Shapley Meets Uniform: An Axiomatic Framework for Attribution in Online Advertising

Raghav Singal
Department of IE&OR, Columbia University, rs3566@columbia.edu

Omar Besbes
Graduate School of Business, Columbia University, ob2105@gsb.columbia.edu

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

Vineet Goyal, Garud Iyengar
Department of IE&OR, Columbia University, vg2277@columbia.edu, garud@ieor.columbia.edu

May 22, 2019

One of the central challenges in online advertising is attribution, namely, assessing the contribution of individual advertiser actions including e-mails, display ads and search ads to eventual conversion. Several heuristics are used for attribution in practice; however, there is no formal justification for them and many of these fail even in simple canonical settings. The main contribution in this work is to develop an axiomatic framework for attribution in online advertising. In particular, we consider a Markovian model for the user journey through the conversion funnel, in which ad actions may have disparate impacts at different stages. We propose a novel attribution metric, that we refer to as *counterfactual adjusted Shapley value*, which inherits the desirable properties of the traditional Shapley value while overcoming its shortcomings in the context of our application. Furthermore, we establish that this metric coincides with an adjusted "unique-uniform" attribution scheme. This scheme is efficiently implementable and can be interpreted as a correction to the commonly used uniform attribution scheme. We supplement our theoretical developments with numerical experiments inspired by a real-world large-scale dataset.

Key words: Digital economy, online advertising, attribution, Markov chain, Shapley value, causality

1. Introduction

With the rise of the Internet, the digital economy has become a trillion dollar industry accounting for over 6% of the U.S. GDP (Hagan 2018). A multitude of retailers reach consumers through online sources with the goal of acquiring new customers and managing relationships with the existing ones. Such wide range of economic activities taking place in the digital world has enabled data collection at a massive scale, allowing retailers to better understand customer behavior and enhance service quality using data-driven decisions. With over a billion people with access to the Internet, it is not surprising that advertising has moved to the digital space. The global digital marketing sector witnessed a growth of 21% in 2017, which increased its market size to USD 88 billion (Sluis

2018). Key decisions in this space include allocating budget to various ad channels and media (e-mail, display media platforms, and paid search for instance) and optimizing tactical decisions in each channel. Such tactical decisions may be driven by the ad publishers; see, e.g., Balseiro et al. (2014), Hojjat et al. (2017) and Lejeune and Turner (2019) or by the advertisers themselves. For example, in display or search advertising, this would entail bidding for ads to push one's product towards the desired customer demographic at the right time; see, e.g., Iyer et al. (2014), Balseiro et al. (2015), Balseiro and Gur (2017) and Baardman et al. (2019). Furthermore, due to the digital nature of such online ad exchanges, advertisers can access user-level information before placing a bid, which motivates the importance of understanding the *state-specific* value of an ad. All such tactical decisions require a deep understanding of the value of showing an ad via a specific channel at a given time. How can an advertiser assign or *attribute* value to the advertising actions taken across the different channels and media? *Attribution* is one of the central questions in online advertising. The value of each channel is an important input to media mix optimization, helps build an understanding of the customer journey, and also helps a company justify its marketing spend (United 2012, Priest 2017). Incorrect understanding of the effectiveness of online channels can result in highly sub-optimal budget allocation, possibly resulting in significant lost revenue (Kireyev et al. 2016). Recognizing the importance of the attribution problem, the Marketing Science Institute identifies attribution as the top most research priority for the period 2016-2018 (Institute 2016). Though multiple heuristics have been proposed and studied in this domain, there is lack of a systematic approach that has both theoretical foundations and is tractable.

Attribution is inextricably linked to causality since it involves quantifying the added value of showing an ad over the baseline value of what *would have* happened if no ad was shown (*counterfactual*). We consider the following view of causality, which comes from the pioneering work of Rubin (1974):

> *"Intuitively, the causal effect of one treatment, E, over another, C, for a particular unit and an interval of time from t_1 to t_2 is the difference between what would have happened at time t_2 if the unit had been exposed to E initiated at t_1 and what would have happened at t_2 if the unit had been exposed to C initiated at t_1."*

In the context of digital advertising, treatments E and C are the advertiser showing, and not showing an ad, respectively. Accordingly, attribution involves capturing the causal effect of an ad where the baseline corresponds to not showing an ad.

Causality is an active research area, and there exist paradigms for capturing the causal effect of a treatment. See, e.g. Pearl (2009), Halpern and Pearl (2005), Chockler and Halpern (2004), Hitchcock (1997), Morgan and Winship (2014), Collins et al. (2004), Eells (1991), Hume (2003) and Rubin (1974). We comment on some of these alternative approaches in our concluding remarks.

In addition to these philosophical works in causality, there also exist data-driven works such as Bottou et al. (2013) that, via importance sampling in a Bayesian network, allow one to estimate counterfactuals without having to collect additional data. In particular, they get a handle on the "what would have happened" scenario using the data that one has *already* collected. As we will see in Section 5.2, we leverage this technique to construct data efficient algorithms.

As with the existing work in attribution, we also consider the setting where an advertiser is interested in understanding the contributions of various ads to a single product they are promoting. Even with a single product, attribution is a challenging problem. It involves distributing the value generated by a network of actions to each individual action. Such a network might have a mix of interaction effects that one needs to account for when decomposing the network value. Capturing such (possibly non-linear) interaction effects is a fundamental challenge.

1.1. Related literature

The attribution problem has been studied widely and in diverse areas. We present a brief overview of existing approaches and refer the reader to existing surveys (Choi et al. 2017, Kannan et al. 2016) for a more complete treatment. Attribution methodologies can be classified into two broad classes: *rule-based* or *algorithmic*. We discuss both the classifications below.

1.1.1. Rule-based heuristics. Rule-based heuristics include approaches such as *last touch attribution* (LTA), *uniform weights*, and *customized weights* (Arensman and Yeung 2016, Priest 2017, Quantcast 2013, 2016). In LTA, an advertiser allocates all the value generated by a user to the last ad directed at her whereas under a uniform weights scheme, all the ads on a conversion path are allocated an equal credit. Ads receive tailored weights under a custom weights scheme. Although such heuristics are transparent and tractable, there is no rationale justifying their appropriateness as a measure for attribution. Metrics such as LTA can be unfair since they do not value the contribution of channels that build product awareness. Uniform or customized weights might appear to be a fix but there is no a priori reason to believe that attribution should be linear.

1.1.2. Algorithmic approaches. The algorithmic approaches can be classified as using either *incremental value heuristic (IVH)* (or *removal effect*) or *Shapley value (SV)* as a measure for attribution.

IVH. IVH computes the change in the eventual *conversion*[1] probability of a user when a specific ad is removed from her path. This is the most common metric for attribution (Abhishek et al. 2012, Anderl et al. 2016, Arava et al. 2018, Archak et al. 2010, Danaher and van Heerde 2018, Kakalejčík et al. 2018, Li and Kannan 2014, Li et al. 2017). In this approach, one calibrates a model that

[1] Conversion refers to the event in which a user buys the underlying product.

predicts the conversion probability as a function of the ad exposure and then, uses the estimated model to compute the incremental value of each ad. The novelty comes from the model proposed to describe user behavior, e.g., an HMM (Abhishek et al. 2012) or a neural network (Arava et al. 2018). Irrespective of the prediction model's level of complexity, attribution is measured via IVH. There exists little (if any) formal justification for why IVH serves as a good approximation to attribution. In fact, in Section 3.2, we show that IVH can result in inappropriate allocations.

SV. SV (Shapley 1953) is a well-accepted concept for assigning credit to individual players in a cooperative game. The value generated by online advertising can be viewed as the outcome of a cooperative effect of the channels and media platforms. Dalessandro et al. (2012) pose attribution as a causal estimation problem and propose SV as an approximation scheme for the causally motivated problem. They also show that SV generalizes the probabilistic model of Shao and Li (2011). Using a stylized model, Berman (2018) shows the use of SV for attribution can be beneficial to the advertiser.

Other works. Attribution has been tackled from various other angles. Jordan et al. (2011) use a Markovian model to motivate a payment scheme that satisfies incentive compatibility for the advertiser and is fair from the publisher's point-of-view. Xu et al. (2014) propose a mutually exciting point process to capture dynamic interactions among various ads. Zhang et al. (2014), Ji et al. (2016), and Ji and Wang (2017) provide an interesting view of attribution via the lens of survival theory. Under a stylized setting, Abhishek et al. (2017) analyze attribution contracts used by an advertiser to incentivize two publishers that affect customer acquisition. Zhao et al. (2018) propose a regression-based relative importance method to compute the marginal contributions. From an empirical perspective, Blake et al. (2015) measure the return on investment of paid search and document the importance of accounting for the counterfactual.

1.2. Our approach and contributions

With almost a decade of research, it remains unclear as to what is an "appropriate" or "best" attribution measure. IVH appears to be the most popular; but, to the best of our knowledge, there is no systematic framework to support it. In fact, as we later discuss, IVH suffers from serious drawbacks in certain settings. On the other hand, SV has a strong theoretical justification and a number of desirable properties such as *efficiency*, *symmetry*, *linearity*, and *null player*; in fact, SV is the unique solution to a cooperative game that has all these properties. However, estimating SV exactly is computationally intractable in general, and one has to resort to approximations (Avrachenkov et al. 2012, Castro et al. 2009, Fatima et al. 2008, Liben-Nowell et al. 2012, Littlechild and Owen 1973, Maleki et al. 2013, Michalak et al. 2013, Owen 1972). In addition, as we show in Section 4.2, SV is not counterfactual in nature. We seek a metric that has the desirable properties of SV, and yet is tractable and able to accommodate counterfactual reasoning.

Our main contributions are as follows. First, we construct an abstract Markov chain model for the user journey through the conversion funnel that generalizes most of the existing Markovian models in the attribution literature. At every period, the user is in one of the finitely many *states*. The advertiser observes the state and takes an *action*. The user transitions to a random state distributed according to a probability mass function that depends on the current state and the advertising action. We propose attributing value to each state-action pair, which is a generalization of the existing approaches that attribute only to advertising actions. This extension allows us to capture state-specific attribution for each action, the need for which is motivated by Bleier and Eisenbeiss (2015). Also, we analytically quantify the value of state-specific attribution and show that attributing at a state-action level allows one to capture a multitude of interactions that are missed otherwise.

Second, we develop a series of intuitive canonical network structures in our Markovian model that serve as a robustness check for any attribution scheme. Using these networks, we show the current attribution metrics (LTA, IVH, and uniform) have serious limitations. Furthermore, we show how one can compute state-specific SV for each action in our Markovian model and highlight that it does not adjust for the counterfactual and hence, is not an appropriate metric for attribution in our setting. To the best of our knowledge, this is the first work in the literature to analyze the various existing attribution metrics using a common framework.

Third, we develop an axiomatic framework for attribution in online advertising, which is the main contribution of this work. In particular, we propose a new metric for attribution in the Markovian model of user behavior, that we refer to as *counterfactual adjusted Shapley value*. We show that our proposed metric inherits the desirable axioms of the classical SV. We also demonstrate its robustness to a mix of network structures that highlighted limitations of the existing metrics. In addition, we establish that our metric admits a crisp characterization under our Markovian model. It coincides with a *unique-uniform attribution scheme* that explicitly adjusts for the counterfactual. In turn, it can be interpreted as a correction to the commonly used uniform attribution scheme. Furthermore, we exploit this characterization to develop simple algorithms to estimate our metric. We show the scalability of our framework by performing numerical experiments on a real-world large-scale dataset.

Outline. The remainder of this paper is organized as follows. In Section 2, we introduce the Markovian model that describes the user journey as a function of ad exposure. We discuss how the existing attribution schemes (LTA, IVH, and uniform) can be captured through the Markovian model and highlight their limitations in Section 3. In Section 4, we showcase the drawback of directly applying SV to measure attribution and then present our novel metric along with the axioms it satisfies. We then characterize our metric as an adjusted unique-uniform attribution scheme in Section 5 and propose simple algorithms to estimate it, followed by numerical experiments in Section 6. In Section 7, we discuss the value of state-specific attribution. We conclude in Section 8 where we also discuss some directions for ongoing and future research.

2. Model

We propose a Markovian model for user behavior where transitions in user's state are stochastic, and are a function of only the current state and the advertiser action. This process ends when the user quits (leaves the system) or converts (buys the product). In Section 2.1, we define the components of the Markov chain (denoted by $\mathcal{M}$). When constructing the Markov chain, we keep most of its elements abstract to showcase its flexibility. In Section 2.2, we shed light on the abstractness through practical examples. We conclude this section by defining the attribution problem over our Markovian model (Section 2.3).

We note that our model builds upon existing works in this literature that have also used a Markovian model to describe user behavior (Abhishek et al. 2012, Anderl et al. 2016, Kakalejčík et al. 2018). In particular, we define our model in abstract terms and therefore, most of the existing models can be seen as special instances of our model.

2.1. Markovian model of user behavior

We first discuss the state space of the Markov chain followed by the arrival process of the users. We then elaborate on the action space of the advertiser. Finally, we define the transition probabilities.

State space. We define $\mathbb{S} := \{s\}_{s=1}^m$ as the set of states excluding the two absorbing states (quit q and conversion c) and $\mathbb{S}^+ := \mathbb{S} \cup \{q, c\}$. In order to highlight the flexibility of our model, we do not yet give a concrete meaning to a state.

Arrival process. External traffic arrives at state $s \in \mathbb{S}$ w.p. λ_s (*initial state probability*). We define the vector $\lambda \in \mathbb{R}^m$ as $[\lambda_s]_{s\in\mathbb{S}}$. We assume no external traffic arrives at c and q.

Action space. We define $\mathbb{A} := \{a\}_{a=1}^n$ as the set of actions the advertiser can take, such as sending an e-mail or showing a display ad. We include the no-ad action $(a = 1)$[2] in $\mathbb{A}$ and define β_s^a to be the probability that an advertiser takes action $a \in \mathbb{A}$ at state $s \in \mathbb{S}$. We denote by β the collection of all β_s^a values and assume it is fixed.

Transition probabilities. We denote by $p_{ss'}^a$ the probability a user moves from $s \in \mathbb{S}$ to $s' \in \mathbb{S}$ in one transition given the advertiser takes action $a \in \mathbb{A}$ at s. Also, for all $(s, s') \in \mathbb{S}^2$, we define $p_{ss'}^\beta := \sum_{a\in\mathbb{A}} \beta_s^a p_{ss'}^a$, which denotes the *average* transition probability. To keep the notation concise, for each $a \in \mathbb{A}$, we define the matrix $P^a := [p_{ss'}^a]_{(s,s')\in\mathbb{S}^2} \in \mathbb{R}^{m\times m}$. Furthermore, $P^\beta := [p_{ss'}^\beta]_{(s,s')\in\mathbb{S}^2} \in$

[2] We are reserving $a = 0$ for future use as it will become transparent in Section 4.2.

$\mathbb{R}^{m\times m}$ and $B^a := \mathrm{diag}([\beta_s^a]_{s\in\mathbb{S}}) \in \mathbb{R}^{m\times m}$ for each $a \in \mathbb{A}$ is a diagonal matrix. Clearly, P^β represents the transition matrix over the *partial* state space $\mathbb{S}$ and $P^\beta = \sum_{a\in\mathbb{A}} B^a P^a$. For all $s \in \mathbb{S}$, we define the vector $p_s^a \in \mathbb{R}^m$ as the s-th row of P^a for all $a \in \mathbb{A}$ and $p_s^\beta \in \mathbb{R}^m$ as the s-th row of P^β. Next, we state the only assumption we make on our problem primitives.

ASSUMPTION 1 **(Absorption)**. *The Markov chain with P^β as its transition matrix is absorbing, i.e., the probability each user will eventually either quit or convert from any state equals 1.*

We note that Assumption 1 is very mild and we validate it on real data in Section 6. Intuitively, it is equivalent to saying that from any state, there exists a positive probability path to one of the absorbing states.

So far, we have not discussed the transitions to and from states q and c. We use the notation p_{sc}^a and p_{sc}^β to denote the action-specific and average one-step transition probabilities from $s \in \mathbb{S}$ to c for all $a \in \mathbb{A}$. (For transitions from $s \in \mathbb{S}$ to q, replace the index c by q.) Since both q and c are absorbing states, the transitions from them are self-loops w.p. 1. We define $p_c^a := [p_{sc}^a]_{s\in\mathbb{S}} \in \mathbb{R}^m$ for each $a \in \mathbb{A}$ and $p_c^\beta := [p_{sc}^\beta]_{s\in\mathbb{S}} \in \mathbb{R}^m$. (Note that p_s^a and p_s^β for $s \in \mathbb{S}$ correspond to the probabilities of *leaving* s whereas p_c^a and p_c^β correspond to the probabilities of *entering* c.) Next, we shift our focus to three quantities of interest for this Markov chain.

