Time-weighted attribution of revenue to multiple e-commerce marketing channels in the customer journey

David A. Wooff[a*] and Jillian M. Anderson[b]

[a] *Durham University, Department of Mathematical Sciences, Stockton Road, Durham DH1 3LE, UK*
[b] *Summit Media Ltd, Albion Mills, Albion Lane, Willerby, Hull, HU10 6DN, UK*

July 7th 2013

Abstract

In this paper we address statistical issues in attributing revenue to marketing channels. We describe the relevant data structures and introduce an example. We suggest an asymmetric bathtub shape as appropriate for time-weighted revenue attribution to the customer journey, provide an algorithm, and illustrate the method. We suggest a modification to this method when there is independent information available on the relative values of the channels. We compare the revenue attributions suggested by the methods in this paper with several common attribution methods.
Keywords: Multi-touch; Asymmetric bathtub weighting; Beta distribution; Online marketing; Path to conversion; Clickstream; Digital marketing; E-commerce.

1 Introduction

This paper concerns statistical analysis of the routes to online purchase – known as conversion – by customers at a retail internet site. Prior to conversion, consumers typically visit several websites, including multiple visits to the final retail site, for purposes including searching, browsing and knowledge building [Moe, 2003]. A typical example might begin with a customer searching for a product, narrowing down on product details, using shopping comparison sites to compare prices, checking for availability of vouchers, and so forth. This is the customer journey, also known as the *clickstream*. Retailers use a variety of online marketing channels to raise brand awareness and drive conversions; therefore, it is possible for a consumer to interact with multiple marketing channels prior to conversion. The customer journey is recorded via cookies stored on the consumer's computer. Usually, some fraction of the sale revenue is allocated to steps in the journey. This is known as revenue attribution. Simplistically, these are monetary rewards for sites which funnel customer traffic towards the final retailer. These sites are classified as marketing channels of various kinds, such as display campaigns, direct email advertisements, and social media such as Facebook. More detailed descriptions of the process may be found in Abhishek et al. [2012] and Xu et al. [2012].

In 2012, total spend on digital advertising in the UK alone amounted to £5416 million, with annual growth of around 13% [Internet Advertising Bureau UK, 2013]. In the USA, corresponding spend is presently around $40000 million [Dalessandro et al., 2012]. Around 58% of spend is on pay-per-click (PPC) advertisements via search engines such as Google, Bing, and Yahoo. The remaining spend is on other digital marketing channels. This sector of the economy is already of major importance, and growing, but many aspects are poorly understood, including our area of interest, the customer journey. Industry evidence is that around 65% of conversion journeys contain more than one visit to the final retail site, and about 81% contain interactions with more than one marketing channel. Understanding the true value of each kind of marketing channel leads to better budget planning, to identification of crucial steps in the journey, and to improved exploitation of emerging channels. One issue is thus to apply a measure of value to the various marketing channels. In the UK, this is the weighted attribution problem. In the USA it is better known as the multi-touch attribution problem.

*Corresponding author. Email: d.a.wooff@durham.ac.uk

Existing methods for attributing conversions to marketing channels range from the simplistic to detailed algorithms. The most basic methods attribute the conversion to a single step in the journey, typically the first step in the journey ("first click wins") or the last step prior to conversion ("last click wins"). By only acknowledging a single channel within a conversion journey we underestimate the importance of channels which might only appear as intermediate in the journey, but which may in fact be crucial to the conversion. Multi-channel attribution assigns a proportion of the conversion revenue to each step in the journey. A recent survey suggests that 30% of retailers use single-source attribution, 34% use a multiple-source method, and 11% use an algorithm-based approach [Osur et al., 2012]. Algorithm-based methods aim to use a more scientific approach to assigning value to each visit by considering both converting and non-converting journeys in order to determine the probability of each channel leading to a conversion; for example, see Shao and Li [2011], Abhishek et al. [2012], Xu et al. [2012]. Many current multi-channel models are subjective with weights assigned on a marketer's experience rather than data analysis. Algorithm-based methods may require large amounts of data and the ability to track, store and analyse data at the consumer level, entailing a higher cost, but normally with a payoff in understanding.

There is no industry standard for attributing revenue and no single measure exists for comparing the many different methods available. Dalessandro et al. [2012] recommends these properties of a good attribution model: (1) *fairness* – attribution should be based on the channel's ability to influence conversions; (2) *data-driven* – attribution should be based on statistical principles, but should also utilise a retailer's knowledge of the marketplace; (3) *interpretability* – the attribution model should be transparent and sufficiently simple to be understood and implemented by all. We propose methods which satisfy these criteria, and which also takes into account temporal features in the journey. We propose a method for dealing with attribution when we have no information about the relevance of different channels to conversion behaviour, and a modification of the method when we do have such information. A companion paper [Wooff and Anderson, 2013] considers the problem of obtaining such information through mining of data arising from sequential analysis of the customer journey.

In Section 2 we describe the relevant data structures and introduce an example. In Section 3 we suggest an asymmetric bathtub shape as appropriate for time-weighted revenue attribution to the customer journey, provide an algorithm, and illustrate the method. In Section 4 we suggest a modification to this method when there is independent information available on the relative values of the channels. In Section 5, we compare the revenue attributions suggested by the methods in this paper with several common attribution methods.

