

Fixed point label attribution for real-time bidding

Martin Bompaire

Criteo AI Lab, m.bompaire@criteo.com

Antoine Désir

Technology and Operations Management, INSEAD, antoine.desir@insead.edu

Benjamin Heymann

Criteo AI Lab, b.heyman@criteo.com

This version: August 3, 2023

Problem definition: Most of the display advertising inventory is sold through real-time auctions. The participants of these auctions are typically bidders (Google, Criteo, RTB House, Trade Desk for instance) who participate on behalf of advertisers. In order to estimate the value of each display opportunity, they usually train advanced machine learning algorithms using historical data. In the labeled training set, the inputs are vectors of features representing each display opportunity and the labels are the generated rewards. In practice, the rewards are given by the advertiser and are tied to whether or not a particular user converts. Consequently, the rewards are aggregated at the user level and never observed at the display level. A fundamental task that has, to the best of our knowledge, been overlooked is to account for this mismatch and split, or attribute, the rewards at the right granularity level before training a learning algorithm. We call this the label attribution problem.

Methodology/results: In this paper, we develop an approach to the label attribution problem, which is both theoretically justified and practical. In particular, we develop a fixed point algorithm that allows for large scale implementation and showcase our solution using a large scale publicly available dataset from Criteo, a large Demand Side Platform. We dub our approach the Fixed Point Label Attribution (FiPLA) Algorithm.

Managerial implications: There is often a hidden leap of faith when transforming the advertiser's signal into display labelling. DSP providers should be careful when building their machine learning pipeline and carefully solve the label attribution step.

Key words. Online advertising, attribution, autobidding

1. Introduction

Digital advertising has been growing continuously since its inception and sustains a large part of the internet as we know it. It is indeed the main source of revenue for several tech giants (Google and Facebook in particular), and has attracted billions of dollars from advertisers over the last two decades. In 2019, the global digital ad spending exceeded 300 billion USD (Enberg 2019). What makes digital marketing so attractive is that marketers can harness the power of data to make

informed budget allocation decisions, and show the right banner to the right customer at the right time.

In this paper, we take the point of view of a *bidder* (Google, Criteo, RTB House, Trade Desk for instance) who participates in a real-time bidding exchange on behalf of an advertiser (Balseiro and Candogan 2017). Usually, such a bidder is rewarded when an ad is won, clicked on and followed in a short period of time by a conversion, typically a sale or any kind of action on the advertiser website such as creating an account or putting an item in a basket. Whenever the bidder receives a display opportunity, he has to estimate its potential reward and then, whatever the auction mechanism, calibrate his bid accordingly. The prediction of this potential reward is usually the result of a supervised learning algorithm that requires *labeled* training samples, where the inputs are vectors of features \mathbf{x} , which represent the display opportunity (these features encode contextual and user information) and the labels are the generated rewards. Table 1 shows a typical dataset.

Table 1 A typical dataset used by the bidder to estimate the potential reward of future display opportunities.

Features (\mathbf{x})					Reward
User ID	Timestamp	Clicked	Display Env.	User Funnel	Conversion
1	10	0	Desktop	Not exposed	0
1	11	1	Desktop	Exposed	1
1	12	0	Desktop	Engaged	0
2	20	1	Mobile	Not exposed	0
2	21	1	Mobile	Engaged	1
2	22	0	Desktop	Engaged	0
2	23	0	Desktop	Engaged	0
3	30	0	Mobile	Not exposed	0

Using a machine learning (ML) terminology, the reward corresponds to the label and this is in fact a typical supervised learning problem, where any ML algorithm can be used to predict the reward as a function of the display \mathbf{x} .

This seemingly straightforward and widely adopted approach unfortunately fails to recognize a crucial aspect of the data: a conversion is the result of a sequence of interactions with a user. Indeed, the rewards, which are decided by the advertiser, are in practice tied to a particular user and given not for a single display but for a sequence of interactions. An important step, which, to the best of our knowledge, has been overlooked is to account for this mismatch and split, or attribute the rewards at the display granularity. We dub this step the **label attribution**. Table 2 illustrates this step.

Table 2 Label attribution step: the reward is spread among the user’s displays instead of being arbitrarily assigned to one of them.

User	Features (\mathbf{x})			Reward	Valuation (y)
ID	Timestamp	Clicked	Display Env. ...	Conversion	Label
1	10	0	Desktop	1	→ 0.1
	11	1	Desktop		→ 0.7
	12	0	Desktop		→ 0.2
2	20	1	Mobile	1	→ 0.5
	21	1	Mobile		→ 0.3
	22	0	Desktop		→ 0.1
	23	0	Desktop		→ 0.1
3	30	0	Mobile	0	→ 0

Existing approaches (Chapelle 2014) omit this important step of converting the raw data given by the advertiser into a dataset that can be fed to a ML pipeline and attribute, somewhat heuristically, the reward given by the advertiser entirely to the last display. We refer to this approach as the *last touch* approach. Among the technical challenges a bidder is facing, recognizing this step and choosing a label attribution mechanism is often an understated problem.

In this paper, we formalize this missing step and propose a principled and scalable solution to predict the value of a given display opportunity that accounts for the fact that the reward given by the advertiser are aggregated at the user level.

1.1. Connection to attribution in online advertising

This problem is related to the well-studied question of attribution in online advertising, i.e., assigning conversions credits to individual marketing interactions. It is usually posed from the point of view of the advertiser who needs to decide how to split the credit between different advertising channels and is a fundamental question that informs the media mix optimization or the understanding of a customer’s journey. It has recently been identified as a top research priority by the Marketing Science Institute (Lemon 2016).

Similarly, in deciding how to split the reward given by the advertiser for sequences of display opportunities, the bidder also needs to perform his own attribution. To differentiate between these two levels of attribution, we refer to the bidder’s task as the *label attribution* problem. In fact, the reward received by the bidder is the result of the advertiser’s attribution. However, for our purposes, we assume that the attribution mechanism used by the advertiser, which impacts the reward of the bidder, is fixed and exogenous for the bidder. This leader-follower assumption is mild, since the advertiser attribution rule is almost always fixed in practice.

1.2. Motivating example

We next present an example to illustrate the limitations of ignoring the label attribution step and motivate the need for a more principled approach. Just like the last touch attribution used by advertisers, using a last touch approach for the label attribution fails to recognize the potential effect of early displays moving a customer up the conversion funnel and therefore fails to acknowledge their contribution into generating the reward. This leads to undervaluing early display opportunities and overvaluing later ones.

To make this even more concrete and build some intuition, assume for simplicity that there are only two types of display opportunities. Displays A influence users' behavior and increase their conversion probability. On the other hand, displays B do not influence the conversion probability at all but artificially always appear before a conversion, as is common for search ads (Blake et al. 2015). In such situation, the last touch label attribution wrongly allocates all the reward to displays B. Over time, the bidder will undervalue displays A and overvalue displays B. This in turn will lead to fewer conversions and therefore also lower the reward attributed to the bidder. This example highlights that the label attribution is a crucial step to align the bidder strategy with the given reward. It showcases the limitations of the current last touch approach and motivates the development of alternate methods of label attribution. We revisit this example in our numerical experiments in Section 5 and 7.

We end this example by noting that the motivation for using a last touch label attribution rule in practice seems to be driven by the misconception that if the advertiser uses a last touch attribution rule, which is a prevalent heuristic (Ji and Wang 2017, Diemert et al. 2017), then the bidder should in turn use a last touch label attribution rule. Since our example does not depend on the particular advertiser attribution rule, it demonstrates that this belief is wrong.

1.3. Contribution

The label attribution problem. In order to accurately bid in any auction, a bidder needs to learn a valuation function V to evaluate each display opportunity. The main challenge is that rewards are given at the user level, i.e., for sequences of displays or timelines, and cannot be used directly as labels. Consequently, we introduce in this paper the label attribution problem, which consists of finding a mapping μ that splits the reward for every user into labels for each individual display opportunity. These labels can then be used to learn the valuation function V using any appropriate ML algorithm. Our first contribution is to identify and formalize this label attribution problem which has not, to the best of our knowledge, been studied before. In particular, existing approaches implicitly assume that the label attribution step is done using a last touch method.

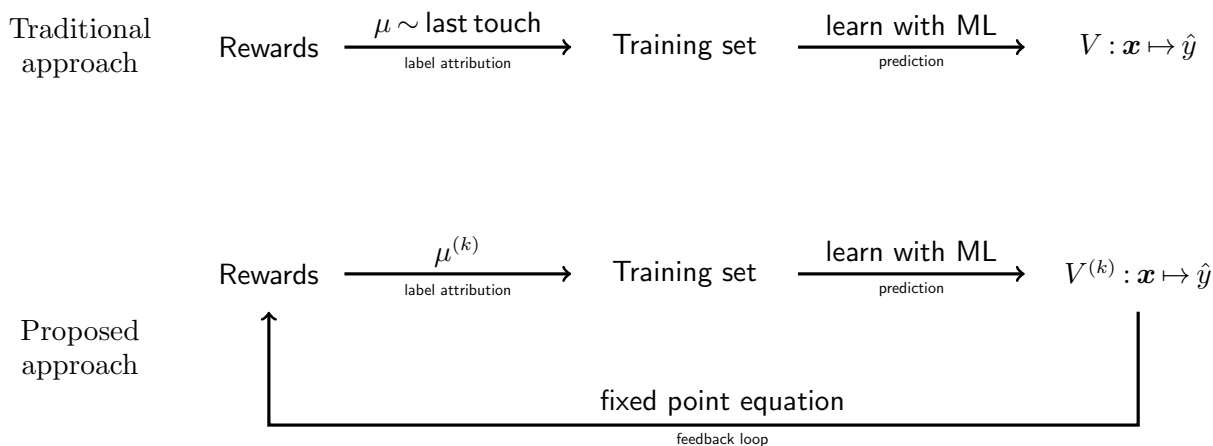


Figure 1 The traditional label attribution approach uses the last touch heuristic. By contrast, our proposed approach is an iterative algorithm that alternates between updating the attribution mapping $\mu^{(k)}$ and the prediction function $V^{(k)}$. The update of the mapping $\mu^{(k)}$ relies on a fixed point relation. Our framework can be implemented by adding a feedback loop to an existing ML pipeline.

A practical algorithm. A natural way to value a display is through its marginal contribution, i.e., the lift it provides to an existing sequence of displays. We call such a valuation *additive* as it decomposes the reward additively over the sequence of displays and next propose a practical algorithm to learn such an additive valuation. In particular, we argue that the label attribution needed to learn such valuation must satisfy a fixed point equation. We leverage this relation to propose an iterative algorithm that simultaneously learns the label attribution and valuation. We refer to our algorithm as the Fixed Point Label Attribution (FiPLA) Algorithm. Our approach is very practical as it can easily be embedded into an existing ML pipeline by adding a fixed point feedback loop.

More precisely, our approach is iterative and alternates between an update step and a predict step. In each iteration k , we maintain two mappings: a label attribution mechanism $\mu^{(k)}$, which is used to split the rewards, and a valuation mechanism $V^{(k)}$, which transforms any vector of features into a predicted reward. In the first step of each iteration, we update the mapping $\mu^{(k)}$ using the fixed point relation and the learned mapping $V^{(k)}$. We then, in a second step, split the rewards from the raw data using the label attribution $\mu^{(k)}$ and learn a new mapping $V^{(k+1)}$ using any ML algorithm. Figure 1 summarizes our framework.

