1016/j.ins.2017.04.015>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.