2 Preliminary processing of data

2.1 Data Collection

We suppose that web analytics tools have been used to collect information about a customer's journey subsequent to a conversion. Pixel tracking is used to record each visit a user makes to a website. The marketing channel and time of each visit is recorded, along with conversion details such as sale type, sale ID, revenue). Visits may be categorised at the marketing channel level (e.g. direct, PPC, organic search) or at a more granular level (e.g. search term, keyword, category). We make no inferences regarding journeys which may be artificially shortened via users either deleting or refusing permission to store cookies. Impressions, or views, of an advertisement may also be included in the journey sequence and assigned a value in a similar manner to visits.

A visit duration window is applied to multiple visits from the same channel, that is, subsequent visits are not recorded if they occur within a given timeframe thereby reducing the influence of click fraud and user behaviour (e.g. page refresh, navigation confusion). Industry standards set the visit duration window at 10 minutes for marketing channels. Furthermore, a maximum time between a visit and conversion is imposed, and will be referred to as the cookie window. The choice of cookie window is subjective, but guided by industry expertise. For retail, 31 days is commonly employed.

We exclude, for now, journeys reaching a terminus such as site registration or booking an appointment. We assume where necessary that abandonment of a journey without conversion is final within a given time period. This is an approximation as some customers do continue their journeys after long breaks. Note that some journeys which do not end in online conversion may end in offline conversion, with customers visiting a store to purchase a product identified online. This is presently excluded from our analysis.

Mathematically we will view the different possible visits as nodes in a sequence. The definition of what constitutes a visit source depends on the requirements of the retailing company. Sometimes this will be at a very fine level of detail, such as named weblinks. At other times the sources may have been classified by the retailer into a smaller number of channel categories such as 'direct', 'email', etc, as it deems appropriate. This is the case in the example we discuss in Wooff and Anderson [2013].

2.2 Data Processing

Suppose we observe a sequence of customer visits to a retail website made at times $T_1, T_2, \ldots, T_k$. We make an assumption that visits that occur further back in time than a specified amount T_{max} are not relevant to the current conversion. Analysis of the journey database allows a retailer-specific T_{max} to be set. The journey lengths, $T_k - T_1$, of all journeys in the database are analysed, with the 90^{th} percentile chosen as T_{max}. Journeys where $T_k - T_1 > T_{max}$ are truncated at the visit T^*, where $T^* \leq T_{max}$.

We also make the assumption that time gaps larger than a specified amount T_Δ imply separate journeys. Thus, if any adjacent times satisfy $T_j - T_i > T_\Delta$, we end one journey at T_i and start another at T_j. All transition times $(T_j - T_i)$ within the journey database are analysed, with T_Δ set at the 90^{th} percentile. For the purposes of this article we consider only one value for T_Δ, however, it is understood that T_Δ may vary depending on the sequence of marketing channels. Journey fragments prior to T_i are not considered in this article.

A maximum number of visits V_{max} might also be imposed, in that journeys with number of visits exceeding V_{max} are assumed to be due to tracking discrepancies and are removed from the analysis. The choice of V_{max} may be based on analysis of the journey database.

Imposing a T_{max} and T_Δ results in left-censoring of the data. The main implication is that data concerning the first click is lost. In analysing such data, the implicit assumption is that T_1 is either genuinely the start of a new journey, or a click made in the same journey but with the preceding click so distant in time that it is deemed irrelevant. For analysis of journeys which end in conversion, the use of a time gap threshold may result in early parts of the journey being discarded. For data where conversion behaviour is an outcome, journeys might be separated into non-converting and converting fragments, and the correlation between the two may be lost.

2.3 Example

INSERT Table 1 about here

Consider the fragment of data shown in Table 1. Data are taken from a sample of customer conversions made on a leading multichannel retail website. Each journey has a starting time T_1, and a number of visits in sequence with time recorded. Also shown is the amount of conversion, the revenue attributed to each journey. These data are reported to two decimal places, but shown rounded in the table. The focus of analysis for this data set is the route to conversion. A maximum journey length T_{max} of 30 days was used, and visits made before T_{max} are removed. A time gap threshold of $T_\Delta = 14$ days was also used, and fragments of any journey with at least such a time gap were discarded. Each customer journey is analysed separately and only time since start of journey is assumed relevant. As such, we fix $T_1 = 0$ for each journey. A maximum number of visits in the journey was also set at $V_{max} = 11$; journeys with more than 11 visits were removed. More than 95% of journeys in the database contained 11 visits or fewer. It is, of course, possible to explore the implications of different choices of T_{max} and so forth, but this is outside the scope of this paper.