Numerical experiments. We illustrate the benefits our approach through three different numerical experiments.¹

¹ All our codes are available at <https://github.com/criteo-research/robust-label-attribution>

1. First, we explore what the FiPLA Algorithm learns in simple synthetic settings. We showcase that our approach can correctly capture critical aspects of typical user behaviors such as a decreasing marginal return effect of advertising. More importantly, we highlight the limitations of the last touch approach.

2. We then test our approach on a large scale dataset. This real dataset exhibit a large feature space and we discuss how the FiPLA can be implemented in such a setting. We compare the output of the FiPLA Algorithm to the typical last touch approach both qualitatively and using an independently proposed performance metrics.

3. Finally, we simulate a real bidding environment and measure the benefits of using our approach on the bidder’s revenue. In our simulated settings, the profit associated with using the FiPLA Algorithm is more than 3.5 times higher than using a last touch approach, even when we allow for shading. This shows that the potential benefits of carefully addressing the label attribution step can be significant.

A principled approach. We give some theoretical underpinning for our approach and mathematically formalize the relation between the fixed point equation and the additivity property for a slightly stylized model. In particular, under some assumptions, we show that the additivity property defines a unique label attribution mechanism and that the sequence of mapping $V^{(k)}$ converges to the additive valuation. We further show that this intuitive label attribution enjoys several interesting properties. First, we show that bidding according to this additive valuation is myopic optimal, i.e., optimal if there are no other future opportunities. Moreover, we show that our proposed approach is distributionally robust, i.e., it does not depend on how many times each set of features has been observed. This is very useful from a learning perspective.

Agenda. The paper is organized as follows. We begin with a review of the literature in Section 2. We then introduce the notations that allow us to properly frame the label attribution problem in Section 3. We present our proposed learning framework in Section 4. In Section 5, we illustrate our method on a series of synthetic experiments. We then test our framework on a large publicly available dataset in Section 6. In Section 7, we evaluate the impact on revenue of our proposed approach. The theoretical underpinning of our method is relegated to Appendix A.

2. Related Work

We survey the immediate related literature and refer to Choi et al. (2020) and Wang et al. (2017) for more background on display advertising.

Most of the advertising is done through auctions, and we thus take the point of view of a bidder in this paper. Regardless of the auction mechanism, the bidder has to estimate the value he could generate from each display opportunity and then calibrate his bid accordingly. For example, in

a second price setting, for a given display opportunity, the standard in the industry is to bid the estimated average value earned with this specific impression (Perlich et al. 2012, Wang et al. 2017). This is referred to as *value based bidding* in Xu et al. (2016). We would like to emphasize that this behavior is *myopically* designed in the sense that the future policy is disregarded or equivalently that the bid is computed as if this was the last opportunity to generate a conversion. Our approach to the label attribution problem adopts a similar philosophy and offers a principled myopic optimal approach. Without loss of generality, we assume that the bidder is compensated proportionally to the reward, i.e., the advertiser’s attribution, that he received. Note that the bidder might need to compute some shading factor to adapt to the competition and the auction mechanism (Balseiro and Gur 2019, Balseiro et al. 2021) or solve a broader optimization problem to reach certain requirements given a budget constraint ?. However, these computations are beyond the scope of this paper.

Realizing the need for more principled solutions to the advertiser attribution problem, recent works have started proposing algorithmic approaches to go beyond rule-based heuristics such as last touch. Using different generative models, Danaher (2018) and Anderl et al. (2016) both propose a counterfactual approach. The notion of Shapley value (Berman 2018, Singal et al. 2019, Dalessandro et al. 2012) has also been used to split the credit fairly among different channels. Johnson et al. (2017), Lewis et al. (2011), Lewis and Rao (2015) all study models related to causality and advertiser attribution. Their approach is based on econometric parametric models.

Our research is inspired by the multi-touch attribution (MTA) literature (Shao and Li 2011, Ji et al. 2016, Arava et al. 2018). This stream of work aims at designing machine learning algorithms that assign fractional credits to different touchpoints from the advertiser’s standpoint. They also aim at solving the advertiser attribution problem but focus on the implementability of the solution rather than its theoretical justification. Despite using a different formalism, our work relates to the survival analysis approach used in Laub et al. (2015). Moreover, the fixed point-based algorithm of Section 4 shares some characteristics with the EM algorithm in Zhang et al. (2014). Similarly, Ji and Wang (2017) relies on a Weibull model and is very close to our work. However, their choice of parametrization hides the fixed point property identified in Zhang et al. (2014).

On top of the advertiser attribution problem, some researchers have also considered the bidder’s problem as we do and tried to improve the buying process using machine learning pipelines. Chapelle et al. (2014) presents a bidding architecture, which relies on an implicit last touch label attribution. Chapelle (2014) provides a technique to deal with the delay between displays and sales. Diemert et al. (2017) proposes to model the advertiser attribution to increase bidding performances. Their experimental results motivate our approach. More recently Bompaire et al. (2020) provides a framework based on reinforcement learning and causal inference.

Finally, we would like to highlight that the question of attribution is not specific to online advertising. For instance, recent work in interpretable machine learning tries to attribute the prediction of a machine learning model to individual features of the input (Dhamdhere et al. 2018). In a very different context, Flores-Szwagrzak and Treibich (2020) develops a method to disentangle individual performance from team productivity. In each of these examples, the goal is to attribute some aggregate output to individual components of the input. Although Flores-Szwagrzak and Treibich (2020) develops an iterative algorithm, which is similar in spirit to our approach, their problem of splitting credit among team members is static in nature, while time is an important aspect of our problem. Another important difference is that we propose a machine learning implementation of our algorithm, which is able to leverage the presence of context features.

3. The Label Attribution Problem

In this section, we formalize the label attribution problem. We begin by introducing the model and some notations.

3.1. Model and notations

In practice, a bidder is given a dataset of T observations $\{\mathbf{x}^1, \dots, \mathbf{x}^T\}$, where for each $t \leq T$, \mathbf{x}^t represents a display opportunity. Due to our focus on display advertising, we refer to an element \mathbf{x} as a display or display opportunity in the paper. However, these could in principle represent any types of marketing interactions. We let $\mathcal{X} \subseteq \mathbb{R}^d$ be a set of potential display opportunities. Each $\mathbf{x} \in \mathcal{X}$ is a vector that encodes all the features associated with a display that the bidder can use to evaluate the display opportunity. For instance, this could include the type of products or creatives as well as contextual information such as whether or not the user has already seen an ad. In particular, we can also capture customer heterogeneity by encoding the user’s features within \mathbf{x} . In general, any feature that is useful for the bidder can be encoded in the display set \mathcal{X} .

Reward. Additionally, each display \mathbf{x} is associated with a reward. Most often, this reward represents whether the display lead to a conversion or not. Importantly, we assume that this reward is given for a subset of displays. For ease of exposition, we assume here that a group of display corresponds to the interaction with a specific user. However, nothing prevents the bidder to define a different grouping. More specifically, we assume that each display is associated with a user u^t . We denote by \mathcal{U} the set of users. Each user $u \in \mathcal{U}$ is associated with a reward r_u . By doing this, we recognize that the reward r_u is generated by the set of displays $\mathcal{X}_u = \{\mathbf{x}^t : \forall t, u^t = u\}$. From the bidder perspective, the value of the reward comes from a black-box decided by the advertiser. In the case where the bidder and advertiser are the same, these rewards could represent sales or any objective the advertiser is optimizing for. However, in general, the bidder and advertiser are different actors.

Valuation. The bidder needs to estimate the value for each new display opportunity. We capture this using a valuation $V : \mathcal{X} \rightarrow \mathbb{R}^+$, which gives the value of a display opportunity. The bidder’s challenge is to appropriately learn a valuation V from the reward samples, given at the user level.

3.2. Label attribution

Since the rewards are given for each user, the reward data cannot be used directly to learn V . In order to train a ML algorithm, a training set has to be constructed from the reward data to learn V . This key step, which we call the label attribution, is often overlooked. It is nevertheless crucial, as different label attribution schemes lead to different valuation functions. One of our key contribution is to formalize this step.

DEFINITION 1 (LABEL ATTRIBUTION). A **label attribution** μ is a mapping from $\mathcal{X} \times \mathcal{U}$ to \mathbb{R}^+ that satisfies the following properties.

1. $\forall u \in \mathcal{U}$ and $\forall \mathbf{x} \notin \mathcal{X}_u$, $\mu(\mathbf{x}|u) = 0$.
2. $\forall u \in \mathcal{U}$, $r_u = \sum_{\mathbf{x} \in \mathcal{X}_u} \mu(\mathbf{x}|u)$.

For a label attribution μ and display $\mathbf{x} \in \mathcal{X}$, $\mu(\mathbf{x}|u)$ represents the label of the display \mathbf{x} for user u , i.e., the part of the reward r_u attributed to display \mathbf{x} . The first property states that we can only split the reward of a user among its displays. The second property states that the attribution has to be balanced and does not attribute more than the reward r_u . Just like the advertiser’s attribution, the label attribution is splitting credit. However, the split is done at the more granular level of the displays. Indeed, the bidder needs to predict the value of a display opportunity in order to calibrate his bid in real-time auctions. With this definition, a typical flow can be described as follows:

1. **Label attribution step.** The label attribution μ is used to create a training set

$$S = \{(\mathbf{x}^1, y^1), \dots, (\mathbf{x}^T, y^T)\},$$

where for each t , $y^t = \mu(\mathbf{x}^t|u^t)$.

2. **Prediction step.** Some ML algorithm, typically a logistic regression, is trained to predict the display labels y from the corresponding vectors of features \mathbf{x} and the resulting mapping constitutes the estimated valuation \hat{V} .

As a concrete example, we can revisit the typical last touch approach with our notation. In particular, the last touch label attribution can be written for all $(\mathbf{x}, u) \in \mathcal{X} \times \mathcal{U}$ as

$$\mu^{\text{LT}}(\mathbf{x}|u) = \begin{cases} r_u, & \text{if } \mathbf{x} = \mathbf{x}_u^{\text{LT}}, \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathbf{x}_u^{\text{LT}} \in \mathcal{X}_u$ is the last display shown to user u .

We would like to end this section by emphasizing a key difference between the valuation and the label attribution. On the one hand, the label attribution is backward looking, i.e., it is a function of the entire set of displays a user has been exposed to. On the other hand, the valuation is forward-looking and only has access to the current display opportunity. Indeed, the valuation is to be used for bidding purposes.

4. The Additive Valuation

In this section, we propose a new approach to solve the label attribution problem, which is more principled and can efficiently be implemented in practice.

4.1. Additivity and fixed point equation

We begin by arguing for an intuitive property that we wish our valuation to satisfy. In particular, a very natural way to evaluate the next display opportunity is through its marginal contribution. This method has recently received attention in the industry (Perlich et al. 2012, Wang et al. 2017). It is also sometimes referred to as lift-based valuation (Xu et al. 2016). We formalize this property using a notion of additivity for our valuation function.

DEFINITION 2 (ADDITIVITY). Let V be a valuation. We say that V satisfies the additivity property if for all $u \in \mathcal{U}$ and $\mathbf{x} \in \mathcal{X}_u$,

$$V(\mathbf{x}) = r_{u'} - r_u,$$

where $\mathcal{X}_{u'} = \mathcal{X}_u \cup \{\mathbf{x}\}$.