Expected number of visits. We define $F^\beta \in \mathbb{R}^{m\times m}$ as the *fundamental matrix* of $\mathcal{M}$, i.e., its (i,j)-th entry equals the expected number of visits to state j if the initial state is i. Results in Markov chain theory (Grinstead and Snell 2012) imply that $F^\beta = (I - P^\beta)^{-1}$ and Assumption 1 ensures that F^β exists and is finite.

Effective arrival rate. Let μ_s^β denote the *effective* arrival rate into state $s \in \mathbb{S}$ and $\mu^\beta := [\mu_s^\beta]_{s\in\mathbb{S}} \in \mathbb{R}^m$. It is easy to show that $\lambda^\top + (\mu^\beta)^\top P^\beta = (\mu^\beta)^\top$; hence, $(\mu^\beta)^\top = \lambda^\top F^\beta$.

Eventual conversion probability. We define h_s^β as the probability of eventually being absorbed in c from state $s \in \mathbb{S}$ and the vector $h^\beta := [h_s^\beta]_{s\in\mathbb{S}} \in \mathbb{R}^m$. It is easy to show that $h^\beta = P^\beta h^\beta + p_c^\beta$. Thus, $h^\beta = F^\beta p_c^\beta$. We set $h_q^\beta = 0$ and $h_c^\beta = 1$.

2.2. Model discussion

We have so far defined the state and action spaces in abstract terms. Now, we provide some possible mappings from physical quantities to these.

As a first possibility, the state may be composed of past interactions with the user, e.g., the number of visits to the product website, number of e-mails received, number of e-mails opened, number of display ads seen, number of display ads clicked, etc. The number of states may become too large, resulting in potential tractability issues. However, one can quantize the counts to control the size of the state space. The advertiser's action space includes doing nothing (no ad), or sending an e-mail, or showing a display ad, etc.

As another possibility, the state space may correspond to the widespread conversion funnel used in the marketing literature which captures the journey of the user from the state of being unaware of the product to eventually becoming interested and finally purchasing. Given the fine granularity of data available to advertisers, we believe the advertiser can infer such states using an appropriate statistical model. The action space remains the same as above.

We note that the definitions of state and action spaces can be customized as per the needs of the advertiser (since our attribution framework is developed for the abstract Markovian model). In fact, the state space should very much be context-specific and driven by the user features that are relevant to a particular advertiser.

2.3. The attribution problem

The total value[3] generated by the network equals $\lambda^\top h^\beta$. The goal of an attribution model is to allocate this value to the underlying *agents* (state-action pairs)[4] in the system while accounting for the counterfactual. For any given $(s,a) \in \mathbb{S} \times \mathbb{A}$, we use π_s^a to denote the corresponding attribution. Naturally, a desirable property is that the total value in the system should be credited back to the agents, i.e., $\sum_{s\in\mathbb{S}} \sum_{a\in\mathbb{A}} \pi_s^a = \lambda^\top h^\beta$.

3. Current approaches and limitations

We now define in our context the commonly used attribution schemes (LTA, IVH, and uniform) and illustrate their limitations. At a high-level, both LTA and uniform are "backward looking" in the sense that they split the value of a path after observing its realization (from the initial state to the end). LTA allocates all the value to the last interacting agent whereas uniform allocates it equally to all the agents that appeared in the path. On the other hand, IVH is "forward looking" since it attributes to a given agent based on what will happen in the future. It does not require the knowledge of how the actual path will unfold. All it requires is a model that can output the change in eventual conversion probability that results if the agent is removed from the path.

3.1. Last touch attribution (LTA)

In LTA, all the value generated due to a purchase is attributed to the last state-action pair in the path. Denote by $\mathcal{P}$ a random path (over state-action pairs) sampled uniformly from $\mathcal{M}$. For $(s,a) \in \mathbb{S} \times \mathbb{A}$, define

$$\hat{w}_s^a(\mathcal{P}) := \begin{cases} 1 & \text{if } \mathcal{P} \text{ converts and } (s,a) \text{ is the last agent in } \mathcal{P} \\ 0 & \text{otherwise.} \end{cases}$$

[3] We assume that the sale of one unit of the product generates a value of 1. If the true number is different from 1, the numbers can be scaled accordingly.

[4] One might want to consider actions as the agents. However, we view the state-action pairs as the agents, because the impact of an action is likely to be state dependent.

The attribution to $(s,a) \in \mathbb{S} \times \mathbb{A}$ under the LTA scheme equals

$$\pi_s^{a,\mathrm{LTA}} := \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[\hat{w}_s^a(\mathcal{P})\right].$$

In other words, (s,a) receives complete credit of each *converted* path it is the last agent of. Trivially, LTA attributes exactly all the value generated by the network (*efficient*) and is easy to implement. However, LTA appears "unfair" since the states are not rewarded for moving the user up in the conversion funnel. We demonstrate this limitation next.

Example 1 (LTA is unfair). Consider the network in Figure 1 with self-loop probability $p = 0$ (*line*). Under LTA, all the value goes to the ad action at state m although state 1 brings in all the traffic and the ad action at each state pushes the user towards conversion.

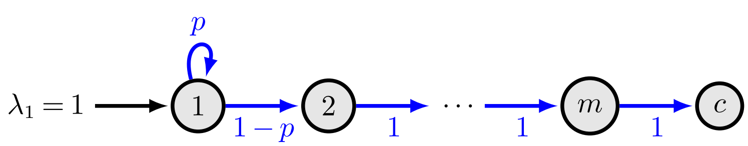

Figure 1 Network for Examples 1, 2, and 4. The action space consists of two actions: show no-ad and show an ad. The lines denote transitions if an ad is shown. For brevity, we do not show the transitions if an ad is not shown (to quit state w.p. 1). The advertiser shows an ad at all states w.p. 1.

3.2. Incremental value heuristic (IVH)

IVH allocates to each state-action pair (s,a) the increase in the eventual conversion probability by taking action a in state s as opposed to the no-ad action. For each $a \in \mathbb{A}$, we first define an auxiliary variable that captures the corresponding forward looking increment *conditioned* on the user being at a given state:

$$z^{a,\mathrm{IVH}} := \underbrace{P^a h^\beta + p_c^a}_{\text{action } a} - \underbrace{(P^1 h^\beta + p_c^1)}_{\text{no-ad action}} = \underbrace{(P^a - P^1)h^\beta}_{\text{eventual}} + \underbrace{(p_c^a - p_c^1)}_{\text{immediate}},$$

where $z^{a,\mathrm{IVH}} = [z_s^{a,\mathrm{IVH}}]_{s\in\mathbb{S}} \in \mathbb{R}^m$. The scalar $z_s^{a,\mathrm{IVH}}$ can be interpreted as the allocation at a "trace" level, i.e., if the advertiser observes a user at state s and decides to take action a, the corresponding allocation would be $z_s^{a,\mathrm{IVH}}$. However, we need to scale this metric for it to be seen as attribution over the entire population. In particular, given that at state s, the effective arrival rate is μ_s^β and action a is taken w.p. β_s^a, the IVH attribution is given by

$$\pi_s^{a,\mathrm{IVH}} := \mu_s^\beta \beta_s^a z_s^{a,\mathrm{IVH}}. \tag{1}$$

Although IVH is tractable and somewhat adjusts for the counterfactual, a serious limitation is that it can distribute more value than the network generates because it pays for both eventual and immediate conversions; consequently, each conversion may be accounted for several times. We now show that this is, indeed, possible.

Example 2 (IVH can over-allocate). Consider the network in Figure 1 with $p = 0$. Under IVH, all states receive an attribution of 1 for the show ad action and 0 for the no-ad action, resulting in a total allocation of m even though the value generated equals 1.

A common workaround to ensure IVH in budget-balanced is to normalize the output by an appropriate constant (*normalized IVH*). Normalizing the numbers in Example 2 results in the ad action at each state receiving $1/m$, which appears reasonable. However, as we illustrate in Example 3, even normalized IVH can be problematic.

Example 3 (Normalized IVH can be unjustified). Consider the network in Figure 2. IVH attributes 1 to state 1 and 1/2 for all other states. Thus, normalized IVH attributes

$$\begin{cases} \frac{1}{1+(m-1)/2} & \text{if state equals 1} \\ \frac{1/2}{1+(m-1)/2} & \text{otherwise.} \end{cases}$$

Clearly, as $m \to \infty$, all the value goes to states 2 to m. This appears inappropriate as state 1 "deserves" at least half the total value since 50% of the users convert immediately after state 1.

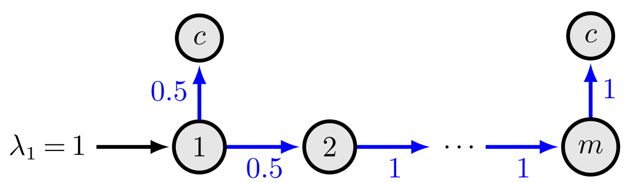

Figure 2 Network for Example 3. The action space consists of two actions: show no-ad and show an ad. Solid blue lines denote transitions if an ad is shown. For brevity, we do not show the transitions if an ad is not shown (to quit state w.p. 1). The advertiser shows an ad at all states w.p. 1.

3.3. Uniform attribution

Under a uniform attribution scheme, the credit generated is attributed equally to all the state-action pairs that are encountered in a path. Denote by $\mathcal{P}$ a random path (over state-action pairs) sampled uniformly from $\mathcal{M}$. For $(s,a) \in \mathbb{S} \times \mathbb{A}$, define

$$\bar{w}_s^a(\mathcal{P}) := \begin{cases} \frac{n_s^a}{|\mathcal{P}|} & \text{if } \mathcal{P} \text{ converts and } (s,a) \in \mathcal{P} \\ 0 & \text{otherwise,} \end{cases}$$

where n_s^a equals the number of times (s,a) appears in $\mathcal{P}$ and $|\mathcal{P}|$ denotes the number of state-action pairs in $\mathcal{P}$ (not necessarily unique). The uniform attribution for $(s,a) \in \mathbb{S} \times \mathbb{A}$ is

$$\pi_s^{a,\text{uni}} := \mathbb{E}_{\mathcal{P} \sim \mathcal{M}} \left[\bar{w}_s^a(\mathcal{P}) \right]. \tag{2}$$

In other words, (s,a) receives an "equal cut" of each *converted* path it contributes to. Such a scheme is efficient and scalable. On the line network, it attributes $1/m$ to the ad action at each state, which seems reasonable. However, it does not account for the counterfactual. Furthermore, it creates undesired incentives, which we show next.

Example 4 (Undesired incentive). Consider the network in Figure 1 with $p < 1$. For a given path with n_1 occurrences of state 1, the uniform attribution scheme attributes $n_1/(n_1 + m - 1)$ to state 1. As p increases, state 1 receives more credit. However, this seems inappropriate since it gives an undesired incentive to the ad action at state 1 to "game the system" via self-loops. Ideally, an attribution scheme should incentivize the ads to push the user towards conversion.

Remark 1 (Unique-uniform). A simple fix to the above drawback is to attribute based on the number of unique agents. For instance, in the example above, the *unique-uniform* scheme would attribute $1/m$ to the ad action at each state. However, this scheme still does not account for the counterfactual, and there is no formal rationale for it. (In the following sections, we will see that adjusting for uniqueness is in fact supported by a strong mathematical rationale and will be part of our prescribed method.)

To summarize this section, the above examples highlight some limitations of existing heuristics but also point to the fact that no single heuristic "dominates the other". In some way, IVH possesses a counterfactual form giving it some practical appeal. At the same time, it is only forward looking, ignoring all past actions and their potential contributions. Uniform, on the other hand, accounts for past actions, but is not counterfactual. Ideally, an attribution scheme would provide the best of both worlds, properly accounting for past actions while also accounting for the counterfactual associated with not advertising.

4. Shapley value (SV)

We now study SV as a measure for attribution. We first present a primer on SV (Section 4.1) followed by a discussion on why a direct application to our context does not suffice (Section 4.2). We then adapt SV accordingly (Section 4.3) and conclude this section by discussing the performance of our proposed measure on all the motivating examples (Section 4.4). In this section, we hope to convince the reader that, leaving computational tractability aside, SV (adjusted for the counterfactual) is an appropriate metric for attribution.

4.1. A quick primer

SV is a well-accepted concept from cooperative game theory that is used to distribute the payoff generated by a *coalition* to the *players* in the coalition (Shapley 1953). Given a finite set $\mathbb{P}$ of players, the *characteristic function* $v(\cdot)$ maps a coalition $\mathcal{X} \subseteq \mathbb{P}$ to the value (a real number) generated by the coalition. The value $v(\emptyset)$ of the empty coalition is normalized to 0. SV distributes the value $v(\mathbb{P})$ of the grand coalition to a player $r \in \mathbb{P}$ as follows:

$$\pi_r^{\text{Shap}} := \sum_{\mathcal{X} \subseteq \mathbb{P} \setminus \{r\}} w_{|\mathcal{X}|} \times \{v(\mathcal{X} \cup \{r\}) - v(\mathcal{X})\},$$

where

$$w_{|\mathcal{X}|} := \frac{|\mathcal{X}|!(|\mathbb{P}| - |\mathcal{X}| - 1)!}{|\mathbb{P}|!}.$$

The attractiveness of SV is rooted in the fact that it is the *only* solution to a cooperative game that has the following four desirable properties:

1. **Efficiency**: $\sum_{r\in\mathbb{P}} \pi_r = v(\mathbb{P})$.

2. **Symmetry**: Consider players $r, r' \in \mathbb{P}$ such that for any $\mathcal{X} \subseteq \mathbb{P} \setminus \{r, r'\}$, r and r' are *equivalent*, i.e., $v(\mathcal{X} \cup \{r\}) = v(\mathcal{X} \cup \{r'\})$. Then, $\pi_r = \pi_{r'}$.

3. **Linearity**: Consider two characteristic functions $v(\cdot)$ and $w(\cdot)$. Linearity states that for all players $r \in \mathbb{P}$, $\pi_r(v + w) = \pi_r(v) + \pi_r(w)$ and $\pi_r(\alpha v) = \alpha\pi_r(v)$ for all $\alpha \in \mathbb{R}$.

4. **Null player**: Suppose player $r \in \mathbb{P}$ does not add any value to any coalition, i.e., for all $\mathcal{X} \subseteq \mathbb{P} \setminus \{r\}$, $v(\mathcal{X} \cup \{r\}) = v(\mathcal{X})$. Then, $\pi_r = 0$.

4.2. Direct application

SV is a natural candidate for decomposing network value because the state-action pairs can be viewed as players participating in a cooperative game to achieve a common goal of converting the users. In particular, given the action intensities β, the network generates a value of $\lambda^\top h^\beta$ and the state-action specific SV $\pi_s^{a,\text{Shap}}$ represents the credit allocated to each state-action pair $(s, a) \in \mathbb{S} \times \mathbb{A}$. We first define the underlying components (players, coalition, and characteristic function) in our context and then discuss an important drawback of such a naive application.

Player. We view each state-action pair $(s, a) \in \mathbb{S} \times \mathbb{A}$ as an underlying player in the cooperative game. Consequently, our SV will not just be action-specific, but also state-specific.

We recall that we consider the no-ad action $(a = 1)$ as one of the actions in $\mathbb{A}$. We do so to respect the fact that the no-ad action can generate value (either in terms of network effect or in terms of direct conversions) and hence, we should attribute value to the players corresponding to the no-ad action $(\{(s, 1)\}_{s\in\mathbb{S}})$ as opposed to considering no-ad action as the "default action" that fills up the empty coalition. It should be easy to see that using the no-ad action as the "default action" would allow the other actions to free ride on the value generated by the no-ad action, and hence, undermine the contribution due to the no-ad action. Accordingly, such a view motivates the need of a "zero-value action" $(a = 0)$ as the action that has no value. In other words, at each state, the zero-value action directs a user to the quit state w.p. 1.

Coalition. A coalition $\mathcal{X} \subseteq \mathbb{S} \times \mathbb{A}$ refers to a collection of players (state-action pairs). We construct a coalition by starting with the *empty coalition*, which corresponds to the Markov chain containing all the states but the zero-value action $(a = 0)$ being taken at each state w.p. 1. If a state-action pair, say (s, a), is added to the empty coalition, then the advertiser takes action a at state s w.p. β_s^a and so on. Therefore, for each coalition $\mathcal{X}$, we obtain a corresponding β, which we denote by $\beta^{\mathcal{X}}$.

