

# Identifying counterfactual probabilities using bivariate distributions and uplift modeling

Théo Verhelst<sup>1\*</sup> and Gianluca Bontempi<sup>2†</sup>

1- Advanced Concepts Team, European Space Agency, Noordwijk, Netherlands

2- Machine Learning Group, Université Libre de Bruxelles, Brussels, Belgium  
WEL Research Institute, Wavre, Belgium

**Abstract.** Uplift modeling estimates the causal effect of an intervention as the difference between potential outcomes under treatment and control, whereas counterfactual identification aims to recover the joint distribution of these potential outcomes (e.g., “Would this customer still have churned had we given them a marketing offer?”). This joint counterfactual distribution provides richer information than the uplift but is harder to estimate. However, the two approaches are synergistic: uplift models can be leveraged for counterfactual estimation. We propose a counterfactual estimator that fits a bivariate beta distribution to predicted uplift scores, yielding posterior distributions over counterfactual outcomes. Our approach requires no causal assumptions beyond those of uplift modeling. Simulations show the efficacy of the approach, which can be applied, for example, to the problem of customer churn in telecom, where it reveals insights unavailable to standard ML or uplift models alone.

## 1 Introduction

Uplift modeling estimates individual-level causal effects (e.g., the impact of showing an ad to a specific user) by training predictors of the potential-outcome probabilities under alternative actions. Such models are typically trained on data from randomized campaigns, where the action assignment is randomized and outcomes are recorded. While running large-scale randomized studies can be costly in domains like medicine, they are comparatively inexpensive in digital settings (e.g., online marketing), where individual-level descriptors and outcomes are readily available. An uplift model trained on these data can then guide future campaigns by targeting users more effectively.

Counterfactual statements concern outcomes under hypothetical alternatives to the realized action. For example: “This customer churned because we did not make a marketing offer; had we made one, they would have stayed.” Counterfactuals are used across domains, from algorithmic recourse [1] to online advertising and customer relationship management [2]. Although uplift and counterfactuals are related, they are formally distinct: the counterfactual distribution characterizes probabilities of all realized/hypothetical outcome combinations, whereas uplift captures the change in outcome probability due to the action. The

---

\*Research carried out during the author’s prior affiliation with the Machine Learning Group.

†This publication benefits from the support of the Walloon Region to Pr. G. Bontempi as part of the funding for the FRFS WELT strategic axis.

counterfactual distribution is strictly more informative yet typically harder to estimate. As noted by [2], the two notions coincide under monotonicity (absence of negative effects).

Without structural assumptions, counterfactual probabilities are not identifiable from data alone; however, one can derive bounds, a task called *partial counterfactual identification* [3]. One line of literature derives such bounds under causal assumptions by combining experimental and observational data [3, 4, 5]. In a complementary direction, [6] use uplift scores learned from randomized campaigns to derive bounds on counterfactual probabilities and propose two point estimators: one based on an independence assumption between the potential outcomes, and one given by the midpoint of the bounds. We go further and adopt a Bayesian approach, fitting a parametric joint distribution over the uplift scores; posterior inference in this model yields a posterior distribution over counterfactual probabilities, from which point estimates can be derived. Compared to the baseline estimators of [6], our approach accommodates a broader class of joint distributions (and thus richer dependence structures) over the potential outcomes, which in turn yields more precise estimates of counterfactual probabilities. This uplift-based approach is supported by a mature literature on uplift modeling [7] and is particularly effective when combined with modern machine learning models. Furthermore, it is attractive in applications where graph-based assumptions are hard to justify but large-scale randomized campaigns are feasible.

The paper provides the first peer-reviewed presentation of four Bayesian counterfactual estimators introduced in extended form in [8]: the bivariate beta (BB) model and three extensions (NBB, GBB, and NGBB). We assess these estimators in two simulation settings and compare them to the independence and midpoint baselines of [6].