In other words, the value associated with \mathbf{x} is the incremental reward. Indeed, in the definition, users u' and u have been exposed to the same set of display ads except for \mathbf{x} , which only user u' has seen. The additivity property hinges on a temporal marginality, i.e., we are quantifying the marginal contribution with respect to the previous actions, and is in line with several works on the additivity effect of ads (Shao and Li 2011, Ji and Wang 2017, Zhang et al. 2014, Dalessandro et al. 2012). This is different than a counterfactual marginality, used for instance in Danaher (2018), Anderl et al. (2016), Bompaire et al. (2020), which quantifies the marginal contribution with respect to an alternative action. The counterfactual marginality depends on the policy used in the future, while the temporal marginality does not.

To learn such an additive valuation, we leverage a fixed point equation between the reward, the valuation and the label attribution. In particular, in an ideal setting where V satisfies Definition 2 and is known, the label attribution would split the reward proportionally to V , i.e., for all $\mathbf{x} \in \mathcal{X}$ and $u \in \mathcal{U}$,

$$\mu(\mathbf{x}|u) = \frac{V(\mathbf{x})}{\sum_{\mathbf{x}' \in \mathcal{X}_u} V(\mathbf{x}')} \cdot r_u. \quad (1)$$

It is immediate to see that this is a valid label attribution. In Appendix A, we prove, using a slightly more abstract setting, that a valuation satisfies Definition 2 if and only if there exists a label attribution that satisfies Equation (1) thereby formalizing the connection between proportional label attribution and additive valuation (see in particular Proposition 1). The theoretical justifications are relegated to the Appendix and we focus here on how to practically learn an additive valuation. In particular, we next show how to leverage the fixed point equation (1) to design a learning algorithm. Before describing our algorithm, we introduce a measure of convergence, which is motivated by Definition 2.

4.2. A measure of convergence

As explained in the previous section, we want to develop an algorithm that learns an additive valuation, i.e., that is able to capture the lift in the reward provided by showing an ad. We use this intuition to propose a measure of convergence for our algorithm. More precisely, for every valuation V , we define \mathcal{L}^{add} as

$$\mathcal{L}^{\text{add}}(V) = \frac{1}{T} \cdot \sum_{u \in \mathcal{U}} \left[r_u \cdot \ln \left(\sum_{\mathbf{x}: u^t = u} V(\mathbf{x}) \right) - \sum_{\mathbf{x}: u^t = u} V(\mathbf{x}) \right].$$

The function \mathcal{L}^{add} can be thought of as a criterion to measure how well some valuation V satisfies the additivity property. We use this loss to measure the progress of our algorithm, which we describe in the next section. Note that this loss resembles a likelihood function. However, \mathcal{L}^{add} is not per se a likelihood function since we do not have an underlying model whose parameters we are trying to find. In Appendix B.3, we show that under some assumptions, a valuation that satisfies Definition 2 minimizes $\mathcal{L}^{\text{add}}(V)$ over all possible valuations.

4.3. Fixed point label attribution (FiPLA) algorithm

In Section 4.1, we introduce the fixed point equation (1) relating the valuation, the label attribution and the reward. However, in practice, we do not know V . For that reason, we propose a practical procedure that simultaneously learns V and μ . The proposed algorithm is iterative and maintains, in each iteration k , a label attribution $\mu^{(k)}$ and a valuation $\hat{V}^{(k)}$. Each iteration k consists of three steps.

- The first two steps are a Label Attribution step and a Prediction step. These are similar to the one described in Section 3.2.
- The third step is a fixed point step that exploits Equation (1) to update the label attribution $\mu^{(k+1)}$ using the valuation computed in the previous step $\hat{V}^{(k+1)}$.

Algorithm 1 summarizes the procedure, which we refer to in the rest of the paper as the Fixed Point Label Attribution (FiPLA) Algorithm. Compared to existing approaches, the main novelties

ALGORITHM 1: Fixed Point Label Attribution (FiPLA) Algorithm

Input: A dataset of displays $\{\mathbf{x}^1, \dots, \mathbf{x}^T\}$, a set of users \mathcal{U} and associated rewards $r_1, \dots, r_{|\mathcal{U}|}$ and an initial attribution μ^0

Output: A valuation \hat{V}^{FP}

while $|\mathcal{L}^{\text{add}}(\hat{V}^{k+1}) - \mathcal{L}^{\text{add}}(\hat{V}^k)| \geq \epsilon$ **do**

1. *Attribution step.* Generate training dataset by attributing rewards at the display level with the label attribution μ^k

$$\mathcal{S}^{(k)} = \{(\mathbf{x}^1, y^1), \dots, (\mathbf{x}^T, y^T)\},$$

where for each t , $y^t = \mu^k(\mathbf{x}^t | u^t)$.

2. *Prediction step.* Train any *ML algorithm* on the training set $\mathcal{S}^{(k+1)}$. Update $\hat{V}^{(k+1)}$ to the resulting mapping.

3. *Fixed point update.* Update the label attribution $\mu^{(k+1)}$ with the fixed point characterization

$$\mu^{(k+1)} = (\mathbf{x}^t | u^t) \mapsto \frac{\hat{V}^{(k+1)}(\mathbf{x}^t)}{\sum_{\mathbf{x}' \in \mathcal{X}_{u^t}} \hat{V}^{(k+1)}(\mathbf{x}')} \cdot r_{u^t}.$$

end

of our approach are (1) explicitly doing the label attribution (Step 1) instead of implicitly using a last touch label attribution, and (2) adding a fixed point update (Step 3). As also illustrated in Figure 1 in the Introduction, our approach suggests a change in paradigm whereby instead of estimating the valuation from a fixed internal procedure, such as last touch attribution, the optimal pair of valuation and label attribution are rather iteratively and jointly updated until convergence. Moreover, this feedback loop is a simple change that can easily be implemented in an existing ML pipeline.

Finally, we provide theoretical justifications to our algorithm and prove in Proposition 2 of Appendix A that the procedure converges, under some assumption. This is done by relating our iterative procedure to a majorize-minorize algorithm.

5. Synthetic Experiments

In this section, we present some numerical experiments using synthetic data to build some intuition. The goal of this controlled experiment setup is to illustrate the benefits of using the FiPLA Algorithm over the last touch approach.

5.1. Experimental setups

We present three different settings, each with a different generative process.

1. **Constant valuation.** In this setup, we construct user timelines as follows. At each time step, the user is exposed to an ad and converts with a constant probability α . Additionally, independent

of everything else, at each time step, there is also a constant probability of β that the user leaves the system. For each user timeline, we let the reward be equal to the number of conversions. In this case, the only feature that we use is the number x of impressions, which in this case is also the length of the user timeline. Note that, in expectation, the reward grows proportionally to the number of impressions and we expect $V^{\text{FP}}(x)$ to be constant equal to α .

2. Decreasing marginal return. In this second experiment, the probability of getting a reward is still increasing with the number of ads seen, but we now assume that it exhibits a decreasing marginal return effect. In particular, we use the same setup as previously except in the way we compute the reward. More precisely, for this setting, we associate a reward of 1 to a user’s timeline if the user converted at least once. In this case, the reward grows exponentially to one as the number of impressions increases and we expect $V^{\text{FP}}(x)$ to decay like $(1 - \alpha)^x$. This aims at capturing a decreasing marginal return effect of advertising (Chapelle et al. 2014, Diemert et al. 2017, Zhang et al. 2014).

3. Two display types. In this third setup, we revisit the motivating example from Section 1.2 and illustrate a multi-dimensional feature space. We assume that in each time step there is an impression for a type A display with probability 0.5 and an impression for a type B display otherwise. Building on Setup 1, we assume that the user converts with a constant probability α_A after a display of type A and α_B after a display of type B . Independent of everything else, there is also a constant probability β that the user leaves the system in each time step. Let x , resp. y , be the number of displays of type A , resp. type B , a user has been exposed to. Additionally, we let $z \in \{A, B\}$ denote the type of the current display opportunity. We apply the FiPLA Algorithm with $\mathcal{X} = \{(x, y, z) : x \in \mathbb{N}, y \in \mathbb{N}, z \in \{A, B\}\}$. In this case, generalizing the first setup, we expect $V^{\text{FP}}(x, y, A)$ (resp. $V^{\text{FP}}(x, y, B)$) to be constant equal to α_A (resp. α_B) for all values of x and y .

5.2. Results

In each experiment, we generate timelines for 3M users and compare the valuation function \hat{V}^{LT} learned using a last touch approach with the valuation function \hat{V}^{FP} learned with the FiPLA Algorithm. Figure 2 shows how the loss progresses with the number of iterations for Setup 1. This is to illustrate that the number of iterations needed to run the FiPLA Algorithm in practice remains very reasonable. Even though we do not show these plots for all our experiments, the number of iterations remains similar throughout all the experiments of this section.

Setup 1. Figure 3a shows the learned valuations $\hat{V}^{\text{LT}}(x)$ and $\hat{V}^{\text{FP}}(x)$ for Setup 1. We observe that the FiPLA Algorithm correctly learns the underlying true valuation since $\hat{V}^{\text{FP}}(x)$ is constant equal to $\alpha = 0.1$. On the other hand, we see that $\hat{V}^{\text{LT}}(x)$ predicts a higher valuation when x is higher. This is precisely what we eluded to in the motivating example where the last touch approach tends


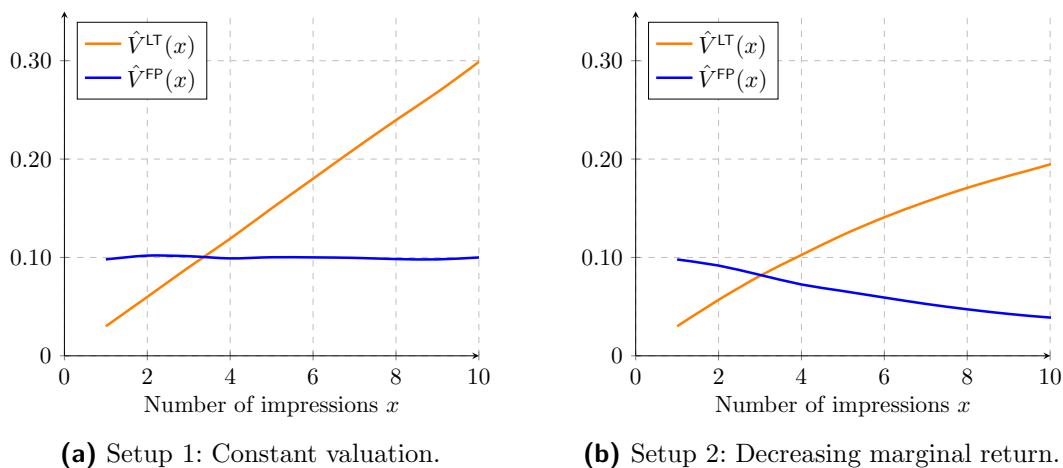


Figure 2 Convergence of the FiPLA Algorithm for Setup 1.

to overestimate the value of impressions that appear later in the user journey and underestimate those that appear earlier. Indeed, if the label attribution is driven by the last touch philosophy, then an impression is evaluated through its probability of being the last impression. This is why we observe that $\hat{V}^{\text{LT}}(x)$ increases with x whereas the true underlying value of each impression remains constant.



(a) Setup 1: Constant valuation.

(b) Setup 2: Decreasing marginal return.

Figure 3 Learned valuations in synthetic experiments for Setups 1 and 2. Here we use $\alpha = 0.1$ and $\beta = 0.3$. Note that we only show the results for x less than 20 since for larger values of x , we have less than 100 data points in the training set.

