
Multi-touch attribution: A case study in automotive media optimisation

Received (in revised form): 10th May, 2019



Jonathan Prantner

is Chief Analytics Officer and co-founder of RXA, where he leads efforts surrounding applied artificial and machine learning as well as integrating advanced analytics with data visualisation platforms. Jonathan's career has spanned educational research, automotive, consumer packaged goods, travel and healthcare. Jonathan is a celebrated thought-leader and recipient of multiple data science patents whose approach to applied mathematics has pushed analytics to the limits for over two decades.

RXA, 330 E. Liberty St. Ann Arbor, MI 48104, USA
Tel: +1 888 294 1512; E-mail: jp@rxa.io

Abstract Multi-touch attribution is a foundation of proper reporting and media optimisation. In the absence of e-commerce, an alternative to direct sales must be developed in order to properly value the credit of each visitor touched by media. This paper uses an automotive marketing case study to illustrate how a Markov-based approach to attribution can be applied to customer digital interactions to form a causal relationship with sales. Once the relationship between consumer digital behavioural data is combined with the brand's media executions, media optimisation including saturation analysis is enabled.

KEYWORDS: multi-touch attribution, media mix, website analytics, Markov, customer journey, automotive

SITUATION

In today's marketing landscape, one of the hottest topics is multi-touch attribution (MTA). MTA examines each individual visitor on a website and the individual chain of advertisements they have seen, along with search queries and other traceable interactions. The attribution process then attempts to slice up the value of the user's visit and estimate the extent to which each media touch point in the user's journey contributed to their actions.

There are two main approaches used in building out an MTA system: the Markov approach and the Shapley Value approach.¹ The Markovian approach utilises a stochastic process framework where the probability of a consumer moving from

one state to another is the main driver of the system; the Shapley Value approach is derived from game theory, where the marginal contribution of each touch point is estimated. Both approaches have their advantages and limitations. The Shapley Value approach has been widely used in recent years due to the fact that the sequence of events is not a consideration, making it more computationally straightforward. Despite these advantages, the Shapley Value method is limited in the number of channels it can address because there is an exponential growth in computation with each new variable added. The Markov approach has the benefit of considering the order of interactions and being able to consider more granularity in

the touch points included. Regardless of the method used, MTA results provide insight into the actual value of touch points in a consumer's journey.

Once the value of each touch point is understood, then discussions around media return on investment (ROI, optimal mix and multi-channel strategies can begin. The inherent problem with this process is that digital media optimisation and attribution reporting assumes a point of sale on the website. Media are then judged on their ability to drive consumers to a trackable purchase. For a number of industries, however, this is simply not how it works. In a field such as hospitality and travel, consumers rely upon both third-party and brand sites for information as well as direct conversion. For highly considered purchases, consumers will utilise a brand's website for information, comparison and decision-making purposes. This is particularly true across financial products, home builders and in the automotive space. By examining such data properly, it becomes clear that consumers are providing considerable information about their interest and likelihood to purchase.

In general, consumer behavioural data can be a gold mine for marketers. Through their behaviour, consumers are constantly providing brands with feedback and a truer measure of their intentions than can be gathered from surveys. In this mindset, one can view consumers' behaviours as an indication of purchase intent that can be used to quantify the value of a visit. Once the dollar value of each visit is known, that information can serve as the basis for attribution value without a directly trackable purchase.

In the automotive sector, the lack of e-commerce is consistent across the industry. Consumers are able to submit leads to be contacted by a dealer and for years this has become the proxy to sales. The issue with this process is that only a fraction of consumers actually submit leads. In many consumers' minds, car dealers do not have

the best reputation for interacting with potential customers, so consumers are often hesitant to provide contact information. Thus, attempting to optimise media by using leads as a proxy for sales not only uses just a fraction of customers, it also means that the data sample is almost certainly biased. Consumers who are willing to provide details to a dealer are not representative of all car shoppers.