The data subset contains visits from a number of channels which may be split into varying degrees of granularity. Natural search channels may be split by search partner (e.g. Google, Bing) or category (e.g. brand, non-brand). Affiliate channels may be categorised according to type (e.g. cashback, voucher codes); this is particularly important for understanding the value of marketing campaigns within the context of attribution and budget forecasting. Visits via individual comparison sites are also included. Finally, for account optimisation, PPC visits may be split at the keyword level, where keyword can be broadly interpreted as meaning a search word or phrase. Identifying keywords which have a strong influence on likely final conversion is a crucial aspect of digital marketing performance. Visits which are not classified into a specific channel are classed as "unlisted referrers" and could be excluded from the attribution model, or assigned a weight of zero; for discussion see Section 5.

This sample of data exhibits features typical of the problem. Journeys vary in length of time. Significant time can be spent on one visit, or the journey can be relatively time-homogeneous. There are two two-step journeys. Instances where successive visits are within the same minute as the previous click (for example, see journey 3 in Table 1) represent visits either by a different channel or search query and are not to be interpreted as page refresh errors. Single visit journeys are assigned revenue and removed from the attribution database after data cleaning.

3 Naive time-weighted Revenue allocation

Suppose we observe the customer journey $X_{(1)} \to X_{(2)} \to \ldots \to X_{(k)}$, $2 \leq k \leq S$, with conversion at node $X_{(k)}$ resulting in revenue R, and where S is some truncating choice. Suppose we visit node $X_{(i)}$ at time T_i, so that the journey begins at T_1 and ends at T_k. Suppose also that we have no information concerning the relative importance of nodes in the journey. The problem is to attribute the revenue to the nodes in the journey, or equivalently to value each node. There are many views as to how we might do this. One is to attribute all revenue to the last node in the journey, known as *last click wins*. This corresponds to the view that the journey itself is irrelevant and that the customer would have arrived at node $X_{(k)}$ irrespective of starting point. Another view is to attribute all revenue to the first node in the journey, known as *first click wins*. This corresponds to the view that once the journey has started at X_1 the journey will end inexorably with a conversion at node $X_{(k)}$. A third view is that all nodes in the journey count equally towards the final conversion, in which case revenue might be attributed equally to each node. There are many other views which suggest that clicks closer to conversion should have a higher weighting. These lead to weights based on monotonically rising functions, for example positive linear and exponential.

In discussion with digital marketing experts, none of these views is felt to be reasonable. Instead, they suggest the following plausible structure. We value recent clicks highly, especially the most recent click. We value the initiating click highly, but less highly than the last click. We value intervening clicks not highly if they are quite distant in time, and less than the initiating click. We regard clicks close in time to the last click as being highly relevant. This suggests that the shape of value which we wish to allocate to clicks in the journey might have an asymmetric bathtub shape, with the rim of the bath lower at the left-hand side. Such bathtub shapes are common in survival analysis, through representing hazard functions. We now consider how to construct such a shape for this application. A simple asymmetric bathtub shape, constructed using a beta distribution, is shown in Figure 1 for Journey 10 of Table 1.

INSERT Fig 1 about here

3.1 Theory

The beta distribution is of the form $f(x) = kx^{a-1}(1-x)^{b-1}$, $0 < x < 1$, where k is a normalising parameter which is of no interest in this context. The parameter choices $0 < a < b < 1$ lead to asymmetric U-shaped distributions with a higher rim at the right-hand side. Other parameter choices can lead to J-shaped and unimodal distributions. Although the distribution is defined on the interval $(0, 1)$, it is trivial to transform journey time (T_1, T_k) to $(0, 1)$ and back again. In fact we will transform not to $(0, 1)$ but to $(\epsilon, 1-\epsilon)$ to avoid infinities at the asymptotes. Experience shows that a good choice is $\epsilon = 0.01$. Smaller values of ϵ imply steeper behaviour at the asymptotes, with the consequence that the last click will be valued relatively more than the penultimate click. The minimum of the distribution occurs at

$$\gamma = \frac{a-1}{a+b-2}, \tag{1}$$

so that $\theta = f(\gamma)$ will be the smallest possible weight given to any click.

We need to make choices about the relative values of clicks. Let θ_L be the relative value of the last click in the journey as compared to the first click in the journey. Let θ_F be the relative value of the first click in the journey as compared to θ, potentially the value assigned to the least valuable click in the journey.

The choices of θ_F and θ_L will depend on context. In discussion with our marketing collaborator, it was felt appropriate to deem the last click as worth about four times as much as the first click, and the first click as worth about twice the minimum value we would wish to assign. That is, $\theta_L = 4$ and $\theta_F = 2$, so that the last click is worth $\theta_L\theta_F = 8\theta$, eight times as much as the least valuable click. Such choices are unavoidable. For example, the judgement that all clicks should be evenly weighted corresponds to $\theta_F = \theta_L = 1$. Similarly, where there is an attribution which rises linearly in value from first click to last click, the underlying choice is $\theta_F = 1$ and θ_L is proportional to the slope of the chosen line.