Setup 2. Figure 3b shows the the learned valuations $\hat{V}^{\text{LT}}(x)$ and $\hat{V}^{\text{FP}}(x)$ for Setup 2. In this case, recall that the underlying model implies that the true valuation is decreasing convex. In other words, the more ads a user have been exposed to, the less valuable the next ad is. This is exactly what $\hat{V}^{\text{FP}}(x)$ learns as illustrated by the convex decreasing shape. On the other hand, $\hat{V}^{\text{LT}}(x)$

is increasing, clearly failing to capture the decreasing marginal return effect from the generative process.

Setup 3. In this experiment, the feature space has dimension three and we need to learn $V(x, y, z)$, where x is the number of displays of type A (the useful displays), y is the number of displays of type B (the useless displays) and z is the type of the current display opportunity. Figure 4 shows the results of the learned valuations.

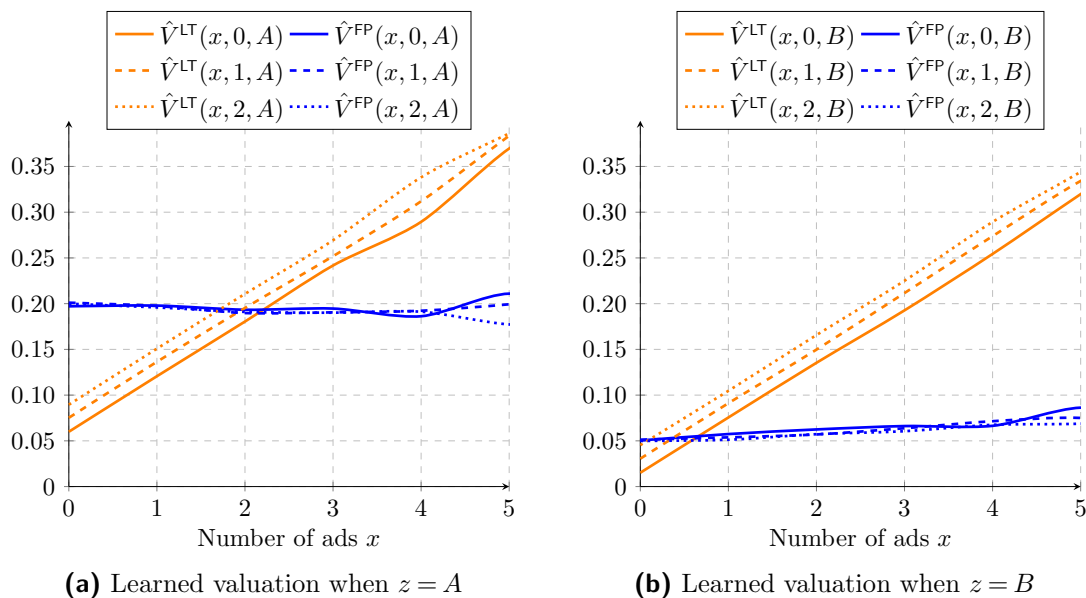


Figure 4 Learned valuations in synthetic experiments for Setup 3. Here we use $\alpha_A = 0.20$, $\alpha_B = 0.05$ and $\beta = 0.1$.

Looking at what the FiPLA Algorithm learns when $z = A$, i.e., the impression is a display of type A, we observe in Figure 4a that $\hat{V}^{FP}(x, y, A)$ correctly recovers the constant valuation from Setup 1. Additionally, we see that the valuation does not change as a function y , which is consistent with the generative model. When $z = B$, we observe in Figure 4b that $\hat{V}^{FP}(x, y, B)$ also recovers the right constant valuation. On the other hand, when the label attribution is done using last touch, we observe in Figure 4a that the valuation $\hat{V}^{LT}(x, y, A)$ increases with x as in Figure 3a. Furthermore, as illustrated in Figure 4b, \hat{V}^{LT} is in fact almost the same whether $z = A$ or $z = B$ and therefore the last touch method does not differentiate the different types of ads.

After seeing that our approach correctly captures critical aspects of typical behaviors and overcomes last touch weaknesses, we take it next to a real dataset.

6. “Criteo Attribution Modeling for Bidding” dataset

In this section, we apply our framework to a real dataset with a large feature space. The feature space size is very high and each input is unique with very high probability, as it is often the case in real data sets. In this type of contexts, we often rely on ML techniques.

6.1. Data description

“The dataset represents a sample of 30 days of Criteo live traffic data. Each line corresponds to one impression (a banner) that was displayed to a user. For each banner we have detailed information about the context, if it was clicked, if it led to a conversion and if it led to a conversion that was attributed to Criteo or not. Data has been sub-sampled and anonymized so as not to disclose proprietary elements” (Diemert et al. 2017). In total, the dataset contains approximately 16M displays shown to 6M users among which 180K have converted. Each display has about 20 features including 10 categorical features describing contextual information with user-specific information in particular. The nature of these categorical features is not disclosed but they are meant to be “used to learn the click/conversion models”. We thus consider them as correct descriptors of the user’s features at the time of the impression. This reflects real settings in which feature engineering is made to transform non structured user’s history data into features that are suitable for a machine learning training. Some other notable features include whether the display was clicked or not, the time since the last click and the number of displays clicked by this user (see Diemert et al. (2017) for a complete description of the Criteo Attribution Modeling for Bidding Dataset).

6.2. Preprocessing and implementation details

Preprocessing. In order to be closer in spirit to the standard last touch benchmark and have stronger signals, we only keep displays that were clicked. Moreover, to mitigate the side effects caused by the fixed window of observation, we remove the users for which some clicks are missing since such information can be inferred from the data. Additionally, to compare our results with existing metrics (see “Performance” paragraph in Section 6.3 below), we need binary rewards. Consequently, we split each user that triggered multiple conversions into multiple users. The resulting dataset contains 16M displays shown to 8M users among which 196K have converted.

Hashing trick. For the prediction step in the FiPLA Algorithm, we use a regularized logistic regression together with a hashing trick described in Chapelle et al. (2014). For completeness, we give more details on this procedure. First, we treat all our variables as categorical features. To do so, we discretize the numerical features and group values into buckets. In our dataset, there are only two numerical features: “time since last click” and “number of clicked displays”. In logistic regression, categorical features are usually handled with one hot encoding where the size of the

encoding is equal to the number of categories. For instance, suppose that there are three types of display, then the feature column corresponding to the display type is transformed as follows.

$$\text{Display A} \rightarrow [1, 0, 0]$$

$$\text{Display B} \rightarrow [0, 1, 0]$$

$$\text{Display C} \rightarrow [0, 0, 1]$$

When the number of possible values for each categorical value explodes, this method is not scalable. This is where the hashing trick is useful since it allows mapping the input to a lower dimensional space of fixed and pre-determined size denoted by m . More precisely, this is done by using a hashing function f that transforms the one-hot encoding vector to a lower dimensional vector. For our experiments, we have chosen $m = 2^{13} = 8,192$. This is similar in spirit to an auto-encoder in ML that learns a lower dimensional representation of the data. In this case, note that the mapping is not learned and there could potentially be better ways of projecting the data to a lower dimensional space. However, this method is standard in the industry and efficient enough to run meaningful experiments. Finally, note among the many implementations of the hashing trick that exist, we use the one implemented in `scikit-learn` (?).

6.3. Results and discussion.

We use a 80/20 split on the user identifiers, where we train our model using 80 % of the dataset and test on the remaining 20 %.

Convergence. We check the convergence of our algorithm using $\mathcal{L}^{\text{add}}(V^{(k)})$ as a criterion. Figure 5a shows that the procedure converges in a few iterations similar to the previous set of experiments. Figure 5b also illustrates that our solution generalizes to the testing dataset. Note that since we use a regularized logistic regression in the prediction step of the FiPLA Algorithm, this introduces a small bias as can be seen by the non-monotonicity of the iterates in Figure 5. Nevertheless, we observe a significant improvement compared to the last touch method.

To put these results in perspective, we additionally tested the sensitivity of our results with respect to the size of the feature space. Indeed, with the hashing trick, we are embedding our feature space into a target space whose size can be controlled. Figure 6 illustrates the value of $\mathcal{L}^{\text{add}}(\hat{V}^{\text{FP}})$ under different feature space sizes, where recall that \hat{V}^{FP} is the output of the FiPLA Algorithm. As expected, $\mathcal{L}^{\text{add}}(\hat{V}^{\text{FP}})$ increases with respect to the feature space size. On the other hand, there is a priori no reason to believe that when increasing the size of the feature space, $\mathcal{L}^{\text{add}}(V^{\text{LT}})$ also increases. Indeed, the last touch method, unlike our valuation, is not designed to maximize \mathcal{L}^{add} . Somewhat reassuringly, we observe that $\mathcal{L}^{\text{add}}(V^{\text{LT}})$ also increases with the feature space size. Interestingly, we observe that the improvement, as measured by \mathcal{L}^{add} , of going from V^{LT} to \hat{V}^{FP} are similar in magnitude to increasing the size of the feature space by a factor five, as illustrated by Figure 6.

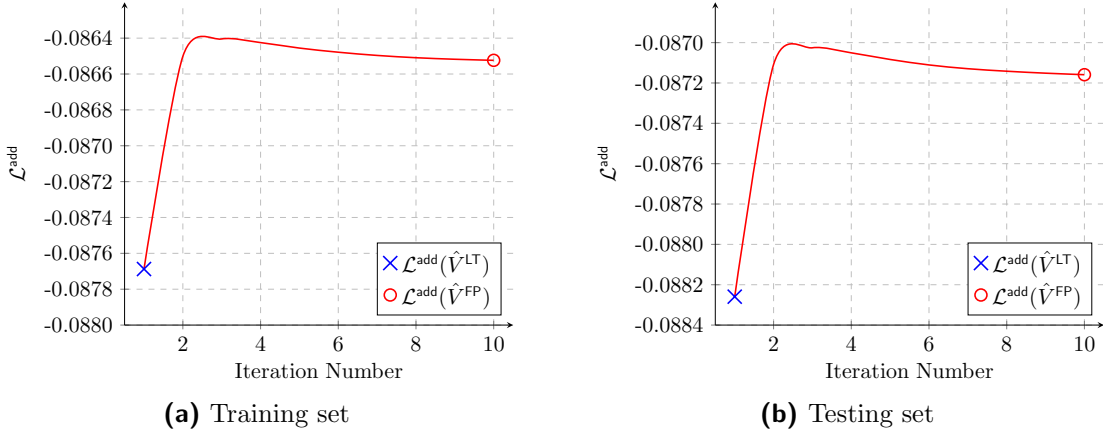


Figure 5 Convergence $\mathcal{L}^{\text{add}}(V^{(k)})$ on the Criteo Attribution dataset.