Characteristic function. Given $\mathcal{X} \subseteq \mathbb{S} \times \mathbb{A}$, we define

$$v(\mathcal{X}) := \lambda^\top h^{\beta^{\mathcal{X}}}.$$

We define the action probabilities collections $\beta^\emptyset$ and β^* such that they correspond to the empty coalition $\emptyset$ and complete coalition $\mathbb{S} \times \mathbb{A}$, respectively. Thus, $\beta^\emptyset = 0$ and $\beta^* = \{\beta_s^a\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$. We satisfy $v(\emptyset) = 0$ and $v(\mathbb{S}\times\mathbb{A}) = \lambda^\top h^{\beta^*}$ is the total network value. Since the value depends on the underlying Markov chain $\mathcal{M}$, we will use the notation $v_{\mathcal{M}}(\cdot)$ when the emphasis on $\mathcal{M}$ is necessary.

Shapley value. The SV for each $(s,a) \in \mathbb{S}\times\mathbb{A}$ is

$$\pi_s^{a,\text{Shap}} := \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a)\}} w_{|\mathcal{X}|} \times \{v(\mathcal{X}\cup\{(s,a)\}) - v(\mathcal{X})\}, \tag{3}$$

where

$$w_{|\mathcal{X}|} := \frac{|\mathcal{X}|!(mn - |\mathcal{X}| - 1)!}{(mn)!}. \tag{4}$$

SV depends on our choice of the characteristic function $v(\cdot)$ and the underlying Markov chain $\mathcal{M}$ and hence, we will use the notation $\pi_s^{a,\text{Shap}}(v)$ and $\pi_s^{a,\text{Shap}}(\mathcal{M})$ when the emphasis on $v(\cdot)$ and $\mathcal{M}$ is necessary. Though this view of attribution inherits the desirability of SV, it suffers from a critical flaw: it does not adjust for the counterfactual (Example 5).

Example 5 (Need for counterfactual). Consider the network in Figure 3. Showing an ad in state 1 should not get any attribution since it provides no *additional* value over the counterfactual action (no ad). However, (3) attributes all the value to the ad action. Furthermore, even LTA and uniform attribution fail this sanity check.

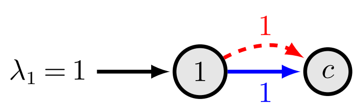

Figure 3 Network for Example 5. The action space consists of two actions: show no-ad and show an ad. Solid blue (dashed red) lines denote transition if an ad is shown (not shown). The advertiser shows an ad w.p. 1.

4.3. Our approach

As we alluded to earlier, we seek a measure that provides the best of two worlds: (1) appropriately captures the contributions of the past actions (capturing the feature of, e.g., uniform and SV)[5] and (2) exhibits counterfactual reasoning (capturing a feature of IVH). To this end, we focus on adapting SV to account for the counterfactual and we do so by adhering to the causality framework of Rubin (1974). We first define a *counterfactual player* followed by our novel definition of *counterfactual adjusted Shapley value (CASV)*. We then show the desirability of CASV.

[5] We will see in Section 5 that uniform and SV are closely linked.

Counterfactual player. Following Rubin (1974), the counterfactual to taking action a at state s w.p. β_s^a is to take the no-ad action (action 1) at state s w.p. β_s^a (instead of β_s^1). Accordingly, for a given player $(s,a) \in \mathbb{S} \times \mathbb{A}$, we denote the counterfactual player as $(s,1)^a$, where the "a" in the superscript captures the dependence on β_s^a.

Counterfactual adjusted Shapley value. The game theoretic setup (player, coalition, characteristic function) is the same as in Section 4.2. For each $(s,a) \in \mathbb{S} \times \mathbb{A}$, we define CASV as

$$\psi_s^{a,\mathrm{Shap}} := \sum_{\mathcal{X} \subseteq \{\mathbb{S} \times \mathbb{A}\} \setminus \{(s,a)\}} w_{|\mathcal{X}|} \times \{v(\mathcal{X} \cup \{(s,a)\}) - v(\mathcal{X} \cup \{(s,1)^a\})\}, \tag{5}$$

where $w_{|\mathcal{X}|}$ is the same as in (4). We use the symbol ψ instead of π to differentiate CASV from SV. The only change we make in going from $\pi_s^{a,\mathrm{Shap}}$ to $\psi_s^{a,\mathrm{Shap}}$ is that we replace $v(\mathcal{X})$ by $v(\mathcal{X} \cup \{(s,1)^a\})$, i.e., CASV captures the value added by a player over its counterfactual.

Intuitively, it is unclear whether CASV under the current cooperative game is equivalent to SV under a different cooperative game. However, it is possible to view CASV as the difference between two SVs:

$$\psi_s^{a,\mathrm{Shap}}(\mathcal{M}) = \pi_s^{a,\mathrm{Shap}}(\mathcal{M}) - \pi_s^{a,\mathrm{Shap}}(\mathcal{M}_s^a), \tag{6}$$

where $\mathcal{M}$ denotes the original Markov chain and $\mathcal{M}_s^a$ denotes the *counterfactual network* for (s,a) where we replace the transition probabilities of (s,a) by those of $(s,1)$. Here, $\pi_s^{a,\mathrm{Shap}}(\mathcal{M}_s^a)$ is the *counterfactual value* of (s,a), i.e., the value generated if no-ad was shown instead of ad a at state s.

Obviously, it is of interest to check how CASV performs with respect to the desirable properties that motivated SV in the first place. However, one needs to be careful when discussing these axioms in the counterfactual context. To be specific, efficiency should now pertain to the redistribution of additional value generated over the counterfactual value. Similarly, the definition of equivalent players should be adjusted (discussed in Remark 2) and null player should correspond to a player with zero value-add. We formalize the desirability of CASV in Theorem 1. We present the proof in Appendix A.

THEOREM 1 **(Axioms)**. *CASV satisfies the following axioms:*

1. ***Counterfactual efficiency****: The sum of CASVs equals the additional value generated over the counterfactual value, i.e.,*

$$\sum_{(s,a) \in \mathbb{S} \times \mathbb{A}} \psi_s^{a,Shap} = v(\mathbb{S} \times \mathbb{A}) - \sum_{(s,a) \in \mathbb{S} \times \mathbb{A}} \pi_s^{a,Shap}(\mathcal{M}_s^a).$$

2. ***Counterfactual symmetry****: If $(s,a) \in \mathbb{S} \times \mathbb{A}$ and $(s',a') \in \mathbb{S} \times \mathbb{A}$ are counterfactual equivalent, i.e.,* *(A)* $v(\mathcal{X}\cup\{(s,a)\})-v(\mathcal{X}\cup\{(s,1)^a\})=v(\mathcal{X}\cup\{(s',a')\})-v(\mathcal{X}\cup\{(s',1)^{a'}\})$ *and*

(B) $v(\mathcal{X}\cup\{(s,1)^a,(s',a')\})=v(\mathcal{X}\cup\{(s',1)^{a'},(s,a)\})$

for all $\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}$*, then*

$$\psi_s^{a,Shap}=\psi_{s'}^{a',Shap}.$$

3. ***Linearity****: Consider two characteristic functions* $v(\cdot)$ *and* $w(\cdot)$*. For all* $(s,a)\in\mathbb{S}\times\mathbb{A}$*, we have*

$$\psi_s^{a,Shap}(v+w)=\psi_s^{a,Shap}(v)+\psi_s^{a,Shap}(w)$$

and for all $\alpha\in\mathbb{R}$*,*

$$\psi_s^{a,Shap}(\alpha v)=\alpha\psi_s^{a,Shap}(v).$$

4. ***Counterfactual null player****: Consider a player* $(s,a)\in\mathbb{S}\times\mathbb{A}$ *that has a zero value-add to all coalitions that do not contain* (s,a)*, i.e., for all* $\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a)\}$*,*

$$v(\mathcal{X}\cup\{(s,a)\})=v(\mathcal{X}\cup\{(s,1)^a\}).$$

Then, $\psi_s^{a,Shap}=0$*.*

Furthermore, CASV is the unique solution to satisfy these four counterfactual axioms.

Remark 2 (Counterfactual equivalent players). In Theorem 1, we define $(s,a)\in\mathbb{S}\times\mathbb{A}$ and $(s',a')\in\mathbb{S}\times\mathbb{A}$ to be counterfactual equivalent if they satisfy (A) and (B). To develop an intuitive understanding of these conditions, one can consider the following three conditions for all $\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}$:

(A1) $v(\mathcal{X}\cup\{(s,a)\})=v(\mathcal{X}\cup\{(s',a')\})$,

(A2) $v(\mathcal{X}\cup\{(s,1)^a\})=v(\mathcal{X}\cup\{(s',1)^{a'}\})$, and

(B) $v(\mathcal{X}\cup\{(s,1)^a,(s',a')\})=v(\mathcal{X}\cup\{(s',1)^{a'},(s,a)\})$.

In simple words, consider a coalition that does not contain either of the two given players. Condition (A1) states that adding either of the two players has the same effect on the network value and condition (A2) says so for the counterfactual players. The remaining case of interest is what happens when, on one hand, we add the first player and the counterfactual of the second player whereas, on the other hand, we do the opposite (add the second player and the counterfactual of the first player). Condition (B) states that doing either is equivalent in terms of the characteristic function. Note that we only need conditions (A) and (B) for two players to be counterfactual equivalent, and (A1) and (A2) are sufficient conditions for (A). It is worth mentioning that our definition of counterfactual equivalent players is a generalization of equivalent players from Section 4.1. In particular, equivalent players only need to satisfy condition (A1) and since the notion of a counterfactual player does not exist in the SV context, conditions (A2) and (B) are irrelevant.

We conclude this subsection with a discussion regarding the attribution to the no-ad action. By definition of CASV, the no-ad action at each state receives zero counterfactual credit. However, one can attribute the counterfactual value $\sum_{(s,a)\in\mathbb{S}\times\mathbb{A}}\pi_s^{a,\mathrm{Shap}}(\mathcal{M}_s^a)$ ("residual") to the no-ad action in a post hoc fashion. In other words, the post hoc attribution to the no-ad action at state $s\in\mathbb{S}$ equals

$$\sum_{a\in\mathbb{A}}\pi_s^{a,\mathrm{Shap}}(\mathcal{M}_s^a)=\pi_s^{1,\mathrm{Shap}}(\mathcal{M})+\sum_{a=2}^{n}\pi_s^{a,\mathrm{Shap}}(\mathcal{M}_s^a), \tag{7}$$

where we use the fact that $\mathcal{M}_s^1=\mathcal{M}$. The two terms in (7) have an intuitive interpretation. The first term captures the value generated by the no-ad action due to a positive value of β_s^1 whereas the second term captures the value that *would have* been generated if the no-ad action was used in place of the other actions.

4.4. Revisiting canonical examples

We now revisit the networks discussed so far and show that our CASV measure does not suffer from the same limitations as existing heuristics.

Example 6. For the network in Figure 1 with $p<1$, CASV at every state equals $1/m$ for the ad action and 0 otherwise.

Example 7. For the network in Figure 2, CASV attributes 0 to the no-ad action at all the states and the following to the ad action:

$$\begin{cases}\frac{1}{2}+\frac{1}{2m} & \text{if } s=1\\ \frac{1}{2m} & \text{otherwise.}\end{cases}$$

This seems sensible as half of the paths convert just due to the ad action at state 1 whereas the other half of the paths convert with the help of the ad action at all the m states.

Example 8. For the network in Figure 3, CASV allocates zero credit to the ad action, which seems appropriate. However, CASV of the no-ad action also equals zero, which raises the following question: which player gets credit for the value in the system? The answer is, as discussed at the end of Section 4.3, the no-ad action.

In sum, CASV appears to be an appealing measure for attribution: it has a number of desirable properties and appears robust to various network structures. However, similar to SV, computing CASV using (5) requires an exponential runtime in the number of underlying players. This is an important tractability concern that could render CASV impractical in settings with moderately-sized state and action spaces. In the next section, we focus on representing CASV in a manner that is amenable to being computed efficiently, under our Markovian model.

5. Characterizing counterfactual adjusted Shapley value

We now characterize CASV by exploiting the structure of the underlying cooperative game in order to build an intuitive understanding (Section 5.1) and use the characterization to develop simple algorithms for estimating CASV (Section 5.2).

5.1. Characterization

To characterize CASV, we use the fact that CASV can be expressed as a difference of two SVs (see (6)) and hence, analyze SV first. Proposition 1 connects the coalition-oriented game-theoretic construct of SV to the paths sampled from $\mathcal{M}$. In particular, we show that, under our Markovian setup, SV is (surprisingly) identical to the unique-uniform attribution scheme (motivated in Remark 1). We present the formal proof of Proposition 1 in Appendix B and highlight the intuition here.

PROPOSITION 1 **(SV equals unique-uniform)**. *Consider* $(s,a) \in \mathbb{S} \times \mathbb{A}$ *and the Markov chain* $\mathcal{M}$*. The SV of* (s,a) *as defined in* (3) *equals*

$$\pi_s^{a,Shap} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[w_s^a(\mathcal{P})\right],$$

where

$$w_s^a(\mathcal{P}) := \begin{cases} \frac{1}{u(\mathcal{P})} & \text{if } \mathcal{P} \text{ converts and } (s,a) \in \mathcal{P} \\ 0 & \text{otherwise.} \end{cases}$$

The function $u(\cdot)$ *returns the number of unique players and* $\mathcal{P}$ *is a path over players, i.e., state-action pairs.*

Proof sketch. There are two keys ideas at interplay here. First, to characterize SV, we observe that the characteristic function for a given coalition $\mathcal{X} \subseteq \mathbb{S} \times \mathbb{A}$ can be expressed as an expectation over the paths sampled from the Markov chain $\mathcal{M}(\mathcal{X})$ in which only the players in $\mathcal{X}$ are "active", i.e.,

$$v(\mathcal{X}) = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}(\mathcal{X})}\left[\mathbb{I}_{\{\mathcal{P} \text{ converts}\}}\right].$$

Using this observation allows us to express SV $\pi_s^{a,\text{Shap}}$ as an expectation over the paths from the original Markov chain $\mathcal{M}$:

$$\pi_s^{a,\text{Shap}} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[\mathbb{I}_{\{\mathcal{P} \text{ converts}\}}\ (\star)\right],$$

where

$$(\star) := \sum_{\mathcal{X} \subseteq \{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a)\}} w_{|\mathcal{X}|}\left(\mathbb{I}_{\{\mathcal{X}(\mathcal{P}) \subseteq \mathcal{X}\cup\{(s,a)\}\}} - \mathbb{I}_{\{\mathcal{X}(\mathcal{P}) \subseteq \mathcal{X}\}}\right)$$

and $\mathcal{X}(\mathcal{P})$ denotes the set of unique players in path $\mathcal{P}$. The second key idea is to analyze $(\star)$. We show that $(\star)$ can be seen as the SV of a "carrier game" and hence, admits the following remarkably simple form:

$$(\star) = \begin{cases} \frac{1}{u(\mathcal{P})} & \text{if } (s,a) \in \mathcal{P} \\ 0 & \text{otherwise.} \end{cases}$$

Putting the two pieces together allows us to finish the proof. □

The above result enables one to characterize CASV, which constitutes the main result of this work.

THEOREM 2 **(CASV equals unique-uniform minus counterfactual)**. *Consider* $(s,a) \in \mathbb{S} \times \mathbb{A}$*, the Markov chain* $\mathcal{M}$*, and the counterfactual Markov chain* $\mathcal{M}_s^a$*. The CASV of* (s,a) *as defined in* (5) *equals*

$$\psi_s^{a,Shap} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[w_s^a(\mathcal{P})\right] - \mathbb{E}_{\mathcal{P}\sim\mathcal{M}_s^a}\left[w_s^a(\mathcal{P})\right],$$

where $w_s^a(\cdot)$ *is as defined in Proposition 1.*

There is a striking resemblance between this characterization of CASV and uniform attribution as defined in (2). In particular, the two expressions are identical except the definition of the weight function and the counterfactual adjustment. The weight function in CASV does not provide an incentive for players to "game the system" as it only rewards based on whether a player appears in the path or not (as opposed to the number of times it appears). In other words, if an ad a had to be shown multiple times (say n_s^a) at the *same* state s to make the user move to another state s', then the CASV weight function only rewards it once (as opposed to rewarding it n_s^a times). This seems very reasonable and, in fact, is the simple fix we proposed in Remark 1.

It is noteworthy that the coalition-oriented construct of CASV, which on the surface does not seem to be related to the paths of the underlying Markovian model, reduces to being expressed as a remarkably simple function of such paths. This connection is quite valuable as it helps to gain deeper insights regarding the structure of CASV and hence, build a better understanding. Next, we use this connection to develop simple algorithms to estimate CASV.