Given these assumptions, we now generate parameter values for our beta distribution. We have

$$\theta_L = \frac{f(1-\epsilon)}{f(\epsilon)} \Rightarrow a = b + v, \tag{2}$$

where

$$v = \frac{\log \theta_L}{\log(1/\epsilon - 1)}.$$

Note that $v > 0$ in order to obtain a higher rim at the right-hand side. We have also

$$\theta_F = \frac{f(\epsilon)}{f(\gamma)} = (\frac{\epsilon}{\gamma})^{b+v-1}(\frac{1-\epsilon}{1-\gamma})^{b-1}, \tag{3}$$

where we can re-express γ via (1,2) as

$$\gamma = \frac{b+v-1}{2b+v-2}.$$

This gives a highly non-linear equation in b, which may be solved numerically. The constraints of the numerical solution are that $0 < b < 1 - v$. This follows as we require $a < 1$ in order to obtain a U-shape. An algorithm for attributing revenue to a channel is thus as follows.

(i) Choose θ_F and θ_L. Fix $\epsilon = 0.01$. Compute v.

(ii) Solve (3) for b and determine a via (2).

(iii) For Journey J with revenue R_J to attribute, transform the click times $T_1, T_2, \ldots, T_k$ linearly to $(\epsilon, 1-\epsilon)$. This gives transformed time values $T_1^* = \epsilon, T_2^*, \ldots, T_k^* = 1 - \epsilon$. Evaluate $w_i = f(T_i^*)$ for each transformed time. The proportion of revenue attributed to the channel clicked at time T_i is $w_i^* R_i$, where

$$w_i^* = \frac{w_i}{\sum_{i=1}^k w_i}.$$

There may be journeys for which all recorded click times are the same, perhaps because of rounding. In this case the rescaling to $(0,1)$ fails and it is simplest to give equal weight to all clicks in such journeys.

3.2 Example

INSERT Table 2 about here

INSERT Table 3 about here

For our data set we choose $\theta_F = 2$ and $\theta_L = 4$. Solving with these choices we obtain $a = 0.739$ and $b = 0.437$. The curves obtained are shown in Figure 2 for journeys 2,3,10,25. For journey 10, the weights and revenue attribution are shown in Table 2. The revenue attributions for all journeys are shown in Table 3. Note that attributions must now be accumulated over channels (or at a more granular level depending on purpose); for example the clicks at T_1 and T_2 for a journey could correspond to the same channel. One feature evident in this data set is multiple clicks close in time, and so which attract similar revenues. In principle it is not difficult to provide more sophisticated methods which could take into account subjective judgements concerning clicks close in time. For example, one might wish to discount all but the most recent of a group of clicks occurring in a narrow time range.

INSERT Fig 2 about here

4 Informed revenue allocation

In this section we discuss weighted attribution when we also have information about the relative importance of different nodes. Judgements about relative importance may be made directly. For example, in the context of online marketing a company might wish to value PPC channels more highly than natural search or email marketing. A number of researchers, see for example Shao and Li [2011], Abhishek et al. [2012], Xu et al. [2012], have provided measures of channel value relating directly to probability of conversion. This requires data on converting and non-converting journeys. Where we have data only on converting journeys, we provide a method to infer channel relevance based on sequential data analysis of journey fragments in a companion paper [Wooff and Anderson, 2013].

Whether channel value is inferred or specified, we suppose that the relative values of the n channels are $u_1, u_2, \ldots, u_n$, where $\sum_{i=1}^n u_i = 1$. There are different possible ways of merging weights due to time and weights due to channel value. The simplest is to compound the two sets of weights and then re-normalize. Thus, suppose that $a_{(1)}, a_{(2)}, \ldots, a_{(k)}$ are the weights suggested by time of click for a k-step journey. These weights are derived using the bathtub method of Section 3, the linear method, or any other desired method. Let $u_{(i)}$ be the value of the i^{th} node clicked. The compounded weight for the node clicked on the i^{th} step of the journey is then

$$a^*_{(i)} = \frac{u_{(i)} a_{(i)}}{\sum_{j=1}^k u_{(j)} a_{(j)}}. \tag{4}$$

Thus, an attribution to the node clicked on step i of the journey which is both time-weighted and value-weighted is given by multiplying weight $a^*_{(i)}$ by the revenue for the journey.

5 Comparison of weighted attribution mechanisms

INSERT Table 4 about here

We consider data from a major UK online retailer which are discussed in depth in Wooff and Anderson [2013]. There are 58667 journeys ending in an online purchase. Of these, 27420 are single-click and 31247 have at least two clicks. 17841 journeys have at least three clicks. We limit to the most recent $S = 19$ steps of any journey. Each click is classified as belonging to one of nine channels as shown in Table 4. This shows that a high proportion of single-click journeys for this retailer at this time were branded natural search, NatB.