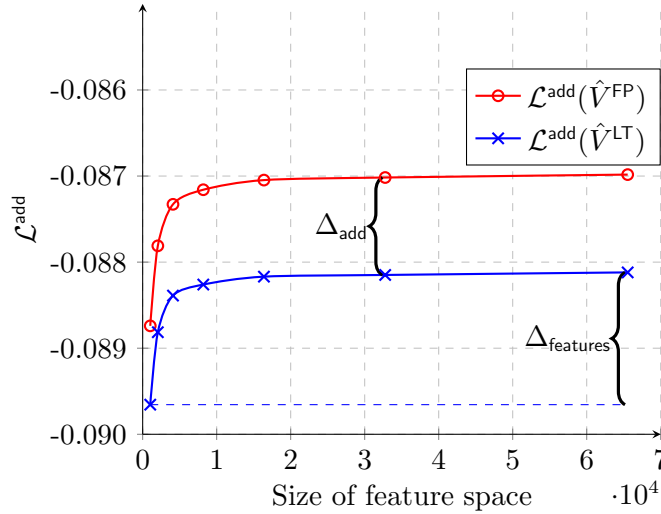
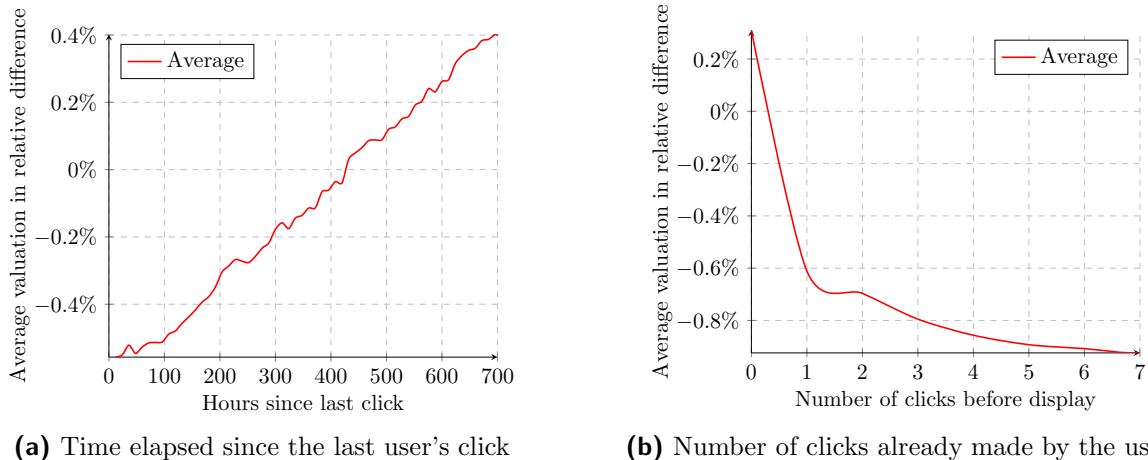


Figure 6 Additivity loss \mathcal{L}^{A} on the test set as a function of the features space size. Δ_{add} represents the improvement from moving from V^{LT} to \hat{V} when $m = 2^{10}$. Δ_{features} represents the improvement of $\mathcal{L}^{\text{add}}(\hat{V}^{\text{LT}})$ when going from $m = 2^{10}$ to $m = 2^{16}$.

Qualitative insights. In the previous paragraph, we show that our method is improving \mathcal{L}^{add} . However, this is more of a sanity check as our method is designed to maximize \mathcal{L}^{add} . We next try to get some insights into how the proposed valuation is capturing some key elements of the problem. We discuss two important effects that have been discussed in the literature and recognized in practice.

The first one is the notion that the effect of a click decays over time. Consequently, it is less valuable to place an ad immediately after another one has been clicked. Several works have tackled this problem by making some assumptions on the underlying mechanism at play and hard-coding this decaying effect (Chapelle et al. 2014, Diemert et al. 2017, Zhang et al. 2014). On the contrary, we do not hard-code this relationship into our algorithm. Nevertheless, our valuation is able to



(a) Time elapsed since the last user’s click (b) Number of clicks already made by the user
Figure 7 Relative variation of the display valuation between \hat{V}^{FP} and V^{LT} .

learn this effect from data. Indeed, we show in Figure 7a the relative difference between \hat{V}^{FP} and V^{LT} as a function of the feature “hours since last click”. We observe that the last touch valuation tends to overestimate the value of an opportunity when a display was just clicked and our approach is able to correct for this effect. This is coherent with the live experiments done in Diemert et al. (2017).

Another effect that has been observed in practice is the diminishing marginal effect of ads. As a user is being exposed to different displays, the incremental benefit of showing an extra ad decreases. This is similar in spirit to the motivating example in the introduction. This is what we explore in Figure 7b where we show the relative difference between \hat{V}^{FP} and V^{LT} as a function of the feature “number of clicks before display”. We observe that, compared to the last touch benchmark, our framework drastically reduces the valuation of a display with the number of preceding clicks.

Performance. Finally, we also evaluate the performance of our approach with a procedure used in Zhang et al. (2014) for Multitouch Attribution models. This procedure consists in mapping the conversion probabilities of each display in a given user’s history to a single user’s conversion probability and was first proposed in Dalessandro et al. (2012) as a generative model. The proposed conversion probability of a history can be written with our notations as follows

$$\mathbb{P}(R_u = 1) = \left(1 - \prod_{t:u^t=u} (1 - V(\mathbf{x}^t)) \right) \cdot \delta^{|\mathcal{X}_u|} \quad (2)$$

Here, $\mathbb{P}(R_u = 1)$ denotes the probability that a user u converts. R_u is a binary reward, and for each display $\mathbf{x}^t \in \mathcal{X}_u$, we identify conversion probability of \mathbf{x}^t with its valuation $V(\mathbf{x}^t)$. Additionally, $|\mathcal{X}_u|$ represents the number of display interactions with user u . The first term in brackets represents the probability that at least one display leads to a conversion assuming zero interaction effects. The second term $\delta^{|\mathcal{X}_u|}$ accounts for the marginally decreasing effect of each ad. Zhang et al. (2014)

fixes δ to 0.95. We use the same value for consistency. This metric aims at measuring how well the conversion for each displays can recover the reward of a timeline. We can then compare two models by evaluating how well they rank the converted users according to Equation (2). A natural metric in this case is the mean average precision metric since it summarizes the Precision-Recall curve used in Zhang et al. (2014) with a single value. For that metric, our method provides an out-of-sample improvement of 20.5% over the last touch attribution approach. This shows that our approach strongly outperforms the last touch valuation on this independent performance metric as well.

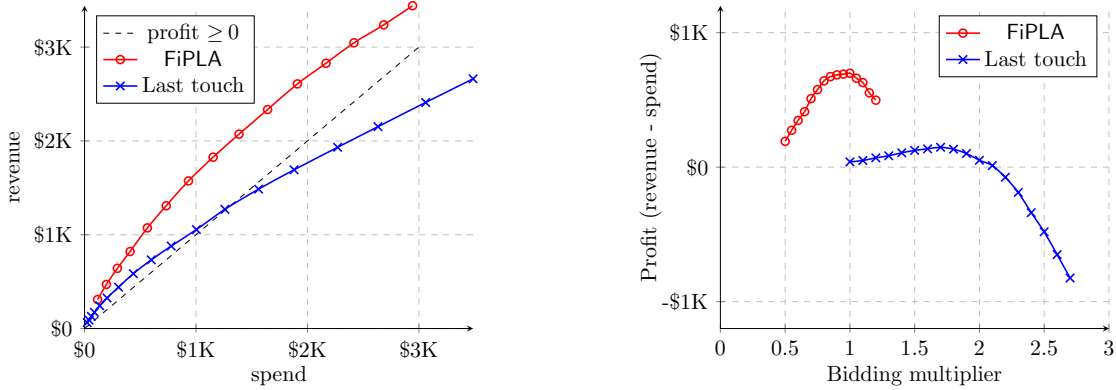
7. Impact on revenue

Properly addressing the label attribution problem helps better quantifying the value of a display opportunity. Ultimately, these values are used in downstream tasks such as bidding in a real-time bidding auction. In this section, we numerically explore the impact of the label attribution on the bidder’s profit. More precisely, we simulate a bidding environment and compare two bidding strategies. These two strategies differ in their solution to the label attribution problem and therefore in the valuation function they use. In particular, one uses the last touch label attribution approach whereas the other one uses the FiPLA Algorithm.

7.1. Experiment setup

Display types. We revisit the synthetic experiments with an additional bidding layer. In particular, we assume that there are two types of displays. Displays of type A lead to a conversion with a constant probability α whereas displays of type B never trigger a conversion. Similar to Setup 2 in Section 5, we associate a reward to a history if the user converts at least once, thus creating an effect of decreasing marginal returns when showing additional ads (of type A). We assume that the display type in each time step is sampled from a Bernoulli distribution of fixed parameter 0.5. As in Section 5, at each time step, there is a constant probability $\beta = 1/4$ that the user leaves the system.

Bidding environment. For each display opportunity, we assume that a second price auction takes place. More precisely, for each time step, we sample the highest bid b^{high} from a Beta distribution whose parameters depend on the display type. We assume that there is stronger competition for the displays of type A . More precisely, for display of type A the competition was generated with a Beta distribution of parameters (5,5), and for displays of type B , the competition was generated with a Beta distribution of parameters (1,9). The bidder computes a bid b using a given bidding strategy. If the b exceeds b^{high} , then the bidder wins the auction and pays b^{high} . On the other hand, if b is below b^{high} , then the bidder loses the auction. For each user’s timeline, the bidder generates a reward if the user converts at least once and incurs a cost for each won auction.



(a) Pareto curve for different multiplier values (b) Profit as a function of the bidding multiplier.
Figure 8 Revenue and spend of both the FiPLA and last touch strategies for different bidding multipliers.

Bidding strategies. We compare two strategies that differ on their label attribution solution and therefore on their estimated valuation function \hat{V} , where recall that the valuation function estimates the value of a display opportunity. This estimated valuation is either \hat{V}^{LT} , when the label attribution step is last touch, or \hat{V}^{FP} when the FiPLA Algorithm is used. We naturally refer to the two strategies as the *last touch* and FiPLA strategies. Both *last touch* and FiPLA Algorithms were trained on the data generated by a bidder that always bids 0.5 on 10,000 users. We then evaluate both strategies on 10,000 new users. In terms of bidding, as often in practice to control budget spend (Balseiro et al. 2021), we do not directly use the value given by V as a bid but allow the strategy to shade or modulate the value by a constant multiplier. For each strategy, we compute the revenue and spend for different values of bidding multiplier allowing us to generate a Pareto curve in the revenue-spend space.

7.2. Results

Figure 8 illustrates the results for different multiplier values. Looking at Figure 8a, we see that for both strategies, a higher spend leads to a higher revenue. Indeed, the multiplier is often adjusted in practice to match the spend with a budget. We observe however that the FiPLA strategy clearly dominates the last touch strategy. Indeed, for all spend values, the revenue generated by the FiPLA strategy strictly exceeds that of the last touch strategy. Moreover, the Pareto curve for the FiPLA strategy is always above the 45 degree line, indicating that, for all the multiplier values, the profit, which is the difference between the revenue and the spend, is always positive. On the other hand, the last touch strategy incurs a negative profit for some multiplier values.

Inspecting the profit as a function of the multiplier, we see on Figure 8b that for the FiPLA strategy, the higher profit is obtained when the multiplier is one, i.e., when the valuation \hat{V}^{FP} , is used directly as a bid. Note that this is not the case for the last touch approach, which needs to be carefully calibrated. Even allowing for the last touch strategy to calibrate a multiplier, we see,

as summarized in Table 3 that the profit generated by the FiPLA strategy is more than 3.5 times higher than the profit generated by the last touch strategy. This illustrates the importance of the label attribution problem and the benefits of properly addressing it.

	No multiplier	Best multiplier
Last touch strategy	\$39	\$148
FiPLA strategy	\$698	\$698

Table 3 Profit obtained by both strategies with and without bidding multipliers. Note that the FiPLA strategy does not benefit from using a bidding multiplier. This can be interpreted as a calibration property, that is not satisfied by the last touch strategy.