To move beyond optimising by leads, marketing analysts and individual car dealerships have come together to develop a broad analysis of dealer digital spend by individual digital marketing channel. The goal was to involve the dealer community to gain insights around consumer behaviour and validate the methodology and results. A coalition of 300 dealers allowed complete access to their website tracking and sales history in order to develop results which would benefit the industry and promote transparency and marketing spend efficacy. This coalition of dealers represented 420,000 consumers, 3.5m consumer interactions and US\$72m in annual media spend.²

SOLUTION

To create a concrete proxy for sales, the study set out to evaluate the impact of each aspect of the dealer's website on the likelihood of purchasing. Across all of the dealerships, consumers' purchase behaviour was followed for all known consumers. A known consumer is someone who had a previous relationship, be it a prior sale or a service visit, with the dealership. Within the known customer base, those who were trackable online had their website behaviour joined with their purchase and service history.

This allowed for the creation of a consumer journey. Some customers will only have a prior sale and then a purchase in the analysis time frame. Others will have repeated service behaviour prior to a sale. These two aforementioned groups could proceed to a sale without any trackable

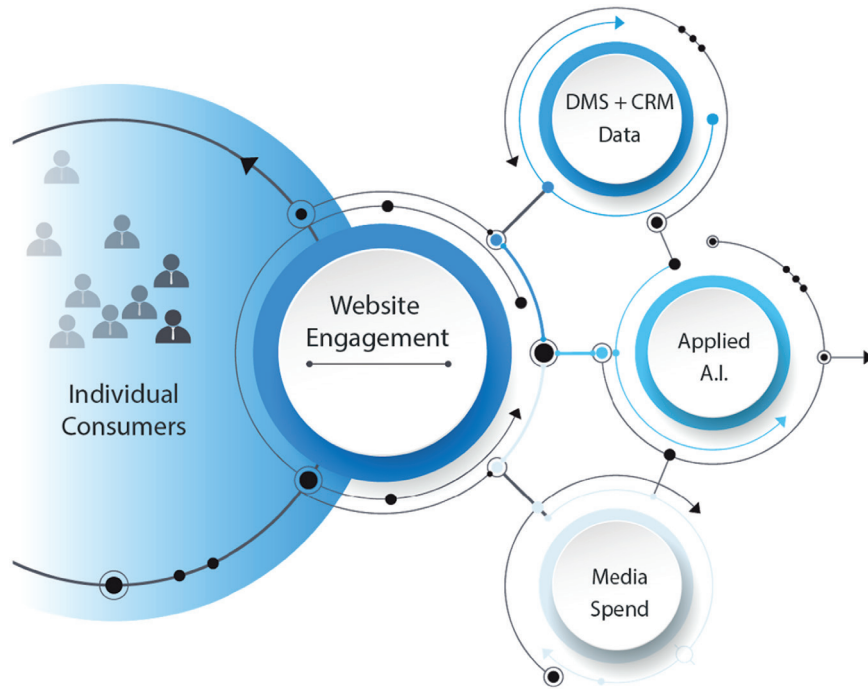


Figure 1: The role of website engagement in the attribution process
Key: DMS - Dealer Management System

web activities. This showcases the purchase (and repair) rates of all customers and those who never visit the dealer's website. These baseline sales and service rates became the benchmark for comparison.

The analysis took a multi-level approach to media optimisation. Where direct attribution was available, a website analysis through removal impact³ was utilised. Later, to address the impact of media execution on sales, a traditional media mix analysis was utilised. This enabled not only a historic analysis but also utilised the relationships uncovered in the models to produce ongoing reporting, ROI and optimisation opportunities. Markov and generalised logistic models were built using a combination of R and Python, drawing upon various libraries and approaches through those platforms.

The website attribution enables the optimal allocation of resources to marketing tactics. It provides the linkage between sales data and media spend on an individual basis (see Figure 1). By understanding the true

value of each touch point on a consumer's likelihood to convert, marketing dollars are invested in media that drive results. The historic analysis of data also provides valuable insight into a consumer's path and begins the conversation around 'what is the next best action for this consumer?'.