## 2 The Bayesian counterfactual approach

We denote random variables in boldface (e.g.,  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\mathbf{t}$ ) and their realizations by  $x$ ,  $y$ , and  $t$ . In uplift modeling,  $\mathbf{y} \in \{0, 1\}$  is the binary outcome,  $\mathbf{t} \in \{0, 1\}$  is the binary treatment/action, and  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$  is the feature vector. We use Pearl’s do-notation for causal interventions [9]: the potential outcome of  $\mathbf{y}$  under the intervention  $\text{do}(\mathbf{t} = t)$  is denoted  $\mathbf{y}_t$ . An uplift model estimates the individual-level effect

$$U(x) = P(\mathbf{y}_0 = 1 \mid \mathbf{x} = x) - P(\mathbf{y}_1 = 1 \mid \mathbf{x} = x). \quad (1)$$

A *counterfactual expression* is any logical expression involving multiple potential outcomes [10]. In the telecom churn setting,  $\mathbf{y} = 1$  denotes churn and  $\mathbf{t} = 1$  a marketing incentive, with features  $\mathbf{x}$ . The joint potential outcomes  $(\mathbf{y}_0, \mathbf{y}_1)$  define four customer types [6]: *sure things* (0, 0), *persuadables* (1, 0), *do-not-disturb* (0, 1), and *lost causes* (1, 1), whose *population-level* and *individual-level* proportions are given by the counterfactual probabilities:

$$p_{ij} = P(\mathbf{y}_0 = i, \mathbf{y}_1 = j) \quad \forall i, j \in \{0, 1\}, \quad (2)$$

$$p_{ij}(x) = P(\mathbf{y}_0 = i, \mathbf{y}_1 = j \mid \mathbf{x} = x). \quad (3)$$

The population-level counterfactual  $p_{ij}$  is the probability that a randomly drawn individual would have potential outcomes  $(\mathbf{y}_0, \mathbf{y}_1) = (i, j)$ , that is,  $\mathbf{y} = i$  under  $\mathbf{t} = 0$  and  $\mathbf{y} = j$  under  $\mathbf{t} = 1$ . Likewise, the individual-level counterfactual  $p_{ij}(x)$  is the same probability conditional on  $\mathbf{x} = x$ .

To estimate the probabilities  $p_{ij}$  and  $p_{ij}(x)$ , we assume access to an uplift model trained on data coming from a randomized campaign. The uplift model provides estimates of the scores

$$z_0(x) = P(\mathbf{y}_0 = 1 \mid \mathbf{x} = x) \quad \text{and} \quad z_1(x) = P(\mathbf{y}_1 = 1 \mid \mathbf{x} = x). \quad (4)$$

We further assume that the scores  $z_0(\mathbf{x}), z_1(\mathbf{x})$  follow a specific joint distribution, whose structure allows us to infer the counterfactual probabilities. Our approach relies on the identities

$$z_0(\mathbf{x}) = p_{10}(\mathbf{x}) + p_{11}(\mathbf{x}) \quad \text{and} \quad z_1(\mathbf{x}) = p_{01}(\mathbf{x}) + p_{11}(\mathbf{x}), \quad (5)$$

which follow directly from the definition of  $z_i(x)$  and  $p_{ij}(x)$  (Eqs. 3 and 4).

Our goal is to model jointly the observable uplift scores  $z_0(\mathbf{x}), z_1(\mathbf{x})$  and the latent counterfactual vector  $[p_{ij}(\mathbf{x})]_{ij}$ , which lies on the simplex by definition (Eq. 3). We do so with the bivariate beta construction of [11]: it introduces latent simplex-valued variables  $[\boldsymbol{\pi}_{ij}] \sim \text{Dir}(m)$  and maps them to a score pair  $(\boldsymbol{\zeta}_0, \boldsymbol{\zeta}_1) \sim \text{BB}(m)$  via the linear identities in Eq. (5). This yields a coherent generative model in which the latent Dirichlet distribution directly represents counterfactual probabilities, while the induced distribution over  $(\boldsymbol{\zeta}_0, \boldsymbol{\zeta}_1)$  has beta marginals suitable for probability scores.