8. Conclusion

In this paper, we introduce the label attribution problem which consists of splitting the reward given to a bidder for each users into labels for each display opportunity. This is an important step needed to create a training dataset for most ML algorithm but it is often overlooked in practice. We formalize this problem and propose an approach to this problem which we call the robust label attribution. Our method is motivated by an intuitive additivity property and enjoys several theoretical structural properties. It is also practical and we show how it can be implemented at the scale required for display advertising.

References

- Anderl E, Becker I, Von Wangenheim F, Schumann JH (2016) Mapping the customer journey: Lessons learned from graph-based online attribution modeling. *International Journal of Research in Marketing* 33(3):457–474.
- Arava SK, Dong C, Yan Z, Pani A, et al. (2018) Deep neural net with attention for multi-channel multi-touch attribution. arXiv preprint arXiv:1809.02230.
- Archak N, Mirrokni VS, Muthukrishnan S (2010) Mining advertiser-specific user behavior using adfactors. *Proceedings of the 19th international conference on World wide web*, 31–40.
- Arjovsky M, Bottou L, Gulrajani I, Lopez-Paz D (2019) Invariant risk minimization. arXiv preprint arXiv:1907.02893.
- Balseiro S, Kim A, Mahdian M, Mirrokni V (2021) Budget-management strategies in repeated auctions. *Operations Research* .
- Balseiro SR, Candogan O (2017) Optimal contracts for intermediaries in online advertising. *Operations Research* 65(4):878–896.
- Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65(9):3952–3968.

- Berman R (2018) Beyond the last touch: Attribution in online advertising. *Marketing Science* 37(5):771–792.
- Blake T, Nosko C, Tadelis S (2015) Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica* 83(1):155–174.
- Bompaire M, Gilotte A, Heymann B (2020) Causal models for real time bidding with repeated user interactions.
- Chapelle O (2014) Modeling delayed feedback in display advertising. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 1097–1105.
- Chapelle O, Manavoglu E, Rosales R (2014) Simple and scalable response prediction for display advertising. *ACM Transactions on Intelligent Systems and Technology* 5(4).
- Choi H, Mela CF, Balseiro SR, Leary A (2020) Online display advertising markets: A literature review and future directions. *Information Systems Research* 31(2):556–575.
- Dalessandro B, Stitelman O, Perlich C, Provost F (2012) Causally motivated attribution for online advertising. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 1–9.
- Danaher PJ (2018) Delusion in Attribution : Caveats in Using Attribution for Multimedia Budget Allocation. *Journal of Marketing Research* 55(April):667–685.
- Dhamdhere K, Sundararajan M, Yan Q (2018) How important is a neuron? arXiv preprint arXiv:1805.12233.
- Diemert E, Galland P, Meynet J, Lefortier D (2017) Attribution modeling increases efficiency of bidding in display advertising. *2017 AdKDD and TargetAd - In conjunction with the 23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2017, 2* (ACM).
- Enberg J (2019) Global Digital Ad Spending 2019. URL <https://www.emarketer.com/content/global-digital-ad-spend>
- Flores-Szwagrzak K, Treibich R (2020) Teamwork and individual productivity. *Management Science* 66(6):2523–2544.
- Golrezaei N, Lobel I, Paes Leme R (2018) Auction design for roi-constrained buyers. *Available at SSRN 3124929* .
- Heymann B (2019) Cost per action constrained auctions. *Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation*, 1–8.
- Hunter DR, Lange K (2004) A tutorial on MM algorithms. *The American Statistician* 58(1):30–37.
- Ji W, Wang X (2017) Additional multi-touch attribution for online advertising. *31st AAAI Conference on Artificial Intelligence, AAAI 2017* 1360–1366.
- Ji W, Wang X, Zhang D (2016) A probabilistic multi-touch attribution model for online advertising. *International Conference on Information and Knowledge Management, Proceedings* 24-28-October-2016:1373–1382.

- Johnson GA, Lewis RA, Nubbemeyer EI (2017) Ghost Ads: Improving the economics of measuring online ad effectiveness. *Journal of Marketing Research* 54(6):867–884.
- Laub PJ, Taimre T, Pollett PK (2015) Hawkes Processes. Technical report.
- Lemon K (2016) Research Priorities 2016–2018. Technical report.
- Lewis RA, Mrao J, Reiley DH (2011) Here, there, and everywhere: Correlated online behaviors can lead to overestimates of the effects of advertising. *Proceedings of the 20th International Conference on World Wide Web, WWW 2011*, 157–166.
- Lewis RA, Rao JM (2015) The unfavorable economics of measuring the returns to advertising. *Quarterly Journal of Economics* 130(4):1941–1973.
- Perlich C, Dalessandro B, Hook R, Stitelman O, Raeder T, Provost F (2012) Bid optimizing and inventory scoring in targeted online advertising. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* .
- Peters J, Bühlmann P, Meinshausen N (2016) Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78(5):947–1012.
- Shao X, Li L (2011) Data-driven multi-touch attribution models. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 258–264.
- Singal R, Besbes O, Desir A, Goyal V, Iyengar G (2019) Shapley meets uniform: An axiomatic framework for attribution in online advertising. *The Web Conference 2019 - Proceedings of the World Wide Web Conference, WWW 2019* 1713–1723.
- Wang J, Zhang W, Yuan S, et al. (2017) Display advertising with real-time bidding (RTB) and behavioural targeting. *Foundations and Trends in Information Retrieval* 11(4-5):297–435.
- Xu J, Shao X, Ma J, Lee Kc, Qi H, Lu Q (2016) Lift-based bidding in ad selection. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Xu L, Duan JA, Whinston A (2014) Path to purchase: A mutually exciting point process model for online advertising and conversion. *Management Science* 60(6):1392–1412.
- Zhang Y, Wei Y, Ren J (2014) Multi-touch Attribution in Online Advertising with Survival Theory. *Proceedings - IEEE International Conference on Data Mining, ICDM 2015-January*(January):687–696.

Appendix

Appendix A: Theoretical grounding for additive valuation

In this section, we provide some theoretical justification for our approach. To be able to prove our results, we first present a slightly more abstract setting.

A.1. Model : from users to history

The goal of the paper is to present a data-driven approach to label attribution. As such, the model description in Section 3 is anchored on a dataset of users. Instead, we move to a slightly more abstract setting where we define a **history** h as a nonempty sequence of elements of \mathcal{X} that represents a sequence of displays. Recall that \mathcal{X} denotes the set of potential displays. We let \mathcal{H} be the set of all possible histories. Unlike the set of users, which we use in the main body of the paper and which is given by the data, we assume that the histories are drawn from some probability \mathbb{P} that denotes the propensity of each history to appear in the historical logs. Additionally, each history $h \in \mathcal{H}$ is associated with an expected reward $R(h) \in \mathbb{R}^+$. One can map a user u to a history h and the corresponding reward r_u to a particular realization of $R(h)$. However, different users with the same history might get different rewards and hence $R(h)$ captures an expected reward. For the rest of this section, we drop the users terminology.

A.2. Notation and assumptions

For any history h , we denote by $|h|$ its size and use the notation $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}]$, where $\mathbf{x}_i \in \mathcal{X}$ for $i \in \{1, \dots, |h|\}$, to denote its individual elements. Furthermore, we write $h \preceq h'$ if h is a sub-history of h' , that is if h is equal to the sequence of the $|h|$ first displays of h' .

For any history h , we let $h_{[1, \dots, i]}$ be the projection of h onto its first i elements, i.e., for all $i \leq |h|$, we have $h_{[1, \dots, i]} = [\mathbf{x}_1, \dots, \mathbf{x}_i]$. With these notations, $h' \succ h$ is equivalent to $|h'| \geq |h|$ and $h'_{[1, \dots, |h|]} = h$. Finally, for all $\mathbf{x} \in \mathcal{X}$ and $h \in \mathcal{H}$, let $h + [\mathbf{x}] = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}, \mathbf{x}]$ denote the history obtained from adding a display \mathbf{x} to a history h . We end this section with two technical assumptions.

ASSUMPTION 1. \mathcal{X} is a finite set.

Assumption 1 is not limiting in practice. Indeed, if some display features are continuous rather than categorical, they can be quantized (Chapelle et al. 2014) in order to satisfy Assumption 1. We also make the following assumption on the underlying probability distribution of histories.

ASSUMPTION 2. If \mathbb{P} has support on a history, then it also has support on all its sub-histories. Namely, if $\mathbb{P}(h) > 0$ for some $h \in \mathcal{H}$, then $\mathbb{P}(h') > 0$ for all $h' \preceq h$.

This is a mild assumption which is verified for various generative models used in the literature such as Markovian models (Archak et al. 2010, Anderl et al. 2016, Singal et al. 2019) or point process models (Xu et al. 2014).

A.3. Valuation and perfect valuation

Using the history notation, for every $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$, we let $V(\mathbf{x}|h)$ denotes the expected value of the display \mathbf{x} given the history h . Note that in practice, making this dependence is not critical as \mathbf{x} often include that information through features such as number of ads seen or number of ads clicked. However, we make this dependence explicit for the purpose of making our theoretical presentation more rigorous and precise. We also revisit the definition of label attribution as follows.

DEFINITION 3 (LABEL ATTRIBUTION). A **label attribution** μ is a mapping from $\mathbb{N} \times \mathcal{H}$ to \mathbb{R}^+ that satisfies the following properties.

1. $\forall h \in \mathcal{H}, \forall i > |h|, \mu(i|h) = 0$.
2. $\forall h \in \mathcal{H}, R(h) = \sum_{i=1}^{|h|} \mu(i, h)$.

For a label attribution μ and history $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}]$, $\mu(i, h)$ represents the label of the i^{th} display \mathbf{x}_i , i.e., the part of the reward $R(h)$ attributed to \mathbf{x}_i .

For a given observed history h and display opportunity \mathbf{x} , recall that $H \succeq h + [\mathbf{x}]$ means that H starts with the sequence of displays of $h + [\mathbf{x}]$, where $h + [\mathbf{x}]$ is the history obtained from adding a display \mathbf{x} to a history h . We next define the ideal value learned by a perfect algorithm.

DEFINITION 4 (ASSOCIATED VALUATION). For every label attribution μ and probability distribution \mathbb{P} , the associated valuation $V_{\mathbb{P}}^{\mu}(\mathbf{x}|h)$ is defined, for all $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$, as

$$V_{\mathbb{P}}^{\mu}(\mathbf{x}|h) = \mathbb{E}_{H \sim \mathbb{P}}[\mu(|h| + 1, H) \mid H \succcurlyeq h + [\mathbf{x}]]. \quad (3)$$

Importantly, note that $V_{\mathbb{P}}^{\mu}$ depends on the label attribution choice μ as well as the distribution \mathbb{P} . Another way to interpret the definition of the associated valuation is that when the bidder observes a display opportunity \mathbf{x} after a history h has already happened, the bidder does not know whether additional opportunities will come in the future. If no other opportunity arises, then the display should be valued $\mu(|h| + 1, h + [\mathbf{x}])$. On the other hand, if more displays are added, then the display should be valued $\mu(|h| + 1, H)$ for some $H \succcurlyeq h + [\mathbf{x}]$.