5.2. Algorithms

Given a set $\mathcal{D}$ of user paths where each user path consists of various state-action-state (s,a,s') tuples, the action-specific transition probabilities and initial state probabilities can be estimated using empirical frequencies[6]. To estimate CASV for player (s,a), one can sample paths from the

[6] We are assuming the advertiser knows β. If not, then it can be estimated similarly.

estimated Markov chains $\mathcal{M}$ and $\mathcal{M}_s^a$ and use Theorem 2 directly. However, this involves sampling from a different Markov chain $\mathcal{M}_s^a$ to estimate CASV for each player, which we address next.

A simple change of measure allows CASV to be expressed as

$$\psi_s^{a,\text{Shap}} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[w_s^a(\mathcal{P})\left(1 - \frac{g_s^a(\mathcal{P})}{g(\mathcal{P})}\right)\right], \tag{8}$$

where $g(\mathcal{P})$ and $g_s^a(\mathcal{P})$ denote the probabilities of observing path $\mathcal{P}$ under Markov chains $\mathcal{M}$ and $\mathcal{M}_s^a$, respectively. The ratio $g_s^a(\mathcal{P})/g(\mathcal{P})$ denotes the *importance weight* and it is easy to show that

$$\frac{g_s^a(\mathcal{P})}{g(\mathcal{P})} = \prod_{s'\in\mathbb{S}^+}\left(\frac{p_{ss'}^1}{p_{ss'}^a}\right)^{n_{ss'}^a(\mathcal{P})},$$

where $n_{ss'}^a(\mathcal{P})$ denotes the number of occurrences of the tuple (s,a,s') in $\mathcal{P}$ for each $s'\in\mathbb{S}^+$. Given (8), it suffices to sample paths from just one Markov chain ($\mathcal{M}$) to estimate CASV for all the players. In fact, each sample can be recycled for multiple players and the scheme can be implemented using a parallel (over both players and samples) architecture. We state this estimation procedure in Algorithm 1, where we assume we have access to the Markov chain $\mathcal{M}$. In order to be concise, we do not discuss the confidence intervals of the estimate of CASV that we get using the importance weights technique. Instead, we refer the reader to Bottou et al. (2013) for a discussion on computing such confidence intervals.

Algorithm 1 Estimating CASV given Markov chain $\mathcal{M}$

Require: Markov chain $\mathcal{M}$ and number of user paths S

1: Initialize: $\psi_s^{a,\text{alg}} = 0$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$
2: **for** $i=1$ to S **do**
3:     Sample path: $\mathcal{P}_i\sim\mathcal{M}$
4:     **if** $\mathcal{P}_i$ converted **then**
5:         **for** $(s,a)\in\mathcal{P}_i$ **do**
6:             $\psi_s^{a,\text{alg}} = \psi_s^{a,\text{alg}} + \frac{1}{u(\mathcal{P}_i)}\left(1 - \prod_{s'\in\mathbb{S}^+}(p_{ss'}^1/p_{ss'}^a)^{n_{ss'}^a(\mathcal{P}_i)}\right)$
7:         **end for**
8:     **end if**
9: **end for**
10: Normalize: $\psi_s^{a,\text{alg}} = \psi_s^{a,\text{alg}}/S$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$
11: **return** $\psi^{\text{alg}} := \{\psi_s^{a,\text{alg}}\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$

It is possible to use a Bayesian version of the approach described above. In particular, one can maintain a belief over the transition probability vector $[p_{ss'}^a]_{s'\in\mathbb{S}^+}$ in the form of a Dirichlet distribution for each $a \in \mathbb{A}$. (One can also maintain a similar belief over the initial state probabilities vector λ and action intensities vector $[\beta_s^a]_{a\in\mathbb{A}}$ for all $s \in \mathbb{S}$ if required.) Using Dirichlet-multinomial conjugacy and real data $\mathcal{D}$, the belief can be updated easily. The posterior belief can be used to generate samples of the Markov chain $\mathcal{M}$ itself. For each sample of $\mathcal{M}$, one can use Algorithm 1 to estimate CASV, resulting in posterior samples of CASV. Naturally, these posterior samples quantify the uncertainty in CASV that arises due to the noise in data and/or lack of data. Furthermore, this Bayesian approach is amenable to parallelization over the samples of $\mathcal{M}$.

6. Numerical experiments

We now shift our focus to evaluating the performance of our framework of attribution on a large-scale real-world dataset and comparing it against competing metrics (LTA, IVH, uniform, and SV). We first discuss the dataset composed of several million real-world user paths (Section 6.1) followed by the underlying Markovian model and its estimation using the real-world data (Section 6.2). We finally present the attribution metrics (computed on data simulated from the estimated model) corresponding to various schemes along with some discussion (Section 6.3). We note that our numerics are for illustrative purposes and in particular, to see how various attribution schemes differ.

6.1. Dataset

Our real-world dataset corresponds to a single product (software) promoted and sold on the Internet by a Fortune 500 company[7]. The dataset contains several million user paths with a few hundred thousand conversions (purchases). Each path starts with a "sign-up", i.e., the user creating an account on the company's website. From the date of the sign-up, we have access to the user's interaction with the company (*touchpoints*) for a period of 8 months[8]. The touchpoints can be classified based on the type of ad they correspond to:

- **E-mail ad**: There are three related touchpoints: (1) advertiser sending an e-mail (ad action), (2) user opening an e-mail (user action), and (3) user clicking on a link in the e-mail (user action).
- **Display ad**: There are two related touchpoints: (1) advertiser showing a display ad (ad action) and (2) user clicking on a display ad (user action).
- **Paid search ad**: There is only one touchpoint corresponding to paid search impressions: user clicking on the paid search ad (user action). Due to the way data is collected, an advertiser usually does not know whether a paid search ad is shown to a user if the user does not click on the ad.

[7] We do not disclose the name of the company and specific statistics related to real-data for anonymity reasons.

[8] Our user paths data suffers from the common issues such as we lose track of a user if he/she clears cookies from the web browser.

There are two additional touchpoints: (1) “sign-up” (user creates an account) and (2) “conversion” (user buys the product). If the user converted within 8 months, we truncate the path at the time of conversion (since we are interested in activity up to first conversion). The quit state is not explicitly observed in the data and we use the following rule to determine a transition to q: if there is no activity (user and advertiser) in the last 30 days and in the future, we mark the next state as q. Not every path ends in the conversion state or the quit state at the end of the 8 months period and we let it be that way (as opposed to forcing it to transition to the quit state).

We discarded certain dubious paths after discussing with our industry partner. In particular, paths with over 100 touchpoints were thrown away with the suspicion of possible bots and paths with any illogical sequence (for example, e-mail opened before being received) were also discarded (possible data recording errors). The occurrence of such paths was very low (less than 1%).

6.2. Markovian model and estimation

Our Markov chain in earlier sections is defined with abstract action and state spaces to showcase the flexibility of our framework. For the numerics, we use the following construction.

The action space $\mathbb{A}$ consists of four actions: (1) no-ad, (2) e-mail, (3) display ad, and (4) paid search *click*[9]. In our dataset, the no-ad action is not observed explicitly and we use the following rule to “implant” no-ad actions in the user paths: if there has been no touchpoint activity (user and advertiser) in the past τ days, we implant the no-ad action. We use $\tau = 10$ in our numerics and perform sensitivity analysis after we present the results (at the end of Section 6.3).

The state space $\mathbb{S}$ is motivated by the widely used conversion funnel in the marketing literature (Strong 1925, Howard and Sheth 1969, Barry 1987, Bettman et al. 1998, Court 2009, Elzinga et al. 2009, Kotler and Armstrong 2010, Mulpuru 2011, Jansen and Schuster 2011, Bruce et al. 2012) with the view that a user goes through the following states during his decision process: (1) unaware, (2) aware, (3) interest, and (4) desire. We use our microscopic user level data to objectively map the touchpoints to these states as follows:

- **Unaware**: User has received no e-mail, no display ad, and clicked on no paid search ad so far.
- **Aware**: At least one ad (e-mail or display) received by the user so far.
- **Interest**: At least one e-mail opened by the user so far.
- **Desire**: At least one ad (e-mail, display, or paid search) clicked by the user so far.

Clearly, a “back” transition (for example, transition from aware to unaware) is impossible in our state space construction but we allow for the possibility of a “jump” (for example, transition from unaware to interest). Furthermore, transitions to quit and conversion states are possible from any

[9] Since we only observe a paid search impression when it is clicked, we use the “click” as an action even though it is a user action (as opposed to an ad action). This is a limitation of the data and affects all attribution metrics, i.e., it is not specific to the framework we propose.

of the four states in $\mathbb{S}$. (The state space can be made more granular depending on the needs of an advertiser and the type of data available. For instance, being "aware" as a result of receiving an e-mail might be different from being "aware" due to seeing a display impression. Furthermore, the "level of awareness" might depend on the number of ads seen and hence, it might be worthwhile to define multiple states capturing different levels of awareness. We kept our construction simple for illustrative purposes.)

Having defined the action and state spaces, we discuss the estimation of the corresponding Markov chain $\mathcal{M}$ using real data. In particular, we need to estimate λ, β, and $\{P^a\}_{a\in\mathbb{A}}$. The touchpoints in the dataset allow us to construct "(s,a,s')" tuples. Since each user path starts with the "sign-up" touchpoint, the initial state is always "unaware" and hence, $\lambda_1 = 1$. To estimate β_s^a for all $(s,a) \in \mathbb{S}\times\mathbb{A}$, we simply compute the empirical frequencies, i.e., we count the number of times action a was taken at state s and divide it by the number of times we observe state s in the dataset. We employ the same technique[10] to estimate $p_{ss'}^a$ for all $(s,a,s') \in \mathbb{S}\times\mathbb{A}\times\mathbb{S}^+$. We acknowledge that our simple estimation scheme could potentially be improved. In particular, there might be undesired endogeneity effects in our data, which might result in biased estimates. For instance, a user who is more inclined to convert might be less likely to receive an ad, which is against the spirit of our β-randomized policy. However, we note that such estimation issues are not peculiar to our framework but they also affect other attribution schemes and hence, we place these challenges outside the scope of this work.

The high-level summary[11] of the parameter estimates is as follows. In terms of actions, e-mail has the highest frequency ranging from 40% to 90% depending on the state followed by the no-ad action (10% to 50%). Combined, e-mail and no-ad actions account for over 85% of the actions at each state. The action intensity of display impression ranges from 2% to 10% and of paid search click ranges from 0.1% to 7%[12]. In terms of transition probabilities, we note that from each state in $\mathbb{S}$, the eventual conversion probability is positive under each action, which validates the absorption assumption (Assumption 1). Furthermore, the self-loop probability for each state-action pair (except for the no-ad action at the unaware state and the impossible self-transitions) is over 90%, indicating "slow-moving" users. For the no-ad action at the unaware state, it is around 50%. In addition, there are intuitive patterns in the estimates. For example, a user is around 5 times

[10] We also experimented with the Bayesian approach discussed at the end of Section 5.2. Due to the enormous number of user paths in our dataset, the posterior variance in the Bayesian approach was extremely small and hence, the results obtained were very similar. We do not present them to be concise.

[11] We do not report the exact numbers to protect the private data of our provider.

[12] Interestingly, the highest (among all the states) action intensity of paid search click was at the first state in the funnel ("unaware"). We believe this is a result of "unaware" users actually knowing the product (from offline channels for instance).

more likely to move from the “aware” state to the “interest” state under the e-mail action as compared to the no-ad action. Moreover, the one-step conversion probabilities are highest for paid search click (5% to 10% depending on the state) followed by display impression (0.2% to 1.5%). As one would expect, the eventual conversion probability h_s^β increases monotonically with $s \in \mathbb{S}$, i.e., as a user goes deeper into the conversion funnel, he/she becomes more likely to buy the product. For instance, a user is around twice more likely to purchase if he/she is in the “desire” state as compared to the “unaware” state.

We also observe some estimates to be counterintuitive. For instance, our estimate of one-step conversion probability under the no-ad action is higher than under the e-mail action for all states except “unaware”. This might be a consequence of the endogeneity issue we alluded to earlier. Furthermore, for all three ad actions (e-mail, display, and paid search click), the one-step conversion probability from the “unaware” state is higher than in the “aware” state. We believe this is primarily a result of the user actually being acquainted with the product (via offline channels for instance).

6.3. Results and Discussion

Given the parameter estimates from real data, we estimate CASV using Algorithm 1, where we set $S = 100,000$ (number of paths to be sampled)[13]. On the $100,000$ sampled paths, it took less than a second to compute the attribution numbers for each scheme[14] on a single core of Intel Xeon E5-2650v2 2.6 GHz processor with 8 GB of RAM. This highlights the scalability of our approach. Figures 4 and 5 display the results. Interestingly, the five attribution schemes output quite different results and we discuss them next. We first discuss the percentage attributions to each action (Figure 4) to give the reader a high-level idea and then, comment on the state-specific attributions to each action (Figure 5).

In terms of attributions to actions (Figure 4), LTA allocates the majority of the share to e-mail and paid search click, which seems appropriate for LTA since e-mail ads are the most prevalent in our dataset and paid search click strongly hints a higher level of customer engagement (since it is a user action as opposed to an ad action) and hence, a relatively high one-step conversion probability. Since LTA does not adjust for the counterfactual, the no-ad action receives little credit. IVH allocates almost all the credit to e-mail, which aligns with the facts that IVH scales with action intensity (see (1)) and e-mail action occurs around 10 times more often than display and paid search click actions. We make a similar observation for uniform attribution. Due to its unique-uniform

[13] We repeated our experiment on multiple seeds and obtained almost identical results, indicating $S = 100,000$ is high enough. Furthermore, one can simply use the paths from the real-world dataset (instead of sampling). We sampled paths via Monte Carlo simulation to protect the private data of our provider.

[14] For LTA, IVH, and uniform, computing attribution is straightforward (see Section 3). To estimate SV, we used the unique-uniform characterization (Proposition 1).

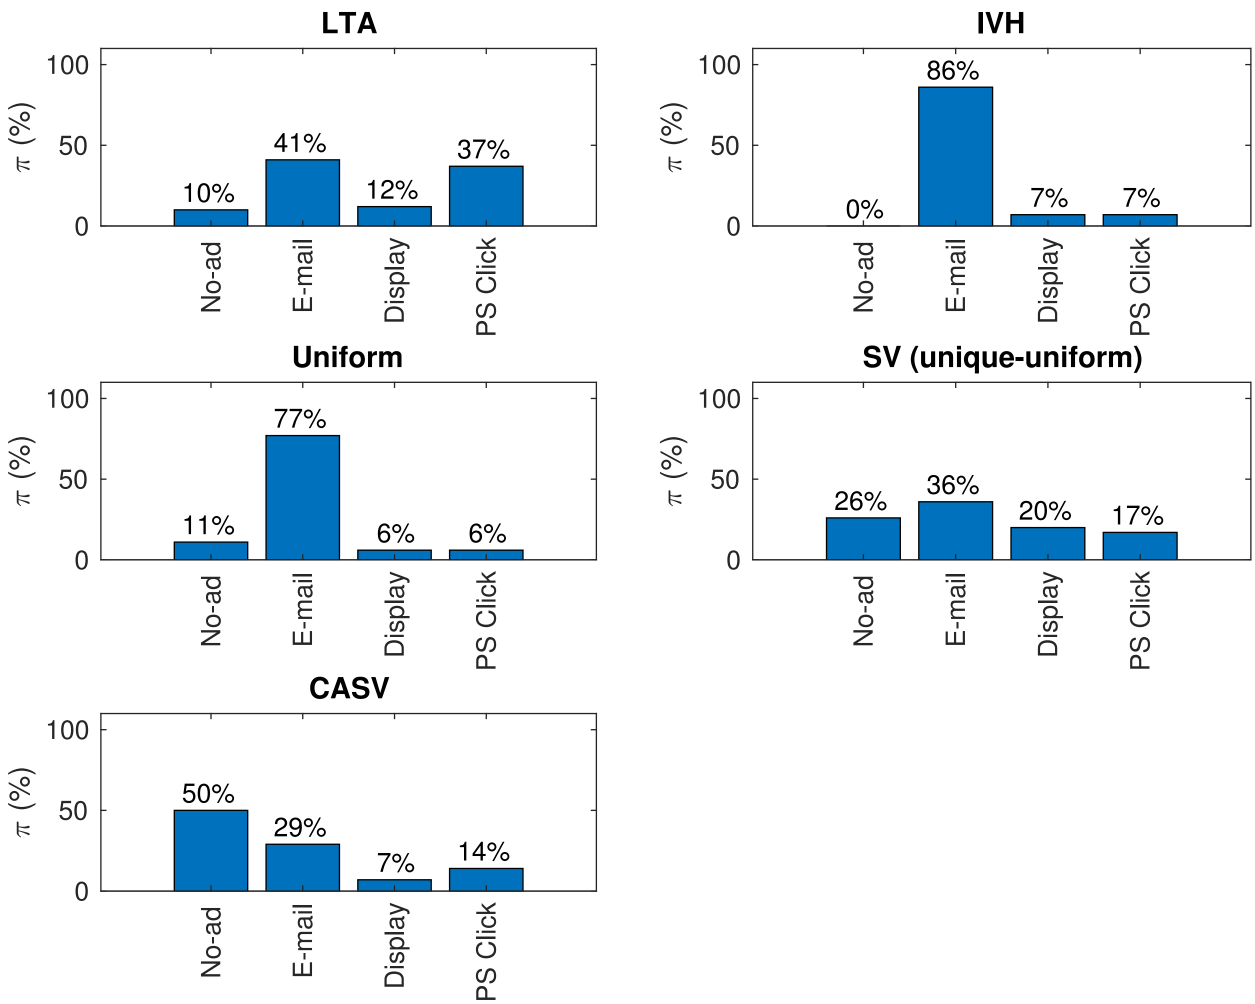

Figure 4 **Attributions to different actions under various schemes with $\tau = 10$. We report the percentage attributions to each action by aggregating over states, i.e., $\sum_s \pi_s^a / \sum_{(s',a')} \pi_{s'}^{a'}$ for each $a \in \mathbb{A}$.**

nature, SV accounts for the fact that players corresponding to e-mail action appear multiple times in certain paths and hence, the credit allocated to e-mail shrinks to roughly half (compared to uniform). This can be seen as a correction to other schemes (IVH and uniform in particular) allocating higher credit to e-mail than it deserves simply because it appears more often in the paths (even though it is "gaming the system" via self-loops). Compared with LTA, SV allocates less value to paid search click, which aligns with the intuition that LTA allocates more credit than appropriate to channels appearing later in the conversion funnel. Finally, CASV corrects SV by adjusting for the counterfactual. As expected, CASV takes away a portion from e-mail, display, and paid search click and allocates it to the no-ad counterfactual. This results in no-ad receiving roughly half of the total network value. The action that is least affected by the counterfactual adjustment is paid search click, which seems reasonable since replacing paid search click (a user