Figure 3 shows the total revenue attributions to eight channels for 58667 journeys for seven attribution methods: (1) the bathtub method described in Section 3 with $\theta_L = 4$ and $\theta_F = 2$; (2) first click wins; (3) last click wins; (4) equal weighting of all clicks – this corresponds to $\theta_L = 1$ and $\theta_F = 1$; (5) linear with last click valued at four times first click – this corresponds to $\theta_L = 4$ and $\theta_F = 1$; (6) exponential with last click valued at four times first click; and (7) the bathtub method additionally weighted according to channel value using metric $\tilde{r}_j$ of [Wooff and Anderson, 2013, Equation (4)], and with weights shown in [Wooff and Anderson, 2013, Table 7]. These weights are then compounding with time using (4). For this online retailer, all attribution methods yield similar results. Of note is that first-click-wins (2) tends to undervalue the **Aff** channel, whereas last-click-wins (3) tends to overvalue it; this is expected as the the nature of affiliate sites is to target consumers at the end of their journey that have already made the decision to buy and to provide a reward (e.g. cashback) for the purchase. Natural search (**Nat** and **NatB**) and PPC (**PPC** and **PPCB**) clicks can be assumed to be part of all stages of the buying journey (browsing, researching and buying) and therefore are expected to be rewarded similarly independent of the attribution model. It should be noted that an exception to this is that the bathub/value method (7) tends more highly to reward the **NatB** channel as it was found to be the most important intermediary channel in a typical journey: see the final column of [Wooff and Anderson, 2013, Table 7], suggesting that **NatB** is perhaps more a navigational click rather than a conversion driver. Different data sets may reveal very different patterns.

INSERT Fig 3 about here

From a marketing perspective, it is more useful to understand the impact of attribution models at a more granular level, especially for PPC (both generic and brand). PPC accounts are optimised at a keyword level, therefore using single-source attribution models tend to reward certain types of keywords only (branded and highly specific keywords). An advantage of a multi-source method is to reward more generic keywords that appear in the browsing and researching phase of the the journey, thereby maintaining (or increasing) spend on these keywords. Reducing spend on generic keywords, which would be valued highly as intermediary nodes, may be detrimental to conversion performance as the link between the source and destination nodes would be removed.

6 Discussion

In this paper we offer a sensible revenue attribution mechanism based on appropriate time-weighting of clicks. We have also shown how the method may be modified when there is separate information available on the quality of visitable channels. There is unavoidably a subjective element in choosing an appropriate shape for time-weighted attribution. This is the same problem faced by Bayesian statisticians in choosing an appropriate prior. This is an uncomfortable fact for major retailers, who often naively expect that there is a single "right" answer. The choice of attribution shape and parameters such as θ_L, the ratio of last click to first click value, depend on the aims of the attribution. If a retailer wishes only to prioritize last-click-wins, then that is the "right" answer for them. Much of the time, what is sought is an appropriate balance between the interests of the retailer and the interests of such as affiliates who may appear in the customer journey.

Acknowledgements

Part of this research was funded by Knowledge Transfer Partnership KTP007499, funded by Summit Media Ltd. and by the UK Technology Strategy Board. We are grateful to Summit Media Ltd for providing data and to colleagues there for providing expertise.

References

V. Abhishek, P. Fader, and K. Hosanagar. The long road to online conversion: A model of multi-channel attribution. `http://dx.doi.org/10.2139/ssrn.2158421`, 2012.

Brian Dalessandro, Claudia Perlich, Ori Stitelman, and Foster Provost. Causally motivated attribution for online advertising. In *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*, ADKDD '12, pages 7:1–7:9, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1545-6. doi: 10.1145/2351356.2351363. URL `http://doi.acm.org/10.1145/2351356.2351363`.

Internet Advertising Bureau UK. 2012 Online Adspend Full Year Results. `http://www.iabuk.net/research/library/2012-full-year-digital-adspend-results`, 2013.

W. W. Moe. Buying, searching or browsing: Differentiating between online shoppers using in-store navigational clickstream. *Journal of Consumer Psychology*, 13(1&2):29–39, 2003.

Ari Osur, Emily Riley, Tina Moffett, Sarah Glass, and Elizabeth Komar. The Forrester Wave Interactive Attribution Vendors Q2 2012. Technical report, Forrester Research, Inc, 2012.

Xuhiu Shao and Lexin Li. Data-driven multi-touch attribution models. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 258–264, 2011.

D. A. Wooff and J. M. Anderson. Inferring marketing channel importance in the customer journey to online purchase. *Durham Research Online*, 2013.

L. Xu, J. A. Duan, and A. B. Whinston. Path to purchase: A mutually exciting point process model for online advertising and conversion. `http://dx.doi.org/10.2139/ssrn.2149920`, 2012.

Table 1: Journeys and conversion revenues for 25 customers, minimum two-step journeys with $S = 11$. Figures given are times of visit rounded to the nearest minute and starting time arbitrarily at $T_1 = 0$ for each customer.

i	T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9	T_{10}	T_{11}	Revenue
1	0	19	70	106	106							869
2	0	113										309
3	0	0	0	37	37	114						50
4	0	0	118									329
5	0	1	7	7	7	122	122					280
6	0	84	137									322
7	0	53	111	144	144	144						196
8	0	13	13	14	137	142	147	148				100
9	0	0	136	149	149	149	149	149				244
10	0	25	77	79	79	79	167	167	167	167	167	378
11	0	0	50	169								494
12	0	172										205
13	0	178	178	178								340
14	0	52	179	179								370
15	0	0	180	180								136
16	0	0	79	181								1289
17	0	33	39	191								160
18	0	198										213
19	0	14	27	99	115	120	125	204				249
20	0	6	139	145	150	153	153	167	206			163
21	0	77	216	216	218	218	218					330
22	0	90	117	121	151	151	241	243	243	247	247	95
23	0	6	126	251	251	251	251					150
24	0	263	263									270
25	0	20	22	23	23	153	247	272				239