For instance, the valuation associated with the last touch label attribution is given by

$$V_{\mathbb{P}}^{\mu^{\text{LT}}}(\mathbf{x}|h) = R(h + [\mathbf{x}]) \cdot \mathbb{P}(H = h + [\mathbf{x}] \mid H \succcurlyeq h + [\mathbf{x}]).$$

If the label attribution is driven by the last touch philosophy, then an opportunity is evaluated through its probability of exactly preceding a conversion, which is precisely what the quantity $\mathbb{P}(H = h + [\mathbf{x}] \mid H \succcurlyeq h + [\mathbf{x}])$ is capturing. It is worth noticing that this valuation depends on the distribution of future displays. This is precisely the caveat observed in the Introduction as this dependence might lead the bidder to overvalue displays that come late in the conversion funnel.

A.4. Additive valuation and fixed point characterization

We can now properly define the additive valuation V^{FP} as follows. For all $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$,

$$V^{\text{FP}}(\mathbf{x}|h) = R(h + [\mathbf{x}]) - R(h). \quad (4)$$

Note that in the paper, additivity is a property one can strive for. However, we cannot properly define an additive valuation per se since the set of users, or realized history, may not be complete. On the other hand,

in this abstract setting where we have a reward function $R(h)$ for all history $h \in \mathcal{H}$, we can formally define the additive valuation V^{FP} . In turn, we can define the additive label attribution μ^{add} through a fixed point equation. In particular, for all $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}] \in \mathcal{H}$ and $i \leq |h|$,

$$\mu^{\text{FP}}(i, h) = \frac{V^{\text{FP}}(\mathbf{x}_i | h_{[1, \dots, i-1]})}{\sum_{j=1}^{|h|} V^{\text{FP}}(\mathbf{x}_j | h_{[1, \dots, j-1]})} \cdot R(h). \quad (5)$$

Note the obvious parallel with Equation (1). We are now ready to formalize the relation between additivity and proportional label attribution through the fixed point equation.

PROPOSITION 1. $V_{\mathbb{P}}^{\mu} = V^{\text{FP}} \iff \mu = \mu^{\text{FP}}$

In words, Proposition 1 states that a label attribution leads to the additive valuation, i.e., the additive valuation is associated to the additive label attribution, if and only if the label attribution and valuation satisfy the fixed point equation (8).

A.5. Convergence of the algorithm

We now formally justify the design of our algorithm. In particular, consider the following iterative procedure:

$$V^{(k+1)}(\mathbf{x}|h) = V_{\mathbb{P}}^{\mu^{(k)}}(\mathbf{x}|h), \quad \forall (\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}, \quad (6)$$

where for all $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}]$ and $i \leq |h|$,

$$\mu^{(k)}(i, h) = \frac{V^{(k)}(\mathbf{x}_i | h_{[1, \dots, i-1]})}{\sum_{j=1}^{|h|} V^{(k)}(\mathbf{x}_j | h_{[1, \dots, j-1]})} \cdot R(h). \quad (7)$$

Recall that for any μ and $h \in \mathcal{H}$, the valuation $V_{\mathbb{P}}^{\mu}(h)$, defined in Equation (3), is the expected reward, where the expectation is conditioned on $H \succcurlyeq h$. Note that Equation (6) corresponds to the prediction step of FiPLA Algorithm in our ideal setting where one can learn perfectly. We also initialize the procedure with some label attribution $\mu^{(0)}$.

PROPOSITION 2. *For any $\mu^{(0)}$ such that $V_{\mathbb{P}}^{\mu^{(0)}}(\mathbf{x}|h) > 0$ for all $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$, the procedure defined in Equations (6)-(7) converges. Moreover, $\mu^{(k)} \rightarrow \mu^{\text{FP}}$ when $k \rightarrow \infty$.*

Note that the technical assumption $V_{\mathbb{P}}^{\mu^{(0)}}(\mathbf{x}|h) > 0$ on the initial valuation is not required to prove convergence but only to prove that the iterates converge to the robust label attribution. Indeed, without this assumption, some coordinates can get stuck at their initial value. Proposition 2 justifies the convergence of our proposed FiPLA Algorithm in this abstract setting.

Remark. Consider a setting where we only have one sample with history $h = (\mathbf{x}_1, \dots, \mathbf{x}_I)$ but do not have data for any subhistory $h = (\mathbf{x}_1, \dots, \mathbf{x}_i)$, $\forall i < I$. In this case, Assumption 2 is does not hold. However, for practical purposes, having access to a single sequence of displays is one setting where we do not have enough data to know what to attribute. Moreover, in such case, the FiPLA Algorithm would be stuck at the initial guess for μ^0 . This means that the algorithm would revert to the heuristic used for initialization, such as any rule-based heuristic like last touch or uniform attribution. This looks like a reasonable fallback: when there is not enough data to learn anything, we do not learn anything beyond the initialization.

A.6. Additional properties

We conclude this section, by showing that the additive valuation has several notable properties. We begin by proving that bidding according to the additive valuation is the only myopic optimal bidding strategy and then show that the additive valuation is the only valuation that does not depend on the probability distribution \mathbb{P} making it very robust from a learning perspective.

A.6.1. Myopic optimality. We first show that bidding according to the robust valuation is myopic optimal, i.e. optimal if the sequence of displays terminates without any future opportunities.

PROPOSITION 3. *V satisfies the additivity property **if and only if** it is myopic optimal, i.e. if in a second-price setting, bidding V is weakly dominant for any competition profile assuming there will be no other future opportunities.*

We acknowledge that this notion of optimality is – as the name suggests – rather narrow: the potential combined effects of the decision at stake (the display to be bought) and the future ones are neglected. Nevertheless, we believe this notion has some merit. In particular, the standard notion of optimality implies that the best decision at a given time depends on what will be done in the future. However, solving the complete optimization problem is in practice often intractable and therefore, one can only hope for approximate solutions. By contrast, as we will discuss next, our proposed strategy does not depend on the probability distribution \mathbb{P} (and hence, on the impact of our bidding strategy in the future), making it more robust. Finally, our aim is also to shift the discussion to principled approaches to label attribution, very much in the same spirit than Singal et al. (2019). Here, we prove that there is a unique valuation which is myopic optimal.

A.6.2. Distributional robustness. We next show that the robust valuation does not depend on the underlying probability distribution \mathbb{P} , i.e., the propensity of each history to appear in the historical logs.

PROPOSITION 4. *V satisfies the additivity property **if and only if** V is an associated valuation which does not depend on \mathbb{P} but only on the reward function R , i.e., there exists a label attribution μ such that $V = V_{\mathbb{P}}^{\mu} = V_{\mathbb{Q}}^{\mu}$ for any other probability distribution \mathbb{Q} over \mathcal{H} that has the same support as \mathbb{P} .*

Surprisingly, the additivity property also leads to distributional robustness. From a learning perspective, this means that this approach alleviates the need of learning the underlying probability distribution \mathbb{P} . In particular, if the environment changes and we receive a new probability distribution on the histories \mathbb{Q} , then the expected reward can still be computed using the formula

$$\mathbb{E}_{H \sim \mathbb{Q}} R(H) = \mathbb{E}_{H \sim \mathbb{Q}} \sum_{i=1}^{|H|} V_{\mathbb{P}}^{\mu}(H_i \mid H_{[1, \dots, i-1]}),$$

even when $\mathbb{P} \neq \mathbb{Q}$. We emphasize that this is not true for any other valuation mechanism. Finally, distributional invariance is often related to causality (Peters et al. 2016, Arjovsky et al. 2019), and it would be an interesting research direction to further explore these connections.

Appendix B: Proofs

B.1. Proof of Proposition 1

We first show that the valuation associated to the additive label attribution is the additive valuation. First, by inspecting Equation (8), it is clear that μ^{add} satisfies Definition 3. Now, for all $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$, we have, using Equation (3),

$$\begin{aligned} V_{\mathbb{P}}^{\mu^{\text{FP}}}(\mathbf{x} | h) &= \mathbb{E}_{H \sim \mathbb{P}} \left[\frac{V^{\text{FP}}(\mathbf{x} | H_{[1, \dots, |h|]})}{\sum_{j=1}^{|h|} V^{\text{FP}}(H_j | H_{[1, \dots, j-1]})} \cdot R(H) \mid H \succcurlyeq h + [\mathbf{x}] \right] \\ &= V^{\text{FP}}(\mathbf{x} | h) \cdot \mathbb{E}_{H \sim \mathbb{P}} \left[\frac{R(H)}{\sum_{j=1}^{|h|} V^{\text{FP}}(H_j | H_{[1, \dots, j-1]})} \mid H \succcurlyeq h \right] \\ &= V^{\text{FP}}(\mathbf{x} | h), \end{aligned}$$

where the last equality follows Equation 4.

We next show the reverse implication. We assume we have a label attribution μ and valuation V such that the fixed point equation is verified, i.e., for all $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}]$ and $i \leq |h|$,

$$\mu(i, h) = \frac{V(h_i | h_{[1, \dots, i-1]})}{\sum_{j=1}^{|h|} V(\mathbf{x}_j | h_{[1, \dots, j-1]})} \cdot R(h).$$

Using Definition 4, we have for all $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$,

$$V_{\mathbb{P}}^{\mu}(\mathbf{x} | h) = \mathbb{E}_{H \sim \mathbb{P}} [\mu(|h| + 1, H) \mid H \succcurlyeq h + [\mathbf{x}]].$$

Consequently, for all for all $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}]$ and $i \leq |h|$, we obtain

$$\begin{aligned} \sum_{j=1}^{|h|} V(\mathbf{x}_j | h_{[1, \dots, j-1]}) &= \sum_{j=1}^{|h|} \mathbb{E}_{H \sim \mathbb{P}} [\mu(j, H) \mid H \succcurlyeq h_{[1, \dots, j]}] \\ &= \sum_{i=1}^{|h|} \mathbb{E}_{H \sim \mathbb{P}} \left[\frac{V(\mathbf{x}_i | h_{[1, \dots, i-1]})}{\sum_{j=1}^{|h|} V(\mathbf{x}_j | h_{[1, \dots, j-1]})} \cdot R(h) \mid H \succcurlyeq h_{[1, \dots, i]} \right] \\ &= R(h), \end{aligned}$$

where we obtain the last equality by interchanging the expectation and sum operators. This concludes the proof.

B.2. Proof of Proposition 2

To prove the convergence of the iterative procedure defined in Equations (6)-(7), we show that it corresponds to the iterates of a majorize-minorize (MM) algorithm. We remind the reader of the MM philosophy and refer to Hunter and Lange (2004) for more details. An MM algorithm aims at finding the maximizer θ^* of a function $f(\theta)$. For that purpose, a surrogate function $g(\theta | \hat{\theta})$ that depends on a current estimate $\hat{\theta}$ and is typically easier to maximize is identified and maximized instead. The key is to choose a surrogate function that minorizes the objective function. A typical iteration of the MM algorithm can be written as follows. Given an estimate $\theta^{(m)}$, we let

$$\theta^{(m+1)} = \arg \max_{\theta} g(\theta | \theta^{(m)}).$$

The procedure is guaranteed to converge to a local maximum of f as long as the following are satisfied:

1. $f(\theta)$ is concave,
2. $g(\theta^{(m)}|\theta^{(m)}) = f(\theta^{(m)})$,
3. $g(\theta|\theta^{(m)}) \leq f(\theta), \forall \theta$.