A Markov model was used to look at each consumer's path to purchase. Every interaction with the dealer became a step in the consumer's journey: oil change, e-mail sent, website visit, tyre replacement, sale. As consumers are exposed to or consume more media, the combination of their past exposure defines their current state. The likelihood that a consumer will move toward the desired result is defined by a transition matrix.⁴ Prior history is retained based on a multiplicative decay of the initial state reduced based on the n-degree of repeated exposure. Thus, the history is utilised to inform the present state to create a Markov chain. This chain moves through time whenever the consumer has a new interaction. The state space is defined by the




Figure 2: Visual example of removal analysis
Key: VDP - vehicle detail page views

various possible combinations of interaction present in the training data.

In this case, as direct attribution is available, an analysis through removal impact was utilised to define the transition matrix. This approach groups customer journeys into distinct paths based on the timing and amount of website interaction.

Each path is then compared against other paths missing one element, thereby creating a series of what-if scenarios to understand how the different combinations impact conversion to a sale (see Figure 2). The differences in these conversion rates are compared when part of the path is removed.

The mechanics of this process lie within the transition matrix. The matrix provides a measure of how each step in a customer's path influences their likelihood to convert. The true benefits of MTA lie in its ability to provide an accurate view of consumer actions and outreach performance. It is from this perspective that the model transverses different hierarchies and granularities to provide actionable insights to buyers, planners and strategists. Separate transition matrices were developed within segments in order to provide insights into customer differences. In this case, four strata were created based on the size of the dealership, with separate matrices developed for each strata.

Once the transition matrices were developed, they were applied against the data to create an attribution of vehicle sales. Now each category (service visit, e-mail opened, website visits with an inventory search) receives a portion of the vehicle sales within each dealership. The benefits of this process are that customer-level attribution infers causation rather than simply

correlation. The difficulty actually lies in the data linkage as it can be difficult to establish linked customer journeys.

The Markov approach was used for the website attribution for two main reasons. The first is that the order of site activity may have an impact on the value of the activity. For instance, a vehicle detail page view that comes directly into the site via a third-party link may have less of an impact on purchase than a vehicle detail page view that follows a visit to the financing section followed by an inventory search. The second reason is that 18 separate touch points were included in the analysis. Under the Shapley Value method, this would create a 262,000 level of interaction, which is computationally unfeasible.

FINDINGS

Given this information, total sales are decomposed into separate buckets (see Figure 3).

The baseline sales are defined as those without an action, or change of state, in the observation period. Then, all other sales have their credit spread among the various touch points. Prior sales and ongoing service visits draw a large portion of sales.

The focus of the analysis is how the remaining portion of sales is distributed across consumer interactions on a dealer's website. The three largest drivers of sales are not surprising. Inventory searches, specials (including incentives) and vehicle detail page views (VDPs) account for the lion's share of sales. The remaining sales are spread among the remaining 13 tracked web activities. The benefit of the removal impact of the Markov process shines in this regard. As the process accounts for a portion of baseline sales as well as sales driven by purchase and repair history, the sales attributed to web activities are treated as incremental.

As the number of sales driven by each activity is known as well as how many activities were performed by consumers, this translates into a per activity value. Inverting