Table 2: Weights and revenue attributions for Journey 10 of the journeys shown in Table 1.

i	Time of click, T_i	Weight, w_i^*	Revenue attributed
1	0.00	0.042	16.03
2	25.05	0.023	8.56
3	77.18	0.022	8.31
4	79.05	0.022	8.36
5	79.09	0.022	8.36
6	79.09	0.022	8.36
7	167.25	0.169	63.71
8	167.26	0.169	63.92
9	167.27	0.170	64.13
10	167.27	0.170	64.13
11	167.27	0.170	64.13
		1.000	378.00

Table 3: Revenue attributions for 25 customer journeys. Figures given are attributions of revenue to the channel clicked at that time, rounded to the nearest integer.

i	T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9	T_{10}	T_{11}	Revenue
1	86	45	52	342	344							869
2	62	247										309
3	6	6	6	3	3	25						50
4	55	55	219									329
5	24	21	15	15	15	95	95					280
6	58	33	231									322
7	14	7	10	55	55	55						196
8	8	5	5	5	10	13	23	32				100
9	11	11	12	42	42	42	43	43				244
10	16	9	8	8	8	8	64	64	64	64	64	378
11	76	76	38	304								494
12	41	164										205
13	26	104	104	105								340
14	39	20	156	156								370
15	14	14	54	55								136
16	198	197	102	792								1289
17	26	14	14	106								160
18	43	170										213
19	30	18	16	16	16	17	17	119				249
20	17	12	10	11	11	11	11	13	66			163
21	17	9	47	48	69	70	70					330
22	5	3	3	3	3	3	10	12	12	21	21	95
23	8	6	4	33	33	33	33					150
24	30	120	120									270
25	26	16	16	16	16	15	30	106				239

Table 4: Single-click-journey probabilities

Channel	Code	Freq	Prob
Affiliates	Aff	3841	0.1401
Banner	Ban	62	0.0023
Price Comparison	Comp	818	0.0298
Listed Referrer	List	96	0.0035
Natural Search (Other)	Nat	1954	0.0713
Natural Search (Brand)	NatB	14081	0.5135
Pay-per-click	PPC	2174	0.0793
Pay-per-click (Brand)	PPCB	2543	0.0927
Unlisted Referrer	Un	1851	0.0675
All		27420	1.0000