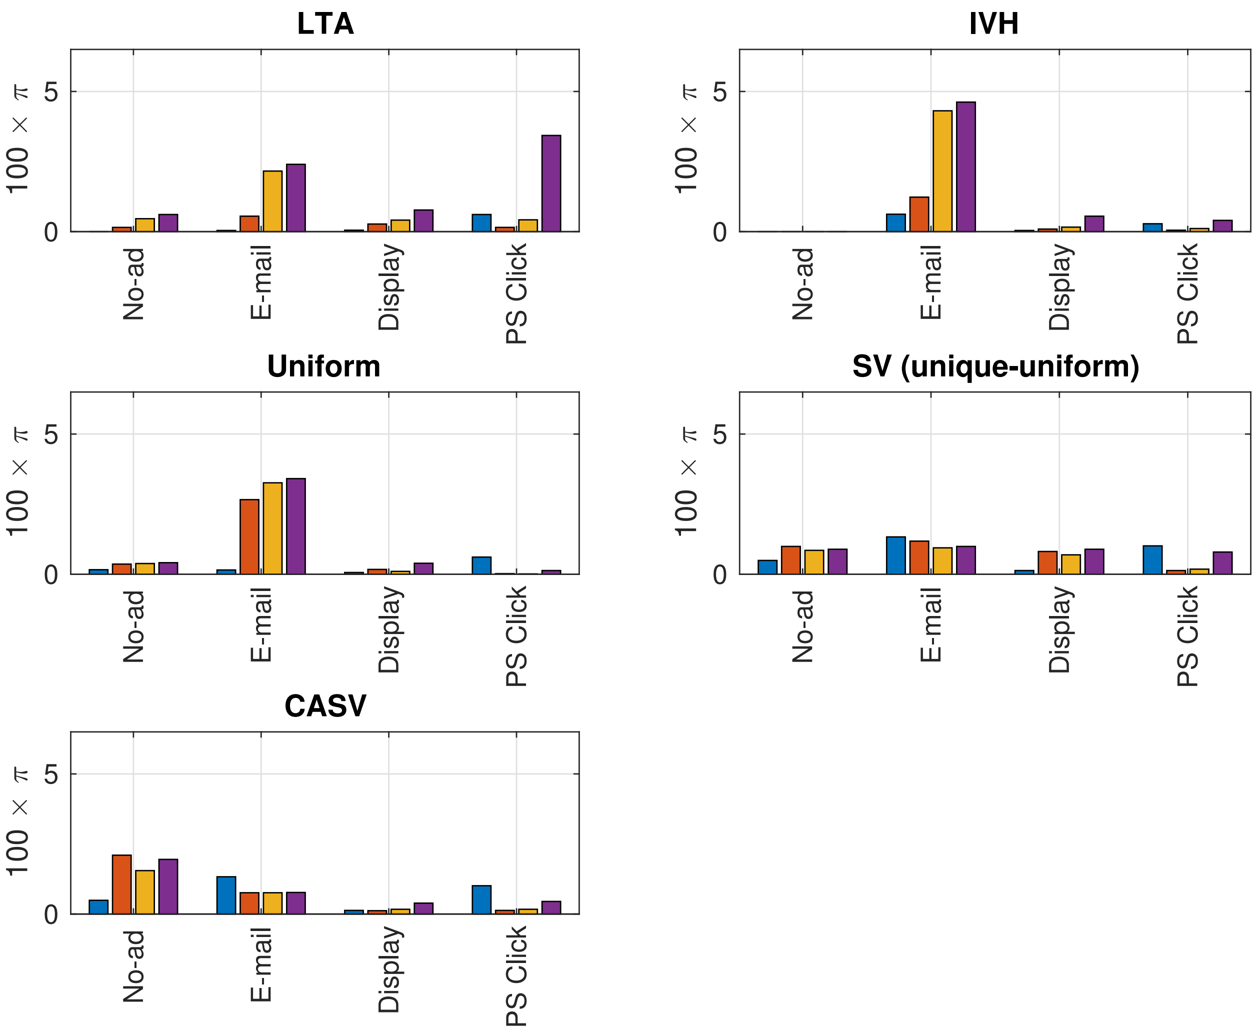

Figure 5 **Attributions to state-action pairs under various schemes with $\tau = 10$. We report the attributions π_s^a for all $(s, a) \in \mathbb{S} \times \mathbb{A}$. The four colors correspond to the four states in $\mathbb{S}$, arranged in the natural order (unaware, aware, interest, desire).**

action) by no-ad (an ad action) can create a big difference in total value of the network and hence, the counterfactual value would be small.

Having a high-level view of the attributions, we briefly discuss the state-action specific attributions (Figure 5). As is evident, each attribution scheme can output quite different allocations to various states for a given action. LTA allocates the most to paid search click at the "desire" state. Similar to the empirical finding of Blake et al. (2015), it is possible that a user in a state of "desire" is already inclined to buy the product and performs a Google search, leading to a paid search click and a conversion. However, in the counterfactual scenario, even if the paid search click did not occur, the user might have bought the product (by clicking on the "organic" link). In fact, the relatively low allocation to paid search click at the "desire" state under CASV provides empirical support to such a view. In terms of e-mail, existing heuristics such as LTA, IVH, and uniform have an increasing pattern in the attribution allocated with respect to the user state whereas CASV allocates the most to the first state. This also aligns with the intuition mentioned above, suggesting that the existing heuristics fail to account for the counterfactual appropriately. Interestingly, all the schemes allocate a considerable amount to paid search click at the "unaware" state. This might be a result of "unaware" users actually being aware of the product (from offline sources for instance) and searching the product (on Google for example) just to make a purchase but clicking the paid search ad in the process. In theory, CASV should adjust for this phenomenon by allocating the "offline value" to the no-ad counterfactual. However, this correction is not reflected in our numerics due to the hidden nature of data corresponding to paid search impressions (recall we only observe them when they are clicked).

We conclude this section by performing sensitivity analysis on τ. Recall that to implant the no-ad action in raw user paths, we used $\tau = 10$. We now experiment with a lower value of $\tau = 7$ (Figure 6) and a higher value of $\tau = 14$ (Figure 7) and analyze the resulting changes in the attribution numbers. It is easy to see that all the discussions we did in the case of $\tau = 10$ (Figure 4) still apply, indicating robustness of high-level insights to the choice of τ. We do see changes in the actual numbers but they align with our expectation. For example, with $\tau = 7$, we implant more instances of no-ad action (as compared to $\tau = 10$) and hence, expect it to receive more attribution (and vice versa with $\tau = 14$).

7. Value of state-specific attribution

In our cooperative game theory setup presented in Section 4, we defined each state-action pair $(s, a) \in \mathbb{S} \times \mathbb{A}$ to be a player (*state-specific model*) and hence, had CASV $\psi_s^{a,\text{Shap}}$ corresponding to each state-action pair. As we mentioned previously, a naive implementation to *exactly* compute $\{\psi_s^{a,\text{Shap}}\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$ using (5) would have an exponential runtime in the size of $\mathbb{S} \times \mathbb{A}$. Interestingly, there is an alternative view that reduces the computational complexity of a naive implementation to be exponential only in the size of $\mathbb{A}$, which might be tractable. In particular, one can define players in the cooperative game to be the actions in $\mathbb{A}$ (*aggregated model*) and compute CASV $\bar{\psi}^{a,\text{Shap}}$ for each action $a \in \mathbb{A}$. Under such a view, a coalition $\bar{\mathcal{X}}$ corresponds to a collection of actions, i.e., $\bar{\mathcal{X}} \subseteq \mathbb{A}$. To be more precise, we define the empty coalition in the same way as we defined earlier, i.e., it corresponds to the Markov chain containing all the states but action 0 being taken at all of them w.p. 1. If an action $a \in \mathbb{A}$ is added to the empty coalition, then the advertiser takes action a at *each* state $s \in \mathbb{S}$ w.p. β_s^a and so on. Under this aggregated model, each action $a \in \mathbb{A}$ receives an attribution of

$$\bar{\psi}^{a,\text{Shap}} := \sum_{\bar{\mathcal{X}} \subseteq \mathbb{A}\setminus\{a\}} \bar{w}_{|\bar{\mathcal{X}}|} \times \left\{ v(\bar{\mathcal{X}} \cup \{a\}) - v(\bar{\mathcal{X}} \cup \{1^a\}) \right\} \quad \text{where} \quad \bar{w}_{|\bar{\mathcal{X}}|} := \frac{|\bar{\mathcal{X}}|!(n - |\bar{\mathcal{X}}| - 1)!}{n!} \tag{9}$$

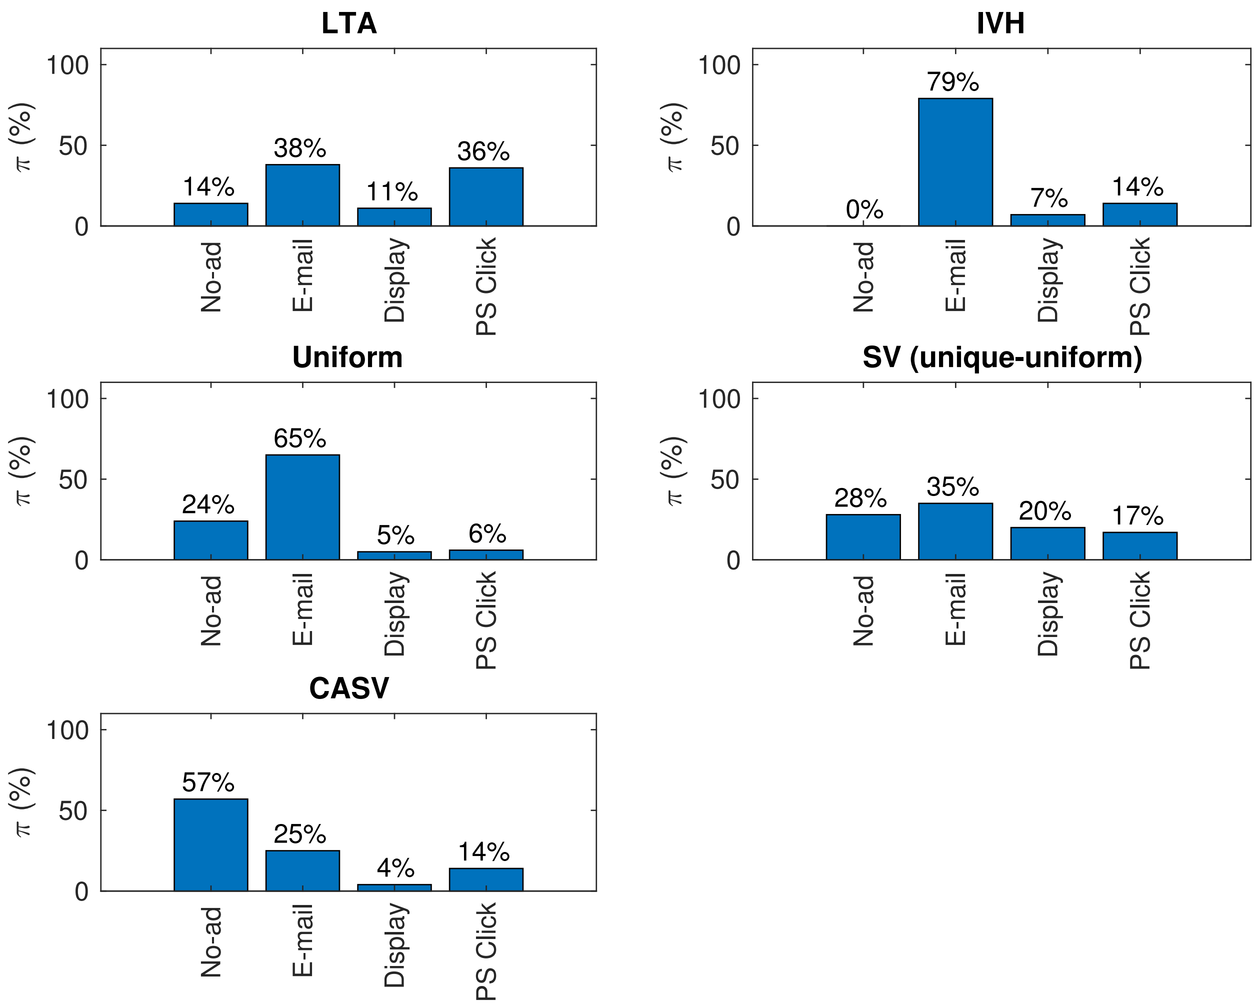

Figure 6 Attributions to different actions under various schemes with $\tau = 7$. We report the percentage attributions to each action by aggregating over states, i.e., $\sum_s \pi_s^a / \sum_{(s',a')} \pi_{s'}^{a'}$ for each $a \in \mathbb{A}$.

and 1^a denotes the counterfactual player, i.e., the no-ad action (action 1) is taken at state s w.p. β_s^a (instead of β_s^1) for all $s \in \mathbb{S}$. Similar to the state-specific model, CASV in the aggregated model can be expressed as $\bar{\psi}^{a,\text{Shap}}(\mathcal{M}) = \bar{\pi}^{a,\text{Shap}}(\mathcal{M}) - \bar{\pi}^{a,\text{Shap}}(\bar{\mathcal{M}}^a)$ where

$$\bar{\pi}^{a,\text{Shap}}(\mathcal{M}) := \sum_{\bar{\mathcal{X}} \subseteq \mathbb{A} \setminus \{a\}} \bar{w}_{|\bar{\mathcal{X}}|} \times \left\{ v_{\mathcal{M}}(\bar{\mathcal{X}} \cup \{a\}) - v_{\mathcal{M}}(\bar{\mathcal{X}}) \right\}, \tag{10}$$

$\mathcal{M}$ denotes the original Markov chain and $\bar{\mathcal{M}}^a$ denotes the counterfactual network for a (in an aggregated sense), i.e., it is identical to $\mathcal{M}$ except that we replace the transition probabilities of (s,a) by those of $(s,1)$ for all $s \in \mathbb{S}$. Furthermore, SV $\bar{\pi}^{a,\text{Shap}}$ for $a \in \mathbb{A}$ in the aggregated model can be characterized as a unique-uniform scheme using the same proof technique as for Proposition 1:

$$\bar{\pi}^{a,\text{Shap}} = \mathbb{E}_{\mathcal{P} \sim \mathcal{M}} \left[\bar{w}^a(\mathcal{P}) \right] \quad \text{where} \quad \bar{w}^a(\mathcal{P}) := \begin{cases} \frac{1}{\bar{u}(\mathcal{P})} & \text{if } \mathcal{P} \text{ converts and } a \in \bar{\mathcal{P}} \\ 0 & \text{otherwise.} \end{cases}$$