Beyond the standard uplift assumptions (randomized treatment assignment, well-defined potential outcomes/SUTVA, and sufficiently calibrated estimates of  $z_t(x)$ ), our approach assumes that the population distribution of the true score pair  $(z_0(\mathbf{x}), z_1(\mathbf{x}))$  is reasonably approximated by the chosen parametric family (BB, or its extensions NBB/GBB/NGBB introduced below), i.e., that population heterogeneity in counterfactual vectors can be captured by a low-dimensional distribution on the simplex and translated into the empirical distribution of uplift scores.

Our approach consists of two phases:

1. **Learning phase:** we train an uplift model on a campaign dataset  $D = \{(x^{(i)}, y^{(i)}, t^{(i)})\}_{i=1}^N$  i.i.d. from  $(\mathbf{x}, \mathbf{y}, \mathbf{t})$ , resulting in estimators  $\hat{z}_t(x^{(i)})$  of the probability score  $z_t(x^{(i)})$  (with  $t = 0, 1$ ). Applying these two estimators on  $D$ , we obtain  $S = \{(\hat{z}_0(x^{(i)}), \hat{z}_1(x^{(i)}))\}_{i=1}^N$ . Class imbalance is mitigated with the EasyEnsemble strategy [12], and the scores are recalibrated using the formula proposed in [13]. Then, we fit a bivariate beta distribution  $\text{BB}(m)$  on the set of scores  $S$  using the method of moments.
2. **Inference phase:** Population-level counterfactuals are estimated using the fitted parameters of the bivariate beta distribution:

$$p_{ij} = \mathbb{E}[p_{ij}(\mathbf{x})] \simeq \mathbb{E}[\boldsymbol{\pi}_{ij}] = \frac{m_{ij}}{M} \quad (6)$$

where  $m_{ij}$  is a parameter of the bivariate beta distribution, and  $M = m_{00} + m_{10} + m_{01} + m_{11}$ .

Individual-level counterfactual probabilities for a new sample  $x'$  are estimated by computing the posterior probability distribution of  $\pi_{00}, \dots, \pi_{11}$  given the observed scores  $\zeta_0 = \hat{z}_0(x')$  and  $\zeta_1 = \hat{z}_1(x')$ . Our proposed point estimator of  $p_{00}(x'), \dots, p_{11}(x')$  is the expected value of this posterior distribution. It is estimated by integrating numerically a formula derived from the PDF of the bivariate beta distribution given in [11]. For more details on this procedure, we refer the reader to [8, Sec. 6.4.3].

We also propose three variations of this approach, with varying degrees of additional complexity. More details on their implementation can be found in [8, Sec. 6.4.4–6].

**Generalized Dirichlet distribution** The bivariate beta distribution ensures that the scores have marginal beta distributions, and it puts constraints on the possible covariance structures between the counterfactuals  $\pi_{ij}$ . We alleviate these constraints by replacing the underlying Dirichlet distribution on  $\pi_{ij}$  by the *generalized Dirichlet distribution* [14], which is a generalization of the Dirichlet distribution with additional parameters. The resulting distribution on the scores  $\zeta_t$  is named the *generalized bivariate beta distribution* (GBB for short).<sup>1</sup>

**Noisy predictions** Because the training set is finite, the estimated uplift scores  $\hat{z}_0(x)$  and  $\hat{z}_1(x)$  are noisy estimates of  $z_0(x)$  and  $z_1(x)$ . Yet, in the inference phase, we initially treat  $z_0(x)$  and  $z_1(x)$  as if they were known exactly. To relax this assumption, we introduce a hierarchical Bayesian extension of the bivariate beta distribution, in which beta noise is added on top of the latent scores  $\zeta_0, \zeta_1$ . The noisy scores are centered on the latent values, with a tunable spread parameter that can be adjusted to reflect the size of the training set. We refer to the resulting distribution as the *noisy bivariate beta distribution*.