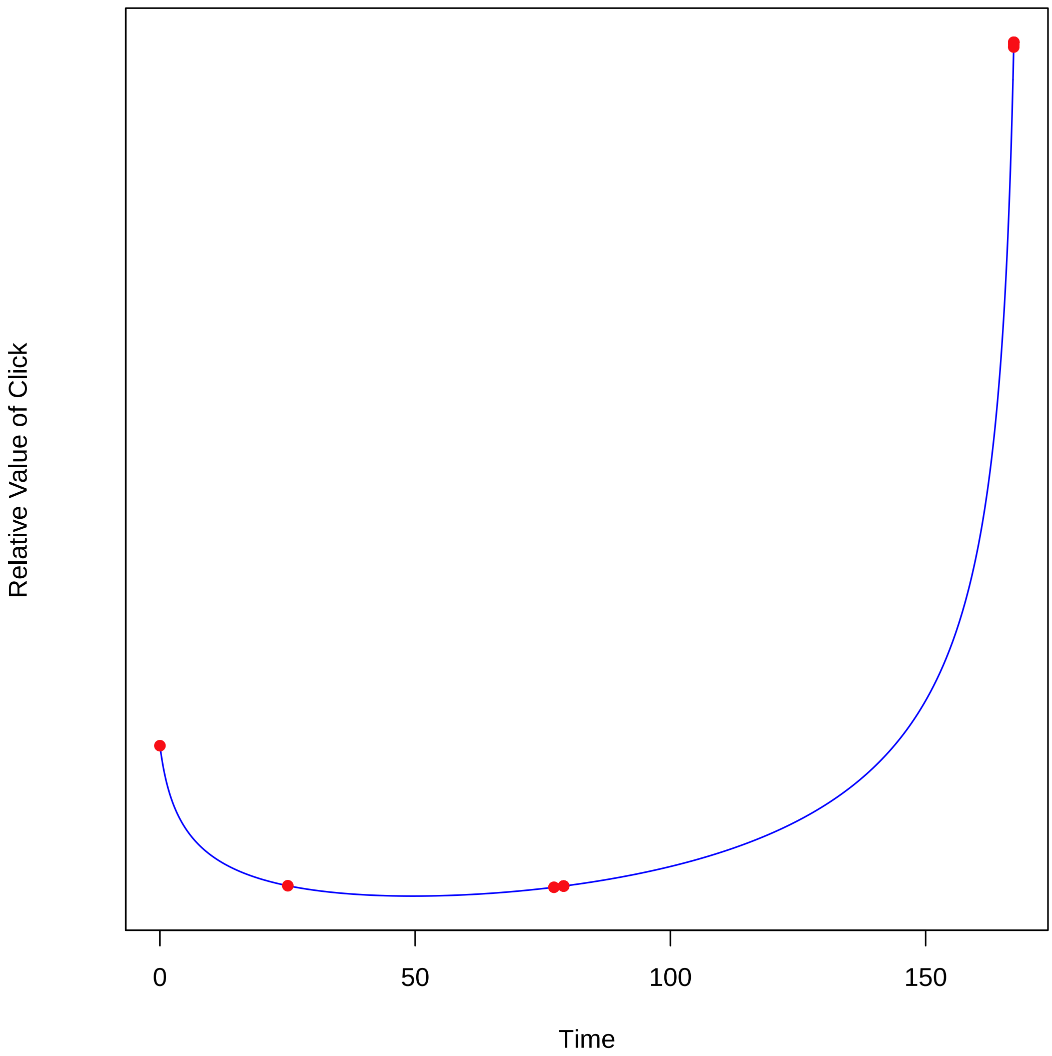

Figure 1: Simple bathtub model for click value, with clicks for Journey 10 marked.

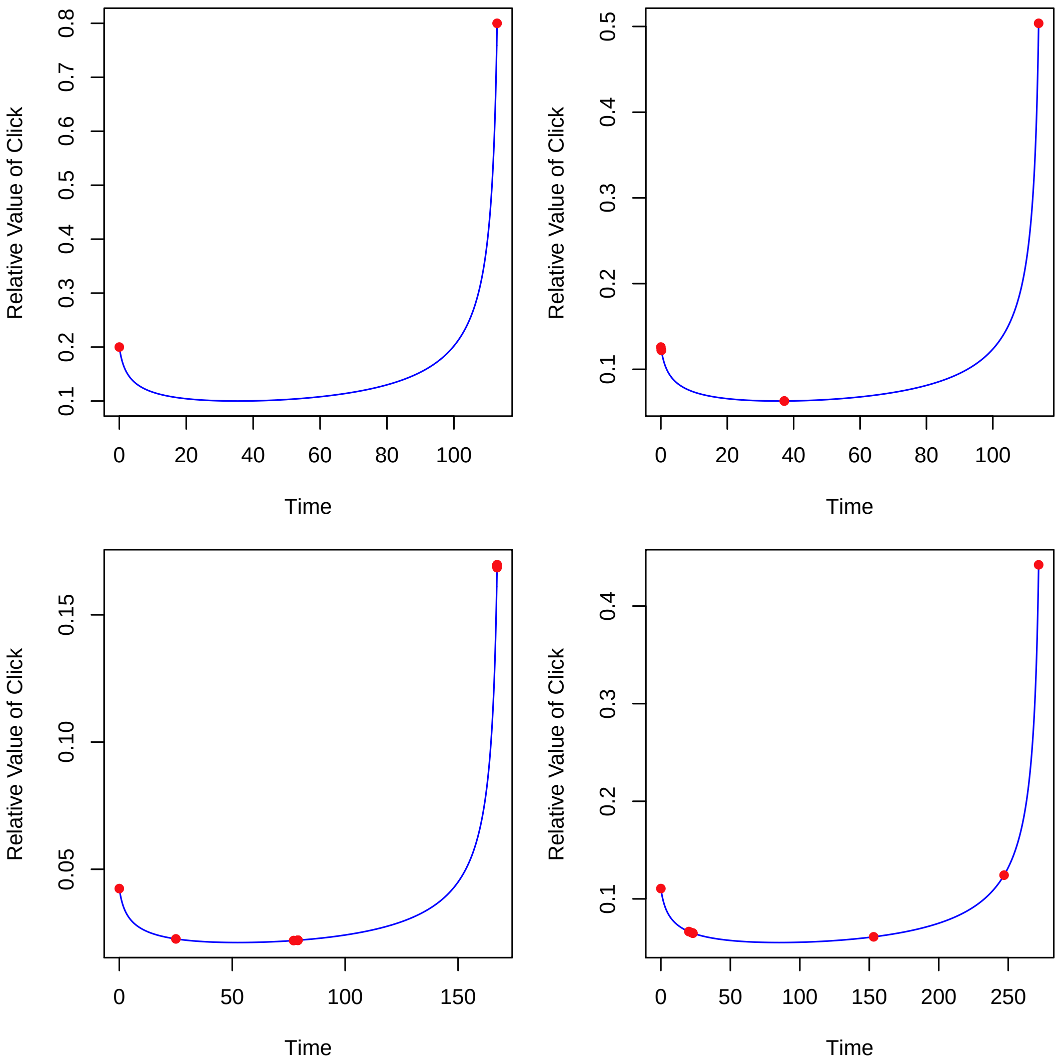

Figure 2: Value of clicks for journeys 2,3,10,25. Beta function parameters are $a = 0.739, b = 0.437$. Last click is worth $\theta_L = 4$ times as much as first click. First click is worth $\theta_F = 2$ times as much as the minimum possible.