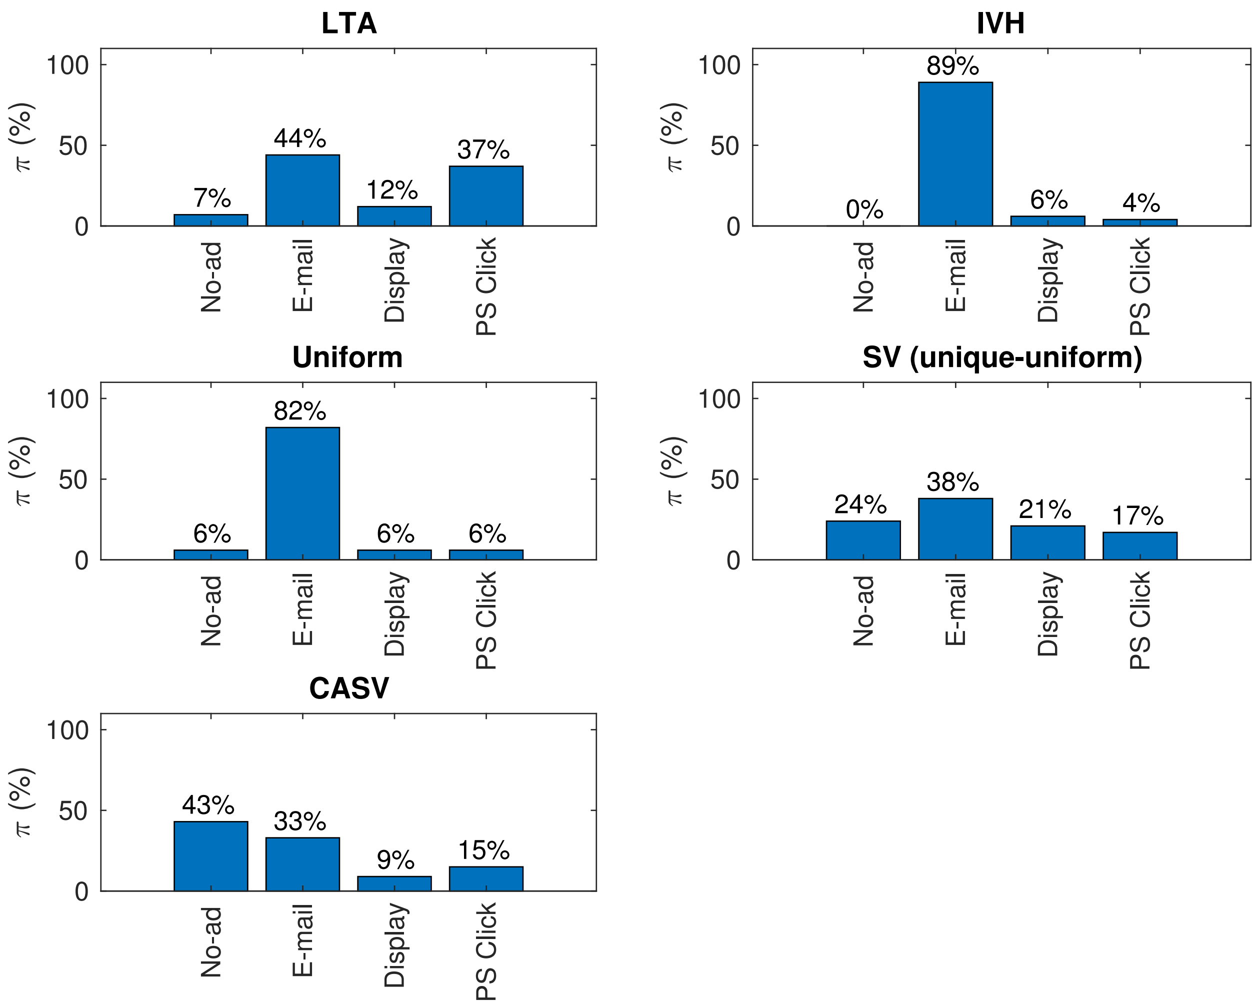

Figure 7 **Attributions to different actions under various schemes with $\tau = 14$. We report the percentage attributions to each action by aggregating over states, i.e., $\sum_s \pi_s^a / \sum_{(s',a')} \pi_{s'}^{a'}$ for each $a \in \mathbb{A}$.**

The function $\bar{u}(\cdot)$ returns the number of unique actions and $\bar{\mathcal{P}}$ denotes the actions in the path $\mathcal{P}$ (over state-action pairs), i.e., $\bar{\mathcal{P}} := \{a : (s,a) \in \mathcal{P}\}$. Thus, the CASV $\bar{\psi}^{a,\text{Shap}}$ for $a \in \mathbb{A}$ in the aggregated model equals

$$\bar{\psi}^{a,\text{Shap}} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[\bar{w}^a(\mathcal{P})\right] - \mathbb{E}_{\mathcal{P}\sim\bar{\mathcal{M}}^a}\left[\bar{w}^a(\mathcal{P})\right],$$

which is analogous to Theorem 2.

If the size of $\mathbb{A}$ is relatively small (which might be the case for certain retailers), CASV can be computed *exactly* in the aggregated model using (9). Accordingly, it is of interest to check whether using the granular state-specific model (which is computationally demanding) provides some quantifiable value over the aggregated model. However, one needs to be careful when comparing the attributions of two models as they are on different "scales". The state-specific model outputs attribution at a state-action level whereas the aggregated model outputs at an action level. Furthermore, decomposing the action level allocation of the aggregated model to the state-action level seems non-trivial. On the other hand, aggregating the state-action level allocations of the state-specific model is straightforward ($\sum_{s\in\mathbb{S}}\psi_s^{a,\text{Shap}}$). Therefore, we will compare $\sum_{s\in\mathbb{S}}\psi_s^{a,\text{Shap}}$ to $\bar{\psi}^{a,\text{Shap}}$ to answer the following question: does computing state-specific attribution and then aggregating over states provide a different answer than directly computing attribution that is not state-specific? Clearly, this is a stronger notion of comparison in the sense that if $\sum_{s\in\mathbb{S}}\psi_s^{a,\text{Shap}}$ differs from $\bar{\psi}^{a,\text{Shap}}$, then any state-specific decomposition of $\bar{\psi}^{a,\text{Shap}}$ would be "inappropriate".

Our unique-uniform characterizations imply that

$$\sum_{s\in\mathbb{S}}\psi_s^{a,\text{Shap}} - \bar{\psi}^{a,\text{Shap}} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[\sum_{s\in\mathbb{S}} w_s^a(\mathcal{P}) - \bar{w}^a(\mathcal{P})\right] - \left[\sum_{s\in\mathbb{S}}\mathbb{E}_{\mathcal{P}\sim\mathcal{M}_s^a}\left[w_s^a(\mathcal{P})\right] - \mathbb{E}_{\mathcal{P}\sim\bar{\mathcal{M}}^a}\left[\bar{w}^a(\mathcal{P})\right]\right],$$

which gives a path-based view of the difference between the two approaches. For instance, in a network with zero counterfactual value, if a path contains x unique state-action pairs with $x-1$ of them corresponding to the same action, the state-specific weight function allocates $1/x$ to each state-action pair and hence, the repeated action receives an attribution of $(x-1)/x$ to account for its state-dependent impact. On the contrary, the aggregated model allocates $1/2$ to both the unique actions. Though useful, such a path-based view does not build an understanding in terms of the type of interactions among players captured (and missed) by the two approaches. To facilitate such a fundamental understanding, we define two types of coalitions (assuming the underlying players are state-action pairs).

Type 1 coalition. In a type 1 coalition (say $\mathcal{X}_1$), if a player (s,a) is in $\mathcal{X}_1$, then the player (s',a) is in $\mathcal{X}_1$ for all $s'\in\mathbb{S}$. In other words, if an action is present at a single state, then it is present at all the states. We denote by $\mathbb{T}_1$ the set of all type 1 coalitions. Type 1 coalitions correspond to the coalitions in the aggregated model. Intuitively, both the aggregated and the state-specific models capture the interactions present in type 1 coalitions. However, it is not clear whether the "weights" assigned to each of these interactions are the same in the two models. (We shed more light on this issue in Theorem 3.)

Type 2 coalition. Any coalition that is a subset of $\mathbb{S}\times\mathbb{A}$ but not a type 1 coalition is defined as a type 2 coalition. In other words, type 2 coalitions have at least one action that is present at one or more states but not at all the states. We define $\mathbb{T}_2$ as the collection of all type 2 coalitions. It should be intuitively clear from (9) that the aggregated model will miss all the interactions present in type 2 coalitions. (We formalize this intuition in Theorem 3.)

We are now equipped to precisely quantify the value of state-specific attribution in terms of the interactions among the underlying players and we do so in Theorem 3, the proof of which is presented in Appendix C.

Theorem 3 **(Value of state-specific attribution)**. *For any action, computing state-specific CASV and then aggregating over states is not the same as computing CASV over the aggregated model. In particular, for all $a \in \mathbb{A}$,*

$$\sum_{s\in\mathbb{S}} \psi_s^{a,Shap} - \bar{\psi}^{a,Shap} = \sum_{\mathcal{X}\in\mathbb{T}_1} \{c^a(\mathcal{X}) - \bar{c}^a(\mathcal{X})\}\{v_{\mathcal{M}}(\mathcal{X}) - v_{\bar{\mathcal{M}}^a}(\mathcal{X})\} + \sum_{\mathcal{X}\in\mathbb{T}_2}\sum_{s\in\mathbb{S}} c_s^a(\mathcal{X})\left\{v_{\mathcal{M}}(\mathcal{X}) - v_{\mathcal{M}_s^a}(\mathcal{X})\right\},$$

where $c^a(\mathcal{X}) := \sum_{s\in\mathbb{S}} c_s^a(\mathcal{X})$,

$$c_s^a(\mathcal{X}) := \begin{cases} w_{|\mathcal{X}|-1} & \text{if } (s,a)\in\mathcal{X} \\ -w_{|\mathcal{X}|} & \text{otherwise} \end{cases} \quad \text{and} \quad \bar{c}^a(\mathcal{X}) := \begin{cases} \bar{w}_{|\bar{\mathcal{X}}|-1} & \text{if } a\in\bar{\mathcal{X}} \\ -\bar{w}_{|\bar{\mathcal{X}}|} & \text{otherwise.} \end{cases}$$

The notation $\bar{\mathcal{X}}$ denotes the actions in $\mathcal{X}$, i.e., $\bar{\mathcal{X}} := \{a : (s,a)\in\mathcal{X}\}$.

We now parse the mathematical result presented above. The first term in the difference corresponds to type 1 coalitions. Such coalitions are captured by both the models. However, it is not true that the two models assign equal weights to them ($c^a(\mathcal{X})$ vs. $\bar{c}^a(\mathcal{X})$). On the other hand, the second term corresponds to type 2 coalitions, which the aggregated model completely disregards (note that there is no $\bar{c}$ weight in the second term). Hence, the aggregated model misses a multitude of interactions that are captured by the state-specific model.

Remark 3. It is easy to show that $c^a(\mathcal{X})$ for any coalition $\mathcal{X}\subseteq\mathbb{S}\times\mathbb{A}$ and action $a\in\mathbb{A}$ equals

$$c^a(\mathcal{X}) = m^a(\mathcal{X})w_{|\mathcal{X}|-1} - (m - m^a(\mathcal{X}))w_{|\mathcal{X}|},$$

where the scalar $m^a(\mathcal{X})$ equals the number of players in the set $\{(s,a)\}_{s\in\mathbb{S}}$ that exist in the coalition $\mathcal{X}$. Furthermore, for $\mathcal{X}\in\mathbb{T}_1$, we have

$$m^a(\mathcal{X}) = \begin{cases} m & \text{if } a\in\bar{\mathcal{X}} \\ 0 & \text{otherwise.} \end{cases}$$

As a result, the coefficient in the first term simplifies to

$$c^a(\mathcal{X}) - \bar{c}^a(\mathcal{X}) = \begin{cases} mw_{|\mathcal{X}|-1} - \bar{w}_{|\bar{\mathcal{X}}|-1} & \text{if } a\in\bar{\mathcal{X}} \\ \bar{w}_{|\bar{\mathcal{X}}|} - mw_{|\mathcal{X}|} & \text{otherwise.} \end{cases}$$

Having quantified the difference between the state-specific and aggregated models, we show that the ratio of the two models can be arbitrarily large, which accentuates the importance of state-specific attribution.

Proposition 2. *The ratio of the aggregated model's attribution to the aggregation of the state-specific model's attribution can be arbitrarily large. Mathematically,*

$$\sup_{\mathcal{M}}\sup_{a\in\mathbb{A}} \frac{\bar{\psi}^{a,Shap}(\mathcal{M})}{\sum_{s\in\mathbb{S}}\psi_s^{a,Shap}(\mathcal{M})} = \infty.$$

Proof. It suffices to show the existence of one such instance of $(\mathcal{M}, a)$. Consider the network in Figure 8. In the aggregated model, the two players ($a = 1$ and $a = 2$) are symmetric since a user can not convert if either one is absent. Accordingly, for $a \in \{1, 2\}$, $\bar{\psi}^{a,\text{Shap}} = 1/2$. On the other hand, in the state-specific model, players $(1, 2)$ and $\{(s, 1)\}_{s=2}^{m}$ are null players and the remaining m players are symmetric and hence, each of them receives an attribution equal to the total value in the system divided by m. Accordingly, $\sum_{s\in\mathbb{S}} \psi_s^{1,\text{Shap}} = 1/m$ and $\sum_{s\in\mathbb{S}} \psi_s^{2,\text{Shap}} = (m-1)/m$. Clearly, $\lim_{m\to\infty} \bar{\psi}^{1,\text{Shap}} / \sum_{s\in\mathbb{S}} \psi_s^{1,\text{Shap}} = \infty$, which concludes the proof. $\square$

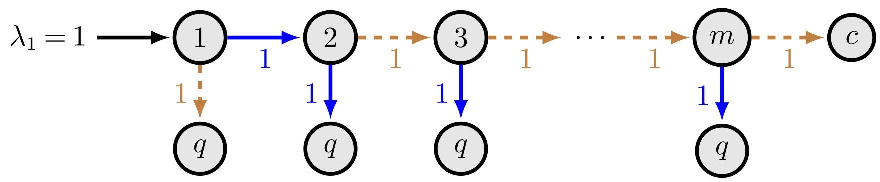

Figure 8 **Network for the proof of Proposition 2. The action space consists of three actions: show no-ad, show ad 1, and show ad 2. Solid blue (dashed brown) lines denote transition if an ad 1 (ad 2) is shown. At state 1, taking action 1 moves the traffic to state 2 w.p. 1 whereas taking action 2 directs the traffic to quit state. At state $s \in \{2, \ldots, m\}$, taking action 1 results in a transition to the quit state whereas action 2 moves the users to the "next" state. For brevity, we do not show the transitions if an ad is not shown (to quit state w.p. 1). The advertiser shows ad 1 at state 1 w.p. 1 and ad 2 at all other states w.p. 1.**

We conclude this section by showing Figure 9, which displays the difference between the two models (state-specific and aggregated) on the simulated dataset from Section 6. The aggregated model underestimates the contributions of no-ad and e-mail actions by around 10%.

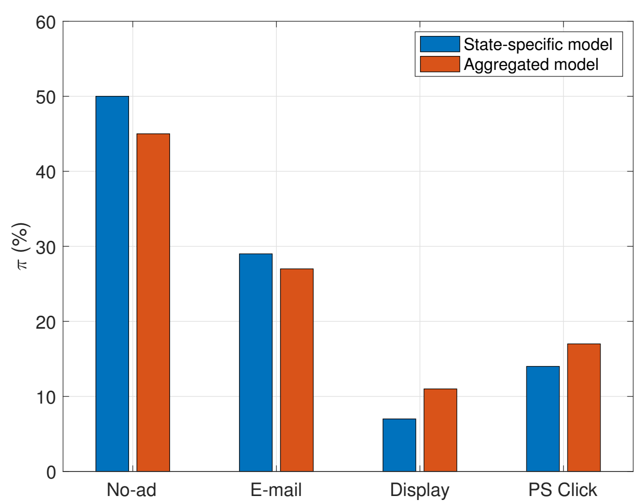

Figure 9 **Value of state-specific attribution on the simulated dataset. For the state-specific model, we report the percentage attributions to each action by aggregating over states, i.e., $\sum_s \psi_s^{a,\text{Shap}} / \sum_{(s',a')} \psi_{s'}^{a',\text{Shap}}$ for each $a \in \mathbb{A}$. For the aggregated model, we report $\bar{\psi}^{a,\text{Shap}}$ for each $a \in \mathbb{A}$.**

8. Conclusions and further research

In this paper, using a Markovian model for user behavior, we propose a new metric (namely, counterfactual adjusted Shapley value) for the attribution problem in online advertising. We establish its theoretical foundations and appropriateness in two ways. First, we show the robustness of the proposed measure over various canonical settings in which the existing metrics fail. Second, we provide an underlying axiomatic framework motivated by game theory and causality that supports our choice. Furthermore, we establish a characterization of the proposed metric as a remarkably simple function of the user paths. In particular, we show that in our Markovian model, the proposed attribution metric is an adjustment to the unique-uniform attribution scheme. Finally, we propose multiple approaches to estimate our metric and show its scalability through numerical experiments on a real-world large-scale dataset, in addition to benchmarking it against existing metrics.