We show that our iterative procedure is a special case of an MM algorithm by letting for all $V : \mathcal{X} \times \mathcal{H} \mapsto \mathbb{R}^+$,

$$f(V) = \mathbb{E}_{H \sim \mathbb{P}} \left[R(H) \cdot \ln \left(\sum_{q \in \{1, \dots, |H|\}} V(H_q | H_{[1, \dots, q-1]}) \right) - \sum_{q \in \{1, \dots, |H|\}} V(H_q | H_{[1, \dots, q-1]}) \right],$$

$$g(V|\hat{V}) = \mathbb{E}_{H \sim \mathbb{P}} \left[\sum_{j \in \{1, \dots, |H|\}} \left\{ \frac{\hat{V}(H_j | H_{[1, \dots, j-1]}) \cdot R(H)}{\sum_{q \in \{1, \dots, |H|\}} \hat{V}(H_q | H_{[1, \dots, q-1]})} \cdot \ln \left(\frac{V(H_j | H_{[1, \dots, j-1]})}{\hat{V}(H_j | H_{[1, \dots, j-1]})} \sum_{q \in \{1, \dots, |H|\}} \hat{V}(H_q | H_{[1, \dots, q-1]}) \right) \right. \right. \\ \left. \left. - V(H_j | H_{[1, \dots, j-1]}) \right\} \right].$$

Since the space of histories is finite, V can always be mapped to a finite dimension vector. We do however keep the function notation for simplicity. First, note that f is strictly concave as a linear combination of concave functions. Second, observe that for all V ,

$$g(V|V) = \mathbb{E}_{H \sim \mathbb{P}} \left[\sum_{j \in \{1, \dots, |H|\}} \left\{ \frac{V(H_j | H_{[1, \dots, j-1]}) \cdot R(H)}{\sum_{q \in \{1, \dots, |H|\}} V(H_q | H_{[1, \dots, q-1]})} \cdot \ln \left(\frac{V(H_j | H_{[1, \dots, j-1]})}{V(H_j | H_{[1, \dots, j-1]})} \sum_{q \in \{1, \dots, |H|\}} V(H_q | H_{[1, \dots, q-1]}) \right) \right. \right. \\ \left. \left. - V(H_j | H_{[1, \dots, j-1]}) \right\} \right] \\ = \mathbb{E}_{H \sim \mathbb{P}} \left[R(H) \cdot \ln \left(\sum_{q \in \{1, \dots, |H|\}} V(H_q | H_{[1, \dots, q-1]}) \right) \cdot \sum_{j \in \{1, \dots, |H|\}} \frac{V(H_j | H_{[1, \dots, j-1]})}{\sum_{q \in \{1, \dots, |H|\}} V(H_q | H_{[1, \dots, q-1]})} \right. \\ \left. - \sum_{j \in \{1, \dots, |H|\}} V(H_j | H_{[1, \dots, j-1]}) \right] \\ = f(V).$$

Third, the inequality $f(V) \geq g(V|\hat{V})$ follows by using the concavity of the logarithm and Jensen inequality on the random variable X which take value

$$\frac{V(H_j | H_{[1, \dots, j-1]})}{\hat{V}(H_j | H_{[1, \dots, j-1]})} \cdot \sum_{q \in \{1, \dots, |H|\}} \hat{V}(H_q | H_{[1, \dots, q-1]})$$

with probability $\hat{V}(H_j | H_{[1, \dots, j-1]}) / \sum_{q \in \{1, \dots, |H|\}} \hat{V}(H_q | H_{[1, \dots, q-1]})$. Finally, for any \hat{V} , we have $\arg \max_V g(V|\hat{V}) = V_{\mathbb{P}}^{\mu_{\hat{V}}}$ by Lemma 1 below, where for all $\hat{V} : \mathcal{H} \mapsto \mathbb{R}$, we define $\mu_{\hat{V}}$ as follows. For all $h = [\mathbf{x}_1, \dots, \mathbf{x}_{|h|}] \in \mathcal{H}$ and $i \leq |h|$,

$$\mu_{\hat{V}}(i, h) = \frac{\hat{V}(\mathbf{x}_i | h_{[1, \dots, i-1]})}{\sum_{j=1}^{|h|} \hat{V}(\mathbf{x}_j | h_{[1, \dots, j-1]})} \cdot R(h). \quad (8)$$

Using Hunter and Lange (2004), we conclude that the sequence of $V^{(k)}$ corresponds to the iterates of an MM algorithm with strictly concave criteria, and hence converges. This concludes the proof of Proposition 2.

LEMMA 1. For any $\hat{V} : \mathcal{H} \mapsto \mathbb{R}$, $\arg \max_V g(V|\hat{V}) = V_{\mathbb{P}}^{\mu_{\hat{V}}}$.

Proof. Since $g(V|\hat{V})$ is concave in V , it suffices to show that $V_{\mathbb{P}}^{\mu\hat{V}}$ satisfies the first order condition. In particular, for any $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$,

$$\frac{\partial g(V|\hat{V})}{\partial V(\mathbf{x}|h)} = \frac{\partial}{\partial V(\mathbf{x}|h)} \left(\sum_t \mathbb{P}(t) \cdot \sum_{j \in \{1, \dots, |t|\}} \left\{ \frac{\hat{V}(t_j|t_{[1, \dots, j-1]}) \cdot R(t)}{\sum_{q \in \{1, \dots, |t|\}} \hat{V}(t_q|t_{[1, \dots, q-1]})} \cdot \ln \left(\frac{V(t_j|t_{[1, \dots, j-1]})}{\hat{V}(t_j|t_{[1, \dots, j-1]})} \sum_{q \in \{1, \dots, |t|\}} \hat{V}(t_q|t_{[1, \dots, q-1]}) \right) - V(t_j|t_{[1, \dots, j-1]}) \right\} \right)$$

Note that h appears in the inner sum if and only if $t \succcurlyeq h + [\mathbf{x}]$. Therefore, we can write

$$\frac{\partial g(V|\hat{V})}{\partial V(\mathbf{x}|h)} = \frac{\partial}{\partial V(\mathbf{x}|h)} \left(\sum_{t \succcurlyeq h + [\mathbf{x}]} \mathbb{P}(t) \cdot \left\{ \frac{\hat{V}(\mathbf{x}|h) \cdot R(t)}{\sum_{q \in \{1, \dots, |t|\}} \hat{V}(t_q|t_{[1, \dots, q-1]})} \cdot \ln \left(\frac{V(\mathbf{x}|h)}{\hat{V}(\mathbf{x}|h)} \sum_{q \in \{1, \dots, |t|\}} \hat{V}(t_q|t_{[1, \dots, q-1]}) \right) - V(\mathbf{x}|h) \right\} \right).$$

Exchanging the derivative and sum operators yields

$$\frac{\partial g(V|\hat{V})}{\partial V(\mathbf{x}|h)} = \sum_{t \succcurlyeq h + [\mathbf{x}]} \mathbb{P}(t) \cdot \left\{ \frac{\hat{V}(\mathbf{x}|h) \cdot R(t)}{\sum_{q \in \{1, \dots, |t|\}} \hat{V}(t_q|t_{[1, \dots, q-1]})} \cdot \frac{1}{V(\mathbf{x}|h)} - 1 \right\}.$$

Setting the above to zero for the first order condition then implies that for all $(\mathbf{x}, h) \in \mathcal{X} \times \mathcal{H}$, we have

$$\begin{aligned} V(\mathbf{x}|h) &= \frac{1}{\sum_{t \succcurlyeq h + [\mathbf{x}]} \mathbb{P}(t)} \cdot \left(\sum_{t \succcurlyeq h + [\mathbf{x}]} \mathbb{P}(t) \cdot \frac{\hat{V}(\mathbf{x}|h) \cdot R(t)}{\sum_{q \in \{1, \dots, |t|\}} \hat{V}(t_q|t_{[1, \dots, q-1]})} \right) \\ &= \mathbb{E}_{H \sim \mathbb{P}} \left[\frac{\hat{V}(\mathbf{x}|h)}{\sum_{q \in \{1, \dots, |H|\}} \hat{V}(H_q|H_{[1, \dots, q-1]})} \cdot R(H) \mid H \succcurlyeq h + [\mathbf{x}] \right] \\ &= V_{\mathbb{P}}^{\mu\hat{V}}(\mathbf{x}|h). \end{aligned}$$

This concludes the proof. \square

B.3. Proof of Proposition 3

Fix a history $h \in \mathcal{H}$ and display $\mathbf{x} \in \mathcal{X}$. Recall that bidding $V(\mathbf{x}|h)$ is myopic weakly dominant if and only if, in a second price setting, bidding $V(\mathbf{x}|h)$ is weakly dominant optimal for any competition profile. For a given competition profile, assume that the highest bid t among competition follows a density g . The payoff of the auction is given by $(R(h + [\mathbf{x}]) - t)$ if the auction is won, i.e., if the bid $V(\mathbf{x}|h)$ exceeds t , and $R(h)$ otherwise. The optimal bid $V(\mathbf{x}|h)$ must therefore maximize the following

$$\int_0^{V(\mathbf{x}|h)} (R(h + [\mathbf{x}]) - t) \cdot g(t) \cdot dt + \int_{V(\mathbf{x}|h)}^{+\infty} R(h) \cdot g(t) \cdot dt,$$

where the first integral corresponds to the set of competition prices where we win the auction and the second term correspond to competition bids above V . In the latter case, we obtain $R(h)$ because of the missed opportunity. Writing the first order condition yields

$$(R(h + [\mathbf{x}]) - V(\mathbf{x}|h) - R(h)) \cdot g(V(\mathbf{x}|h)) = 0$$

for any g , hence, bidding $V(\mathbf{x}|h)$ is myopic weakly dominant if and only if $R(h + [\mathbf{x}]) - V(\mathbf{x}|h) - R(h) = 0$.

B.4. Proof of Proposition 4

By Proposition 1, if V is myopic optimal, then it is associated with the additive attribution given in Equation (4), which does not depend on \mathbb{P} . We therefore only have to prove the reverse implication which we do next.

Assume that V does not depend on \mathbb{P} but only on the reward function R , i.e., there exists an internal attribution μ such that $V = V_{\mathbb{P}}^{\mu} = V_{\mathbb{Q}}^{\mu}$ for any other probability distribution \mathbb{Q} over \mathcal{H} . We begin by showing that there exists a unique valuation which is distributionally robust, i.e., if we let μ its associated attribution, we must have for any probability \mathbb{P} and $h \in \mathcal{H}$

$$V(h) = V_{\mathbb{P}}^{\mu}(h) = \mathbb{E}_{H \sim \mathbb{P}}[\mu(|h|, H) | H \succcurlyeq h].$$

The associated attribution defined by Equation (4) is distributionally robust since the expression of μ does not depend on \mathbb{P} which shows existence. We show unicity by contradiction. Assume that there are two such internal attributions μ^1 and μ^2 . For all \mathbb{P} , h and $j \leq |h|$, we must have

$$V(h_j | h_{[1, \dots, j-1]}) = \mathbb{E}_{H \sim \mathbb{P}}[\mu^1(j, H) | H \succcurlyeq h_{[1, \dots, j]}] = \mathbb{E}_{H \sim \mathbb{P}}[\mu^2(j, H) | H \succcurlyeq h_{[1, \dots, j]}].$$

Consider a probability \mathbb{Q} such that the weight of all elements $H \succcurlyeq h_{[1, \dots, j]}$ is moved to $h_{[1, \dots, j]}$. Then, we have for all h and $j \leq |h|$, $\mu^1(j, h) = \mu^2(j, h)$. Note that since the probability \mathbb{Q} must satisfy Assumption 2, a limiting argument yields the result.