**Generalized Dirichlet distribution with noisy predictions** Observe that the two variations presented above pertain to different aspects of the bivariate beta distribution, and therefore can be used in tandem, resulting in the *noisy generalized bivariate beta distribution*.

### 3 Performance assessment

In this section, we assess our proposed estimator with simulated data. Using simulated data allows us to compare the estimators with the ground truth, which is not feasible with real data. We consider two data-generating processes: one based on a Gaussian distribution and one based on a bivariate beta distribution. In the Gaussian-based simulation, the features  $\mathbf{x}$  are sampled from a multivariate standard normal distribution, and the potential outcomes  $\mathbf{y}_0, \mathbf{y}_1$  are generated via a linear combination of  $\mathbf{x}$  with fixed, randomly chosen weights. In the second simulation, we first fix the parameters of a bivariate beta distribution, then sample

<sup>1</sup>This is a slight abuse of terminology, since the marginals do not necessarily have beta distributions.

Table 1: Mean and standard deviation of the squared estimation error.

	Model	Dirichlet simulation	Gaussian simulation
Population-level	Independence	$(1.79 \pm 5.62) \times 10^{-3}$	$(1.25 \pm 0.32) \times 10^{-2}$
	Midpoint	$(1.16 \pm 3.63) \times 10^{-3}$	$(1.91 \pm 0.55) \times 10^{-2}$
	BB	$(0.87 \pm 1.76) \times 10^{-5}$	$(0.36 \pm 1.24) \times 10^{-5}$
	NBB	$(0.88 \pm 1.75) \times 10^{-5}$	$(0.35 \pm 1.24) \times 10^{-5}$
	GBB	$(0.86 \pm 1.72) \times 10^{-5}$	$(0.78 \pm 3.17) \times 10^{-6}$
	NGBB	$(0.88 \pm 1.75) \times 10^{-5}$	$(1.14 \pm 3.32) \times 10^{-6}$
Individual-level	Independence	$(0.33 \pm 1.12) \times 10^{-2}$	$(1.46 \pm 0.56) \times 10^{-2}$
	Midpoint	$(0.23 \pm 0.85) \times 10^{-2}$	$(2.23 \pm 0.71) \times 10^{-2}$
	BB	$(0.18 \pm 3.11) \times 10^{-2}$	$(0.67 \pm 1.56) \times 10^{-4}$
	NBB	–	$(0.72 \pm 1.45) \times 10^{-4}$
	GBB	$(0.14 \pm 2.78) \times 10^{-2}$	$(0.69 \pm 1.91) \times 10^{-4}$
	NGBB	–	$(0.71 \pm 1.38) \times 10^{-4}$

pairs of scores from this distribution. We add noise using a beta distribution, as in the noisy bivariate beta model. More details of the experimental protocol are provided in [8, Sec. 6.5]. The bivariate beta simulation is more flexible than the Gaussian simulation and satisfies the modeling assumptions of our counterfactual estimators; the Gaussian simulation therefore allows us to evaluate the impact of model misspecification. For both simulations, we generate 5000 samples and repeat this process 30 times with different simulation parameters to obtain average performance statistics for the different models.

We compare our proposed estimators to two baselines introduced in [6]. *(i) Independence baseline:* it assumes that the two uplift scores  $z_0(\mathbf{x}), z_1(\mathbf{x})$  are independent, which yields closed-form individual-level counterfactual estimates. *(ii) Midpoint baseline:* without additional assumptions, the scores  $z_0(\mathbf{x}), z_1(\mathbf{x})$  still partially identify the counterfactual distribution, yielding bounds on  $p_{ij}(\mathbf{x})$ ; the midpoint baseline sets the estimate to the midpoint of these bounds.