We think there are many potential future directions to build upon our proposed framework. It is of interest to develop more canonical settings to further verify the appropriateness of our proposed measure. In particular, while the axiomatic framework provides theoretical justifications for our proposed metric, we still do not have objective measures to compare against the existing approaches. It would also be worthwhile to conduct a rigorous empirical study that addresses the data/estimation issues we have discussed in this work and then, compare various attribution schemes. Since these issues are universal to all the attribution schemes, we believe such a work is of independent interest but outside the scope of this paper.

In this work, we implicitly assume that the various ad actions/channels cooperate with each other towards the common goal of maximizing network value instead of being strategic and maximizing their individual values. Motivated by the works of Berman (2018) and Abhishek et al. (2017), it would be interesting to analyze our framework when the individual channels are strategic. At a high-level, it is unclear if a *cooperative* game theory framework is even appropriate under this strategic setting. For instance, the *unique*-uniform scheme provides each channel an incentive to pose as two smaller but distinct channels.

Furthermore, our framework could be of independent interest in other domains. For instance, there has been a growing interest in the field of “interpretable machine learning”, where the goal is to attribute the output of a prediction model to the individual features of the input (Baehrens et al. 2010, Sundararajan et al. 2017, Lundberg and Lee 2017) or to the components of the prediction model (Dhamdhere et al. 2018, Leino et al. 2018). It is of interest to see if the ideas presented in this work can be used to answer such questions.

It is also worthwhile to explore whether attribution metrics can be constructed from other notions of causality. In this work, we modify Shapley value using Rubin’s definition of causality Rubin (1974). However, there exist several alternative formulations for causality. See, e.g. Pearl

(2009), Halpern and Pearl (2005), Chockler and Halpern (2004), Hitchcock (1997), Morgan and Winship (2014), Collins et al. (2004), Eells (1991) and Hume (2003) for approaches to causality in the computer science and philosophy literature. Chockler and Halpern (2004) proposed metrics such as "degree of responsibility" and "blame" in order to *quantify* the causal effect of one variable on another. However, computing these metrics is computationally expensive. Furthermore, these metrics lack axiomatic support. Also, similar to IVH, these metrics are not budget balanced. In fact, Chockler and Halpern (2004) mention adjusting their notions using Shapley value as a potential future work. Accordingly, though we have unified the existing attribution literature in this work, there is more to explore in terms of the set of counterfactuals to consider.

Acknowledgments

We acknowledge the support of Adobe Digital Marketing Research Award for this project. V. Goyal also acknowledges the support of NSF grants CMMI 1636046 and CMMI 1351838 for this work. We also thank Abhishek Pani, Chen Dong, Manoj Ravi, and Jie Zhang from the Media and Advertising Solutions team at Adobe for useful discussions during the course of this work. Finally, we thank the WWW reviewers for the valuable feedback and pointing us to the responsibility and blame literature.

References

Abhishek V, Despotakis S, Ravi R (2017) Multi-channel attribution: The blind spot of online advertising. Working paper, available at: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2959778` .

Abhishek V, Fader P, Hosanagar K (2012) Media exposure through the funnel: A model of multi-stage attribution. Working paper, available at: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2158421` .

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. Working paper, available at: `https://arxiv.org/abs/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 (ACM).

Arensman J, Yeung W (2016) Move beyond last click attribution in adwords. URL `https://adwords.googleblog.com/2016/05/move-beyond-last-click-attribution.html`.

Avrachenkov K, Cottatellucci L, Maggi L (2012) Confidence intervals for Shapley value in Markovian dynamic games. Technical report (Eurecom), available at: `http://www.eurecom.fr/publication/3602` .

Baardman L, Fata E, Pani A, Perakis G (2019) Learning optimal online advertising portfolios with periodic budgets. Working paper, available at: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3346642` .

Baehrens D, Schroeter T, Harmeling S, Kawanabe M, Hansen K, MÃžller KR (2010) How to explain individual classification decisions. *Journal of Machine Learning Research* 11(Jun):1803–1831.

Balseiro S, Gur Y (2017) Learning in repeated auctions with budgets: Regret minimization and equilibrium. Working paper, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2921446 .

Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61(4):864–884.

Balseiro SR, Feldman J, Mirrokni V, Muthukrishnan S (2014) Yield optimization of display advertising with ad exchange. *Management Science* 60(12):2886–2907.

Barry TE (1987) The development of the hierarchy of effects: An historical perspective. *Current issues and Research in Advertising* 10(1-2):251–295.

Berman R (2018) Beyond the last touch: Attribution in online advertising. *Marketing Science* 37(5):771–792.

Bettman JR, Luce MF, Payne JW (1998) Constructive consumer choice processes. *Journal of Consumer Research* 25(3):187–217.

Blake T, Nosko C, Tadelis S (2015) Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica* 83(1):155–174.

Bleier A, Eisenbeiss M (2015) Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science* 34(5):669–688.

Bottou L, Peters J, Quiñonero-Candela J, Charles DX, Chickering DM, Portugaly E, Ray D, Simard P, Snelson E (2013) Counterfactual reasoning and learning systems: The example of computational advertising. *The Journal of Machine Learning Research* 14(1):3207–3260.

Bruce NI, Peters K, Naik PA (2012) Discovering how advertising grows sales and builds brands. *Journal of Marketing Research* 49(6):793–806.

Castro J, Gómez D, Tejada J (2009) Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research* 36(5):1726–1730.

Chockler H, Halpern JY (2004) Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research* 22:93–115.

Choi H, Mela C, Balseiro S, Leary A (2017) Online display advertising markets: A literature review and future directions. Working paper, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3070706 .

Collins JD, Hall EJ, Hall N, Paul LA (2004) *Causation and counterfactuals* (MIT Press).

Court D (2009) The consumer decision journey. *McKinsey Quarterly* .

Dalessandro B, Perlich C, Stitelman O, Provost F (2012) Causally motivated attribution for online advertising. *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*, 7 (ACM).

Danaher PJ, van Heerde HJ (2018) Delusion in attribution: Caveats in using attribution for multimedia budget allocation. *Journal of Marketing Research* 55(5):667–685.

Dhamdhere K, Sundararajan M, Yan Q (2018) How important is a neuron? Working paper, available at: `https://arxiv.org/abs/1805.12233` .

Eells E (1991) *Probabilistic causality*, volume 1 (Cambridge University Press).

Elzinga D, Mulder S, Vetvik OJ, et al. (2009) The consumer decision journey. *McKinsey Quarterly* 3:96–107.

Fatima SS, Wooldridge M, Jennings NR (2008) A linear approximation method for the Shapley value. *Artificial Intelligence* 172(14):1673–1699.

Grinstead CM, Snell JL (2012) *Introduction to probability* (American Mathematical Society).

Hagan S (2018) Digital economy has been growing at triple the pace of U.S. GDP. URL `https://www.bloomberg.com/news/articles/2018-03-15/digital-economy-has-been-growing-at-triple-the-pace-of-u-s-gdp`.

Halpern JY, Pearl J (2005) Causes and explanations: A structural-model approach. Part I: Causes. *The British Journal for the Philosophy of Science* 56(4):843–887.

Hitchcock C (1997) Probabilistic causation. URL `https://plato.stanford.edu/entries/causation-probabilistic/`.

Hojjat A, Turner J, Cetintas S, Yang J (2017) A unified framework for the scheduling of guaranteed targeted display advertising under reach and frequency requirements. *Operations Research* 65(2):289–313.

Howard JA, Sheth JN (1969) *The theory of buyer behavior* (John Wiley & Sons).

Hume D (2003) *A treatise of human nature* (Courier Corporation).

Institute MS (2016) Research priorities 2016-2018. URL `https://www.msi.org/uploads/articles/MSI_RP16-18.pdf`.

Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. *Management Science* 60(12):2949–2970.

Jansen BJ, Schuster S (2011) Bidding on the buying funnel for sponsored search and keyword advertising. *Journal of Electronic Commerce Research* 12(1):1.

Ji W, Wang X (2017) Additional multi-touch attribution for online advertising. *AAAI*, 1360–1366.

Ji W, Wang X, Zhang D (2016) A probabilistic multi-touch attribution model for online advertising. *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 1373–1382 (ACM).

Jordan P, Mahdian M, Vassilvitskii S, Vee E (2011) The multiple attribution problem in pay-per-conversion advertising. *International Symposium on Algorithmic Game Theory*, 31–43 (Springer).

Kakalejčík L, Bucko J, Resende PA, Ferencova M (2018) Multichannel marketing attribution using markov chains. *Journal of Applied Management and Investments* 7(1):49–60.

Kannan P, Reinartz W, Verhoef PC (2016) The path to purchase and attribution modeling: Introduction to special section. *International Journal of Research in Marketing* 33(3):449–456.

Kireyev P, Pauwels K, Gupta S (2016) Do display ads influence search? Attribution and dynamics in online advertising. *International Journal of Research in Marketing* 33(3):475–490.

Kotler P, Armstrong G (2010) *Principles of marketing* (Pearson Education).

Leino K, Sen S, Datta A, Fredrikson M, Li L (2018) Influence-directed explanations for deep convolutional networks. *2018 IEEE International Test Conference (ITC)*, 1–8 (IEEE).

Lejeune M, Turner J (2019) Planning online advertising using Lorenz curves. *Operations Research (In press)* .

Li H, Kannan P (2014) Attributing conversions in a multichannel online marketing environment: An empirical model and a field experiment. *Journal of Marketing Research* 51(1):40–56.

Li Y, Xie Y, Zheng E (2017) Modeling multi-channel advertising attribution across competitors. Working paper, available at: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3047981` .

Liben-Nowell D, Sharp A, Wexler T, Woods K (2012) Computing Shapley value in supermodular coalitional games. *International Computing and Combinatorics Conference*, 568–579 (Springer).

Littlechild SC, Owen G (1973) A simple expression for the Shapley value in a special case. *Management Science* 20(3):370–372.

Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 4765–4774.

Maleki S, Tran-Thanh L, Hines G, Rahwan T, Rogers A (2013) Bounding the estimation error of sampling-based Shapley value approximation. Working paper, available at: `https://arxiv.org/abs/1306.4265` .

Michalak TP, Aadithya KV, Szczepanski PL, Ravindran B, Jennings NR (2013) Efficient computation of the Shapley value for game-theoretic network centrality. *Journal of Artificial Intelligence Research* 46:607–650.

Morgan SL, Winship C (2014) *Counterfactuals and causal inference* (Cambridge University Press).

Mulpuru S (2011) The purchase path of online buyers. *Forrester Report* 5:12.

Owen G (1972) Multilinear extensions of games. *Management Science* 18(5-part-2):64–79.

Pearl J (2009) *Causality* (Cambridge University Press).

Priest C (2017) Multichannel marketing attribution with datarobot. URL `https://www.datarobot.com/resource/mktattribution/`.

Quantcast (2013) Beyond last touch: Understanding campaign effectiveness. URL `http://info.quantcast.com/rs/quantcast/images/Quantcast%20White%20Paper_Beyond%20Last%20Touch.pdf`.

Quantcast (2016) Guide to marketing attribution. URL `https://www.iabuk.com/sites/default/files/case-study-docs/Quantcast%20-%20Guide%20to%20Marketing%20Attribution%20-%20EN%20-%202016%20-%20print.pdf`.

Rubin DB (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66(5):688.

Shao X, Li L (2011) Data-driven multi-touch attribution models. *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 258–264 (ACM).

Shapley LS (1953) A value for n-person games. *Contributions to the Theory of Games* 2(28):307–317.

Sluis S (2018) Digital ad market soars to $88 billion, Facebook and Google contribute 90% of growth. URL `https://adexchanger.com/online-advertising/digital-ad-market-soars-to-88-billion-facebook-and-google-contribute-90-of-growth/`.

Strong EK (1925) *The psychology of selling and advertising* (McGraw-Hill Book Company, Incorporated).

Sundararajan M, Taly A, Yan Q (2017) Axiomatic attribution for deep networks. *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 3319–3328.

United EDM (2012) Quarterly digital intelligence briefing: Making sense of marketing attribution (in association with adobe). URL `http://success.adobe.com/assets/en/downloads/whitepaper/Adobe-Quarterly-Digital-Intelligence-Briefing-Digital-Trends-for-2013.pdf`.

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. *2014 IEEE International Conference on Data Mining (ICDM)*, 687–696 (IEEE).

Zhao K, Mahboobi SH, Bagheri SR (2018) Revenue-based attribution modeling for online advertising. *International Journal of Market Research* 61(2):195–209.

Appendix A: Proof of Theorem 1

Theorem 1. *CASV satisfies the following axioms:*

*1. **Counterfactual efficiency**: The sum of CASVs equals the additional value generated over the counterfactual value, i.e.,*

$$\sum_{(s,a)\in\mathbb{S}\times\mathbb{A}} \psi_s^{a,Shap} = v(\mathbb{S}\times\mathbb{A}) - \sum_{(s,a)\in\mathbb{S}\times\mathbb{A}} \pi_s^{a,Shap}(\mathcal{M}_s^a).$$

*2. **Counterfactual symmetry**: If $(s,a)\in\mathbb{S}\times\mathbb{A}$ and $(s',a')\in\mathbb{S}\times\mathbb{A}$ are counterfactual equivalent, i.e.,*

(A) $v(\mathcal{X}\cup\{(s,a)\}) - v(\mathcal{X}\cup\{(s,1)^a\}) = v(\mathcal{X}\cup\{(s',a')\}) - v(\mathcal{X}\cup\{(s',1)^{a'}\})$ and

(B) $v(\mathcal{X}\cup\{(s,1)^a,(s',a')\}) = v(\mathcal{X}\cup\{(s',1)^{a'},(s,a)\})$

for all $\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}$, then

$$\psi_s^{a,Shap} = \psi_{s'}^{a',Shap}.$$

*3. **Linearity**: Consider two characteristic functions $v(\cdot)$ and $w(\cdot)$. For all $(s,a)\in\mathbb{S}\times\mathbb{A}$, we have*

$$\psi_s^{a,Shap}(v+w) = \psi_s^{a,Shap}(v) + \psi_s^{a,Shap}(w)$$

and for all $\alpha\in\mathbb{R}$,

$$\psi_s^{a,Shap}(\alpha v) = \alpha\psi_s^{a,Shap}(v).$$

*4. **Counterfactual null player**: Consider a player $(s,a)\in\mathbb{S}\times\mathbb{A}$ that has a zero value-add to all coalitions that do not contain (s,a), i.e., for all $\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a)\}$,*

$$v(\mathcal{X}\cup\{(s,a)\}) = v(\mathcal{X}\cup\{(s,1)^a\}).$$

Then, $\psi_s^{a,Shap}=0$.

Furthermore, CASV is the unique solution to satisfy these four counterfactual axioms.

Proof. Counterfactual efficiency and linearity follow from (6) when used with the efficiency and linearity of SV, respectively. Counterfactual null player follows from (5). For counterfactual symmetry, consider $(s,a)\in\mathbb{S}\times\mathbb{A}$ and $(s',a')\in\mathbb{S}\times\mathbb{A}$ satisfying (A) and (B) and observe that

$$\begin{aligned}
\psi_s^{a,\text{Shap}} &= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a)\}} w_{|\mathcal{X}|}\times\{v(\mathcal{X}\cup\{(s,a)\}) - v(\mathcal{X}\cup\{(s,1)^a\})\} \\
&= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|}\times\{v(\mathcal{X}\cup\{(s,a)\}) - v(\mathcal{X}\cup\{(s,1)^a\})\} \\
&+ \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|+1}\times\{v(\mathcal{X}\cup\{(s,a),(s',a')\}) - v(\mathcal{X}\cup\{(s,1)^a,(s',a')\})\} \\
&= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|}\times\Big\{v(\mathcal{X}\cup\{(s',a')\}) - v(\mathcal{X}\cup\{(s',1)^{a'}\})\Big\} \\
&+ \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|+1}\times\Big\{v(\mathcal{X}\cup\{(s,a),(s',a')\}) - v(\mathcal{X}\cup\{(s',1)^{a'},(s,a)\})\Big\} \\
&= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s',a')\}} w_{|\mathcal{X}|}\times\Big\{v(\mathcal{X}\cup\{(s',a')\}) - v(\mathcal{X}\cup\{(s',1)^{a'}\})\Big\} \\
&= \psi_{s'}^{a',\text{Shap}}.
\end{aligned}$$

To show uniqueness, consider $\{\psi_s^a\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$ such that it satisfies the counterfactual axioms. Motivated by (6), express ψ_s^a as

$$\psi_s^a = \underbrace{\psi_s^a + \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a)}_{=:\pi_s^a} - \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a)$$

for all $(s,a)\in\mathbb{S}\times\mathbb{A}$. To show $\psi_s^a = \psi_s^{a,\text{Shap}}$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$, it suffices to show that $\pi_s^a = \pi_s^{a,\text{Shap}}$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$. We do so by proving $\{\pi_s^a\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$ satisfies the four desirable properties of SV from Section 4.1 (recall $\{\pi_s^{a,\text{Shap}}\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$ is the unique solution to those properties).