Figure 3: Decomposed sales volumes
Key: VDP - Vehicle detail page views

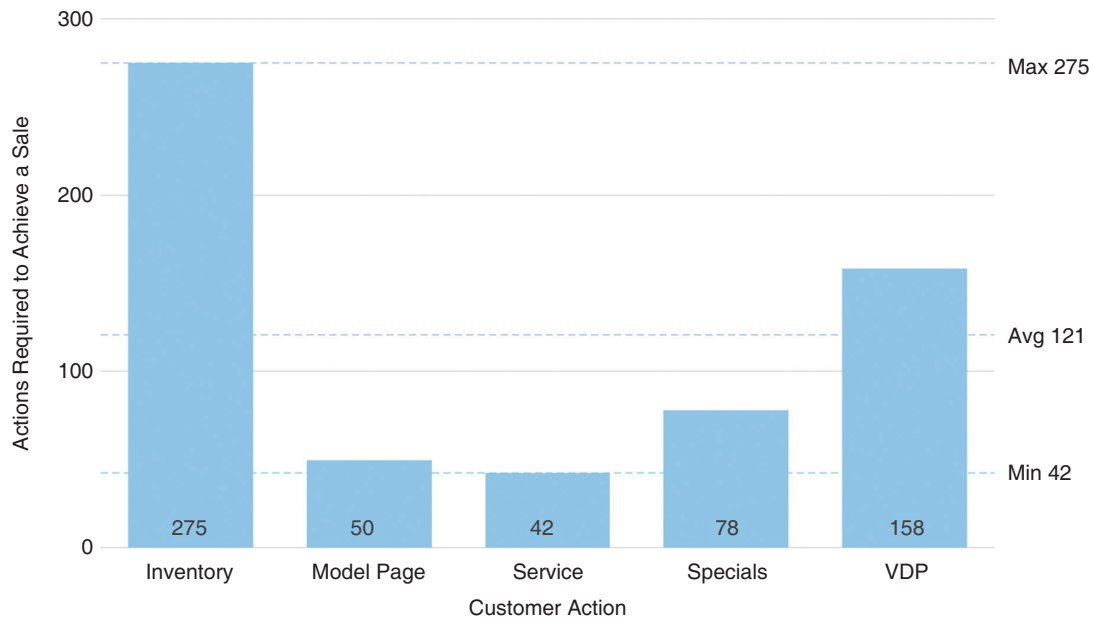


Figure 4: Per activity value

this provides the number of activities required to drive an incremental sale. The top three performers in terms of driving sales (inventory, VDPs and specials) require the largest number of activities to equate to a sale. This implies that these actions are an intrinsic part of the sales process and also that

they are undertaken by most consumers. For this reason, an inventory view in and of itself provides no clear indication of a consumer’s likelihood to purchase; taken in aggregate, however, the number of inventory visits is indicative of overall sales interest (see Figure 4).

Taking things a step further, once the value per website activity is known then tracking of the impact of media becomes a straightforward process. At this point, all media going to the site can receive a dollar value per visit based on the count of the high-value activities identified, weighted by their sales contribution. When combined with media cost, the value delivered can create an attributed ROI by channel. The attributed ROI calculation is simply:

$$\frac{\text{Value} - \text{Cost}}{\text{Cost}}$$

With this formulation, attributed ROI showcases the amount of incremental profit received for each dollar spent. Thus, media with an attributed ROI of less than US\$1 indicates that a dealer is losing money while an attributed ROI of greater than US\$1 shows that a dealer is making money on the media investment. Media data are examined at the most granular level available, the key principle being that media are estimated at the level at which tactical decisions can be made. This may be at the channel level, provider

level or occasionally the placement level. The goal is to balance the robustness of the model with the optimal level of granularity.

MEDIA PERFORMANCE

For the purposes of this study, media were examined at the channel level in order to understand the attributed ROI across each channel (see Figure 5).

This becomes important not only in reporting but also in optimisation efforts. The study results were in line with generally accepted views surrounding digital media. Media falling into the category of awareness has the lowest level of attributed ROI. In the case of acquisition display media, the study showcased the lowest performance. Given that this media type is targeting owners of competitive vehicles, a weak attributed ROI is not surprising. Search engine marketing (SEM) is also a low performer. This is often the case as SEM is used as a marketing tool, a vehicle for conquest and as a navigation tool for consumers. Media cost per click (CPC), is the strongest performer in the awareness