Table 1 reports squared errors for population-level (upper half) and individual-level (lower half) counterfactual estimation. Across both simulations, our Bayesian estimators outperform the independence and midpoint baselines by about two orders of magnitude. For population-level estimation, the generalized models (GBB and NGBB) can improve performance under misspecification (Gaussian simulation) while remaining comparable to BB/NBB under the well-specified Dirichlet simulation. For individual-level estimation, NBB and NGBB are omitted on the Dirichlet simulation due to long inference time; among the reported models, GBB is best on the Dirichlet simulation, whereas BB is marginally better on the Gaussian simulation, and differences between Bayesian variants are small relative to the gap to the baselines.

While the empirical analysis in this paper is restricted to simulated data, the same methods have already been applied successfully to real churn prevention

campaign data in the telecom sector; see [8, Sec. 6] for a detailed evaluation. Together with the performance gains observed on our simulations, this prior evidence indicates that the proposed improvements are promising for practical applications. Inference of counterfactual probabilities exposes causal patterns and untapped intervention opportunities that are otherwise hidden, such as the maximum profits that a campaign could have generated under perfect targeting, or the number of *do-not-disturb* customers [8, Sec. 6.6].

## References

- [1] Amir-Hossein Karimi, Bernhard Schölkopf, and Isabel Valera. Algorithmic recourse: from counterfactual explanations to interventions. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 353–362, 2021.
- [2] Ang Li and Judea Pearl. Unit selection based on counterfactual logic. IJCAI’19, pages 1793–1799. AAAI Press, 2019. Place: Macao, China Number of pages: 7.
- [3] Jin Tian and Judea Pearl. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1):287–313, 2000. Publisher: Springer.
- [4] Scott Mueller, Ang Li, and Judea Pearl. Causes of effects: Learning individual responses from population data. In *Proceedings of the thirty-first international joint conference on artificial intelligence, IJCAI-22*, pages 2712–2718. International Joint Conferences on Artificial Intelligence Organization, 2022.
- [5] Junzhe Zhang, Jin Tian, and Elias Bareinboim. Partial counterfactual identification from observational and experimental data. In *International conference on machine learning*, pages 26548–26558, 2022.
- [6] Théo Verhelst, Denis Mercier, Jeevan Shrestha, and Gianluca Bontempi. Partial counterfactual identification and uplift modeling: theoretical results and real-world assessment. *Machine Learning*, 2023.
- [7] Pierre Gutierrez and Jean-Yves Gérardy. Causal inference and uplift modelling: A review of the literature. In *International conference on predictive applications and APIs*, pages 1–13. PMLR, 2017.
- [8] Théo Verhelst. *Causal and predictive modeling of customer churn: lessons learned from empirical and theoretical research*. PhD thesis, Université libre de Bruxelles, 2024.
- [9] Judea Pearl. *Causality: models, reasoning, and inference*. Cambridge university press, 2009.
- [10] Elias Bareinboim, Juan D Correa, Duligur Ibeling, and Thomas Icard. On pearl’s hierarchy and the foundations of causal inference. In *Probabilistic and causal inference: the works of judea pearl*, pages 507–556. 2022.
- [11] Ingram Olkin and Thomas A Trikalinos. Constructions for a bivariate beta distribution. *Statistics & Probability Letters*, 96:54–60, 2015. Publisher: Elsevier.
- [12] Xu-Ying Liu, Jianxin Wu, and Zhi-Hua Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):539–550, 2009. Publisher: IEEE.
- [13] Andrea Dal Pozzolo, Olivier Caelen, Reid A Johnson, and Gianluca Bontempi. Calibrating probability with undersampling for unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence*, pages 159–166. IEEE, 2015.
- [14] Robert J Connor and James E Mosimann. Concepts of independence for proportions with a generalization of the Dirichlet distribution. *Journal of the American Statistical Association*, 64(325):194–206, 1969. Publisher: Taylor & Francis.