Efficiency: Since $\pi_s^a = \psi_s^a + \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a)$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$ and $\{\psi_s^a\}_{(s,a)\in\mathbb{S}\times\mathbb{A}}$ satisfies counterfactaul efficiency, we get

$$\sum_{(s,a)\in\mathbb{S}\times\mathbb{A}} \pi_s^a = \sum_{(s,a)\in\mathbb{S}\times\mathbb{A}} \psi_s^a + \sum_{(s,a)\in\mathbb{S}\times\mathbb{A}} \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) = v(\mathbb{S}\times\mathbb{A}).$$

Symmetry: Consider $(s,a)\in\mathbb{S}\times\mathbb{A}$ and $(s',a')\in\mathbb{S}\times\mathbb{A}$. We need to show that

$$\underbrace{v(\mathcal{X}\cup\{(s,a)\}) = v(\mathcal{X}\cup\{(s',a')\})\ \forall\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}}_{(\diamond)} \implies \pi_s^a = \pi_{s'}^{a'}.$$

We can use the fact that $\{\psi_s^a, \psi_{s'}^{a'}\}$ satisfy counterfactual symmetry, i.e., if (A) and (B) hold for all $\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}$, then

$$\pi_s^a - \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) = \pi_{s'}^{a'} - \pi_{s'}^{a',\text{Shap}}(\mathcal{M}_{s'}^{a'}).$$

For the purposes of a contradiction, suppose that $(\diamond)$ holds but $\pi_s^a \neq \pi_{s'}^{a'}$. It suffices to show that this results in $\{\psi_s^a, \psi_{s'}^{a'}\}$ violating counterfactual symmetry. Suppose (A) and (B) hold. Given $(\diamond)$, statement (A) implies

$$v(\mathcal{X}\cup\{(s,1)^a\}) = v(\mathcal{X}\cup\{(s',1)^{a'}\})\ \forall\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}.$$

Together with (B), this implies $\pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) = \pi_{s'}^{a',\text{Shap}}(\mathcal{M}_{s'}^{a'})$:

$$\begin{aligned}
\pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) &= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a)\}} w_{|\mathcal{X}|}\times\{v(\mathcal{X}\cup\{(s,1)^a\}) - v(\mathcal{X})\}\\
&= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|}\times\{v(\mathcal{X}\cup\{(s,1)^a\}) - v(\mathcal{X})\}\\
&+ \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|+1}\times\{v(\mathcal{X}\cup\{(s,1)^a,(s',a')\}) - v(\mathcal{X}\cup\{(s',a')\})\}\\
&= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|}\times\Big\{v(\mathcal{X}\cup\{(s',1)^{a'}\}) - v(\mathcal{X})\Big\}\\
&+ \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s,a),(s',a')\}} w_{|\mathcal{X}|+1}\times\Big\{v(\mathcal{X}\cup\{(s',1)^{a'},(s,a)\}) - v(\mathcal{X}\cup\{(s,a)\})\Big\}\\
&= \sum_{\mathcal{X}\subseteq\{\mathbb{S}\times\mathbb{A}\}\setminus\{(s',a')\}} w_{|\mathcal{X}|}\times\Big\{v(\mathcal{X}\cup\{(s',1)^{a'}\}) - v(\mathcal{X})\Big\}\\
&= \pi_{s'}^{a',\text{Shap}}(\mathcal{M}_{s'}^{a'}).
\end{aligned}$$

Hence, we require $\pi_s^a = \pi_{s'}^{a'}$ for $\{\psi_s^a, \psi_{s'}^{a'}\}$ to satisfy counterfactual symmetry, which contradicts $\pi_s^a \neq \pi_{s'}^{a'}$.

Linearity: Since $\pi_s^a = \psi_s^a + \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a)$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$ and both ψ_s^a and $\pi_s^{a,\text{Shap}}(\mathcal{M}_s^a)$ satisfy linearity, it follows that π_s^a does so too.

Null player: Consider $(s,a) \in \mathbb{S} \times \mathbb{A}$. We need to show that

$$\underbrace{v(\mathcal{X} \cup \{(s,a)\}) = v(\mathcal{X}) \ \forall \mathcal{X} \subseteq \{\mathbb{S} \times \mathbb{A}\} \setminus \{(s,a)\}}_{(*)} \implies \pi_s^a = 0.$$

We can use the fact that ψ_s^a satisfies counterfactual null player, i.e.,

$$\underbrace{v(\mathcal{X} \cup \{(s,a)\}) = v(\mathcal{X} \cup \{(s,1)^a\}) \ \forall \mathcal{X} \subseteq \{\mathbb{S} \times \mathbb{A}\} \setminus \{(s,a)\}}_{(\square)} \implies \pi_s^a = \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a).$$

For the purposes of a contradiction, suppose that $(*)$ holds but $\pi_s^a \neq 0$. It suffices to show that this results in ψ_s^a violating counterfactual null player. Suppose $(\square)$ holds. Subtract $v(\mathcal{X})$ from both sides of $(\square)$ to get

$$v(\mathcal{X} \cup \{(s,a)\}) - v(\mathcal{X}) = v(\mathcal{X} \cup \{(s,1)^a\}) - v(\mathcal{X}) \ \forall \mathcal{X} \subseteq \{\mathbb{S} \times \mathbb{A}\} \setminus \{(s,a)\},$$

which combined with $(*)$ yields

$$0 = v(\mathcal{X} \cup \{(s,1)^a\}) - v(\mathcal{X}) \ \forall \mathcal{X} \subseteq \{\mathbb{S} \times \mathbb{A}\} \setminus \{(s,a)\}.$$

This implies $\pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) = 0$ and hence, we require $\pi_s^a = 0$ for ψ_s^a to satisfy counterfactual null player, which contradicts $\pi_s^a \neq 0$.

This completes the proof. $\square$

Appendix B: Proof of Proposition 1

PROPOSITION 1. *Consider $(s,a) \in \mathbb{S} \times \mathbb{A}$ and the Markov chain $\mathcal{M}$. The SV of (s,a) as defined in* (3) *equals*

$$\pi_s^{a,Shap} = \mathbb{E}_{\mathcal{P} \sim \mathcal{M}}\left[w_s^a(\mathcal{P})\right],$$

where

$$w_s^a(\mathcal{P}) := \begin{cases} \frac{1}{u(\mathcal{P})} & \text{if } \mathcal{P} \text{ converts and } (s,a) \in \mathcal{P} \\ 0 & \text{otherwise.} \end{cases}$$

The function $u(\cdot)$ returns the number of unique players and $\mathcal{P}$ is a path over players, i.e., state-action pairs.

Proof. For convenience, we use linearized notation such that $r := (s,a)$ and $\mathbb{P} := \mathbb{S} \times \mathbb{A}$ with the understanding that $\pi_r^{\text{Shap}} := \pi_s^{a,\text{Shap}}$, and $w_r(\cdot) := w_s^a(\cdot)$. The proof is split into three parts.

Step 1: Express $v(\mathcal{X})$ as an expectation over paths:

$$\begin{aligned} v(\mathcal{X}) &= \lambda^\top h^{\beta^{\mathcal{X}}} \\ &= \lambda^\top \left(I + P^{\beta^{\mathcal{X}}} + (P^{\beta^{\mathcal{X}}})^2 + \dots\right) p_c^{\beta^{\mathcal{X}}} \\ &= \mathbb{E}_{\mathcal{P} \sim \mathcal{M}(\mathcal{X})}\left[\mathbb{I}_{\{\mathcal{P} \text{ converts}\}}\right] \\ &= \mathbb{E}_{\mathcal{P} \sim \mathcal{M}}\left[\mathbb{I}_{\{\mathcal{P} \text{ converts}\}} \mathbb{I}_{\{\mathcal{X}(\mathcal{P}) \subseteq \mathcal{X}\}}\right], \end{aligned}$$

where $\mathcal{M}(\mathcal{X})$ denotes the Markov chain in which only the players in coalition $\mathcal{X}$ are "active" and $\mathcal{X}(\mathcal{P})$ denotes the set of unique players in path $\mathcal{P}$.

Step 2: Use the representation from Step 1 in the definition of SV π_r^{Shap} as stated in (3):

$$\begin{aligned}\pi_r^{\text{Shap}} &= \sum_{\mathcal{X}\subseteq\mathbb{P}\setminus\{r\}} w_{|\mathcal{X}|} \times \{v(\mathcal{X}\cup\{r\}) - v(\mathcal{X})\} \\ &= \sum_{\mathcal{X}\subseteq\mathbb{P}\setminus\{r\}} w_{|\mathcal{X}|} \left\{\mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[\mathbb{I}_{\{\mathcal{P}\text{ converts}\}}\left(\mathbb{I}_{\{\mathcal{X}(\mathcal{P})\subseteq\mathcal{X}\cup\{r\}\}} - \mathbb{I}_{\{\mathcal{X}(\mathcal{P})\subseteq\mathcal{X}\}}\right)\right]\right\} \\ &= \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[\mathbb{I}_{\{\mathcal{P}\text{ converts}\}}(\star)\right],\end{aligned}$$

where

$$(\star) := \sum_{\mathcal{X}\subseteq\mathbb{P}\setminus\{r\}} w_{|\mathcal{X}|}\left(\mathbb{I}_{\{\mathcal{X}(\mathcal{P})\subseteq\mathcal{X}\cup\{r\}\}} - \mathbb{I}_{\{\mathcal{X}(\mathcal{P})\subseteq\mathcal{X}\}}\right).$$

Step 3: Observe that for a given path $\mathcal{P}$, the expression in $(\star)$ can be seen as the SV of player r corresponding to a different game with characteristic function defined as

$$v_{\mathcal{X}(\mathcal{P})}(\mathcal{X}) := \mathbb{I}_{\{\mathcal{X}(\mathcal{P})\subseteq\mathcal{X}\}} \text{ for all } \mathcal{X}\subseteq\mathbb{P}.$$

A game with such a characteristic function is referred to as a *carrier game* and it is well-known (see Lemma 2 of Shapley (1953) for instance) that the SV of a player in a carrier game equals $1/u(\mathcal{P})$ if the player is in $\mathcal{P}$ and 0 otherwise. Therefore,

$$(\star) = \begin{cases}\frac{1}{u(\mathcal{P})} & \text{if } r\in\mathcal{P} \\ 0 & \text{otherwise,}\end{cases}$$

which implies

$$\pi_r^{\text{Shap}} = \mathbb{E}_{\mathcal{P}\sim\mathcal{M}}\left[w_r(\mathcal{P})\right].$$

This completes the proof. □

Appendix C: Proof of Theorem 3

THEOREM 3. *For any action, computing state-specific CASV and then aggregating over states is not the same as computing CASV over the aggregated model. In particular, for all $a\in\mathbb{A}$,*

$$\sum_{s\in\mathbb{S}} \psi_s^{a,Shap} - \bar{\psi}^{a,Shap} = \sum_{\mathcal{X}\in\mathbb{T}_1}\{c^a(\mathcal{X}) - \bar{c}^a(\mathcal{X})\}\{v_{\mathcal{M}}(\mathcal{X}) - v_{\bar{\mathcal{M}}^a}(\mathcal{X})\} + \sum_{\mathcal{X}\in\mathbb{T}_2}\sum_{s\in\mathbb{S}} c_s^a(\mathcal{X})\{v_{\mathcal{M}}(\mathcal{X}) - v_{\mathcal{M}_s^a}(\mathcal{X})\},$$

where $c^a(\mathcal{X}) := \sum_{s\in\mathbb{S}} c_s^a(\mathcal{X})$,

$$c_s^a(\mathcal{X}) := \begin{cases} w_{|\mathcal{X}|-1} & \textit{if } (s,a)\in\mathcal{X} \\ -w_{|\mathcal{X}|} & \textit{otherwise}\end{cases} \quad \textit{and} \quad \bar{c}^a(\mathcal{X}) := \begin{cases}\bar{w}_{|\bar{\mathcal{X}}|-1} & \textit{if } a\in\bar{\mathcal{X}} \\ -\bar{w}_{|\bar{\mathcal{X}}|} & \textit{otherwise.}\end{cases}$$

The notation $\bar{\mathcal{X}}$ denotes the actions in $\mathcal{X}$, i.e., $\bar{\mathcal{X}} := \{a : (s,a)\in\mathcal{X}\}$.

Proof. We split the proof into four parts.

Step 1: We use the fact that CASV equals the difference between two SVs

$$\begin{aligned}\psi_s^{a,\text{Shap}}(\mathcal{M}) &= \pi_s^{a,\text{Shap}}(\mathcal{M}) - \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) \quad \forall (s,a)\in\mathbb{S}\times\mathbb{A} \\ \bar{\psi}^{a,\text{Shap}}(\mathcal{M}) &= \bar{\pi}^{a,\text{Shap}}(\mathcal{M}) - \bar{\pi}^{a,\text{Shap}}(\bar{\mathcal{M}}^a) \quad \forall a\in\mathbb{A}\end{aligned}$$

to express the difference $\sum_{s\in\mathbb{S}} \psi_s^{a,\text{Shap}} - \bar{\psi}^{a,\text{Shap}}$ as

$$\underbrace{\left[\sum_{s\in\mathbb{S}} \pi_s^{a,\text{Shap}}(\mathcal{M}) - \bar{\pi}^{a,\text{Shap}}(\mathcal{M})\right]}_{=:\,(\diamond)} - \underbrace{\left[\sum_{s\in\mathbb{S}} \pi_s^{a,\text{Shap}}(\mathcal{M}_s^a) - \bar{\pi}^{a,\text{Shap}}(\bar{\mathcal{M}}^a)\right]}_{=:\,(\bullet)}.$$

Step 2: To analyze $(\diamond)$, we observe the state-specific and aggregated SVs (see (3) and (10)) equal

$$\begin{aligned} \pi_s^{a,\text{Shap}}(\mathcal{M}) &= \sum_{\mathcal{X}\in\mathbb{T}_1\cup\mathbb{T}_2} c_s^a(\mathcal{X}) v_{\mathcal{M}}(\mathcal{X}) \\ \bar{\pi}^{a,\text{Shap}}(\mathcal{M}) &= \sum_{\mathcal{X}\in\mathbb{T}_1} \bar{c}^a(\mathcal{X}) v_{\mathcal{M}}(\mathcal{X}), \end{aligned}$$

where $\mathbb{T}_1$ and $\mathbb{T}_2$ are the sets of type 1 and 2 coalitions, respectively. This implies

$$(\diamond) = \sum_{\mathcal{X}\in\mathbb{T}_1} \{c^a(\mathcal{X}) - \bar{c}^a(\mathcal{X})\}\, v_{\mathcal{M}}(\mathcal{X}) + \sum_{\mathcal{X}\in\mathbb{T}_2} \sum_{s\in\mathbb{S}} c_s^a(\mathcal{X}) v_{\mathcal{M}}(\mathcal{X}).$$

Step 3: We analyze $(\bullet)$ similarly to obtain

$$(\bullet) = \sum_{\mathcal{X}\in\mathbb{T}_1} \{c^a(\mathcal{X}) - \bar{c}^a(\mathcal{X})\}\, v_{\bar{\mathcal{M}}^a}(\mathcal{X}) + \sum_{\mathcal{X}\in\mathbb{T}_2} \sum_{s\in\mathbb{S}} c_s^a(\mathcal{X}) v_{\mathcal{M}_s^a}(\mathcal{X}),$$

where we use the fact that if $\mathcal{X}\in\mathbb{T}_1$, then $v_{\mathcal{M}_s^a}(\mathcal{X}) = v_{\bar{\mathcal{M}}^a}(\mathcal{X})$ for all $(s,a)\in\mathbb{S}\times\mathbb{A}$.

Step 4: Putting steps 1, 2, and 3 together concludes the proof.

$\square$