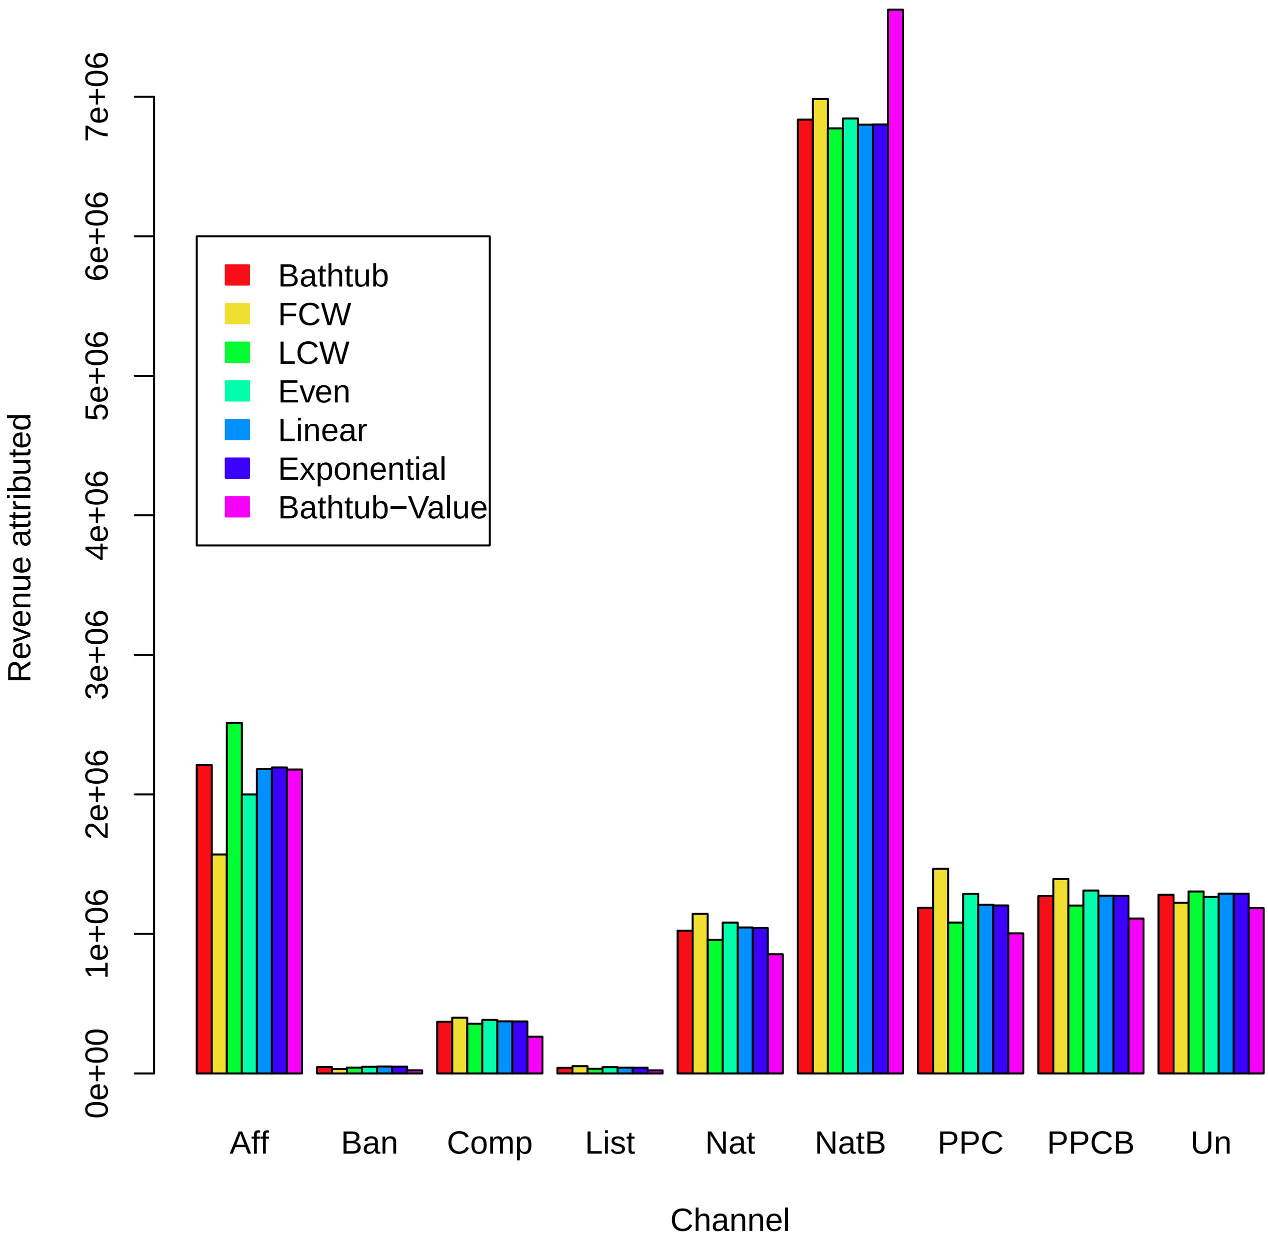

Figure 3: Comparison of total revenue attributions to eight channels for 58667 journeys for seven attribution methods.