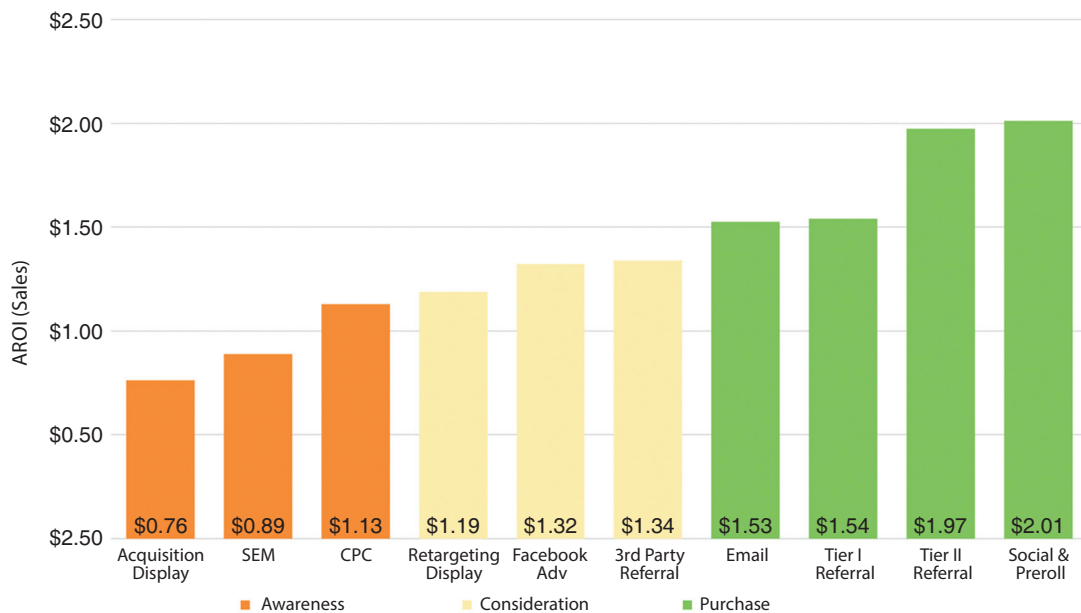


Figure 5: Attributed ROI
Key: SEM - search engine marketing, CPC - Media cost per click

category. The return on CPC is the only awareness medium that saw a positive return.

Moving into the consideration phase of marketing, stronger performers emerge. Display retargeting, Facebook advertising and third-party retargeting all show positive returns. Retargeting media often emerges as a costly digital spend but one that drives a number of consumer actions with high sales contributions. Facebook advertising is an area where spends are starting to become more consistent across dealerships. Third-party referrals have long been a solid driver for dealerships which initiate a handoff from external shopping sites to dealer inventory.

Finally, strongest overall performance is seen in the media purchase phase. Referrals coming from the original equipment manufacturer (Tier I) and regional dealer groups (Tier II) show rather strong results. Much like with third-party referrals, these sources transition consumers from the shopping phase to the purchase phase. E-mail also has strong performance given

the nature of the medium. Well executed e-mail programmes can represent an ongoing relationship between a consumer and a dealer. The highest performer overall is social media (excluding Facebook) and video preroll. While these are stronger performers, one is left to wonder how the level of spending in this channel is impacting performance. If the analysis showcasing strong results is because dealers are only putting a small amount of budget behind these media types and therefore reaching consumers who are most likely to be impacted, then how scalable are these results?

MEDIA OPTIMISATION

As all media can now be tracked in terms of the value delivered and the attributed ROI, the derived value and media spend was integrated into logistic curves to determine optimal thresholds by channel (see Figure 6).

The generalised logistic function was used to model the inflection points, the point of diminishing returns and saturation level.

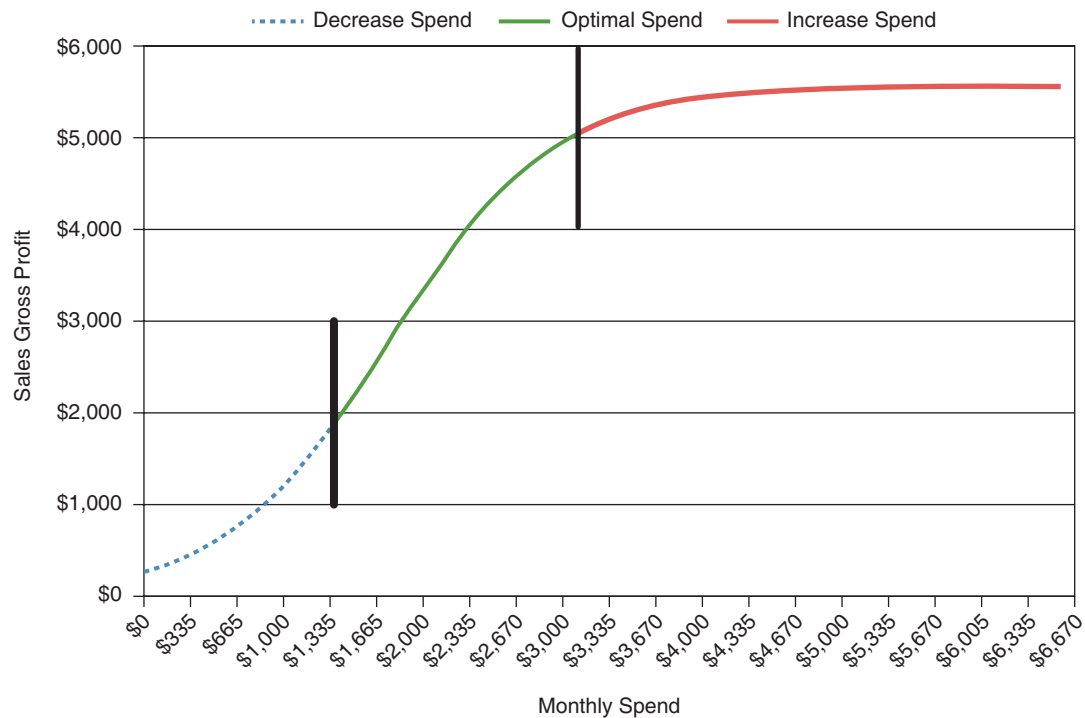


Figure 6: Diminishing returns curve

For each media type, the model estimated the asymptote, inflection point and growth rate. Throughout the model development stages, a classical training and validation approach was utilised to determine model effectiveness.⁵ In Figure 6, the area to the left represents levels of media spending where dealers have not yet broken through the media environment to the point at which they have achieved economies of scale. The area to the right represents levels of media spending where dealers are overspending past the point of diminished return. The area in the middle represents the healthy level of media spending where dealers are receiving solid returns on their investment. Across the study of automotive dealers, there were significant differences in media spending performance across channels, and a varying degree of over and underspending. Half of all dealers are overspending, but across a range of channels.

Recall that acquisition display media, which falls into the category of awareness, had the lowest level of attributed ROI at US\$0.76. While 25 per cent of dealers were found to be overspending on acquisition display media, most dealers were underspending. This presents an opportunity for dealers who are able to optimise the ROI on other better-performing media in to invest in acquisition display media and potentially grow their consumer base. SEM was also a low performer, with an attributed ROI level of US\$0.89. Thirty-four per cent of dealers were overspending on SEM. This is a trend seen across the industry as dealers are hesitant not to purchase all available terms in order to defend against conquest by other dealers. The return on CPC was the only awareness medium that saw a positive return, with an attributed ROI of US\$1.13 and with only 13 per cent of dealers overspending.

In the consideration phase of marketing, display retargeting and Facebook advertising have similar positive returns. The attributed ROIs are US\$1.19 and US\$1.32, respectively.

Twenty-five per cent of dealers are overspending on both media types, but Facebook shows room for growth, with the majority of dealers underspending on this medium. Third-party referrals show an attributed ROI of US\$1.34, with only 19 per cent of dealers found to be overspending.

The true power of diminishing returns curves is in the concept of marginal return. The attributed ROI values given look at the value driven by the media spend in total. However, when the concept of marginal returns is realised, the last dollar spent does not have the same value as either the first dollar or the middle dollar. This allows for a finer understanding of where to invest in media across channels.

IMPACT

Once dealerships employ the findings of the website attribution study, they can determine the incremental value delivered by media. This allows the creation of dashboards, reporting and alerting to be put in place around specific metrics such as ROI or cost per lead (Figure 7). When combined with diminishing returns, dealers can convert traditional media reporting into a prescriptive media optimisation tool. By understanding the incremental performance of each of the media channels, marketing managers can make surgical changes to their media spending. The output of this process serves to guide media teams to move resources around to maximise the ROI of the entire programme. Users could quickly and easily identify which channels and placements for which spending should be decreased as well as the channels and placements where spending should be increased.

Ultimately, media optimisation is achievable even in the absence of e-commerce on a website. Given a known and trackable consumer base, combining offline purchase behaviour with online consumer activities through a Markov



Media Spend and ROI by Location

\$1.74

Total ROI

Location Name	Channel	Combined ROI (Sales & Service)	Imp	Clicks	Cost	Est. Sales Profit	Est. Service Profit	Attributed Sales	Attributed ROs	Cost per Sale	Cost per RO
Location 1											
	Display	\$71.29	1,206	0	\$2.10	\$111.00	\$40.80	0.06	0.48	\$35.00	\$4.38
	Search	\$22.75	564	16	\$44.75	\$777.00	\$285.60	0.42	3.36	\$106.55	\$13.32
Location 1 Total		\$24.92	1,770	16	\$46.85	\$888.00	\$326.40	0.48	3.84	\$97.60	\$12.20
Location 2											
	Display	-\$0.94	1,092,945	1,645	\$8,827.00	\$407.00	\$124.10	0.22	1.46	\$40,122.73	\$6,045.89
	Search	\$19.05	9,376	427	\$1,243.54	\$19,499.00	\$5,428.10	10.54	63.86	\$117.98	\$19.47
Location 2 Total		\$1.53	1,102,321	2,072	\$10,070.54	\$19,906.00	\$5,552.20	10.76	65.32	\$935.92	\$154.17
Location 3											
	Display	-\$0.97	92,728	208	\$819.24	\$18.50	\$6.80	0.01	0.08	\$81,924.00	\$10,240.50
	Search	\$3.14	13,896	654	\$1,893.31	\$6,308.50	\$1,528.30	3.41	17.98	\$555.22	\$105.30
Location 3 Total		\$1.90	106,624	862	\$2,712.55	\$6,327.00	\$1,535.10	3.42	18.06	\$793.14	\$150.20

Figure 7: Ongoing media reporting

process creates a set of per activity values. These activity values enable the calculation of ROI through attributed ROI. These activity values also can be integrated with media spends to model diminishing returns functions. The combination of accurate activity value and marginal returns puts very powerful tools in the hands of marketers who for too long have had little choice other than to market to lead generation.

References

1. Anderl, E., Becker, I., Wangenheim, F.V. and Schumann, J.H. (2014) 'Mapping the customer journey: a graph-based framework for online attribution modeling', *International Journal of Research in Marketing*, Vol. 33, No. 3, pp. 457–474.
2. Harper, J., Wethington, M. and Moore, D. (2018) 'New study reveals that half of auto dealers overspend on marketing, with no measurable return on investment', available at: <https://www.globenewswire.com/news-release/2018/12/17/1668156/0/en/New-Study-Reveals-That-Half-of-Auto-Dealers-Overspend-on-Marketing-With-No-Measurable-Return-On-Investment.html> (accessed 10th May, 2019).
3. Kakalejčík, L., Bucko, J., Resende, P.A.A. and Ferencova, M. (2018) 'Multichannel marketing attribution using Markov chains', *Journal of Applied Management and Investments*, Vol. 7, No. 1, pp. 49–60.
4. Bickenbach, F. and Bode, E. (2001) 'Markov or not Markov — This should be a question', Kiel Institute for the World Economy, Working Paper No. 1086, available at: <https://www.econstor.eu/bitstream/10419/2673/1/kap1086.pdf> (accessed 15th May, 2019).
5. Varma, S. and Simon, R. (2006) 'Bias in error estimation when using cross-validation for model selection', *BMC Bioinformatics*, Vol. 7, No. 1, pp. 7–91.