

Combining marketing mix and attribution models

Attribution modelling has emerged as a more agile, dynamic and granular alternative to the lumbering, macro marketing mix models of the past and moves us on from measuring to managing ROI. But while the science is still raw, the challenge facing the industry is how to integrate the best of both approaches

By Alice Sylvester and Jim Spaeth,
Sequent Partners

As recently as three years ago, most marketers used marketing mix models for assessing the return on media investments. That practice was extremely pervasive – across almost all industries. But there has been a swift and very striking change in marketers' expectations in ROI media performance measurement. Today, mix models are viewed as too slow, too macro and too backward thinking for the boots on the ground.

Instead, the push is for a more detailed, granular and integrated top-down/bottom-up kind of view that tells us what to do differently. The gauntlet has been thrown down. In study after study (close to a dozen separate studies over the past three years), brand marketers have told us they are looking at adapting marketing mix modelling into something more dynamic and useful – more like attribution.

Attribution modelling has emerged as an alternative to lumbering market mix models. Agile and dynamic, attribution modelling resides mostly in the digital

ecosystem, crossing search, display, email, social, and websites – but with the right data, promises to also handle cross-platform ROI analysis. Granularity and speed make attribution very appealing.

Why the sudden adoption of attribution? There are several reasons. First, consumers are spending more time with digital devices and content – and advertisers are spending far more money there. As a result, advertisers realise they need to coordinate their digital and traditional media (at least television) correctly. They need to accurately understand the trade-offs, potential synergies and ultimate returns of these investments at a granular, actionable level.

A second reason for the urgency may be the most important. We have consistently heard that advertising returns and profitability are down. An advertising ROI crisis appears to be emerging. A recent Nielsen Catalina Solutions meta-analysis of 1,400 studies across 450 brands showed that the average return on ad spending was about \$2.50. While that sounds great, note that this is the revenue generated by the ad spending. When you apply typical margins to that figure, you find that the average returns are not profitable in the short term. Attribution

may provide the guidance that marketers finally need to manage, rather than measure, advertising performance.

Although the conversation about marketing mix models and attribution takes place in the same space, namely marketing analytics, they are not the same thing, as one marketer recently suggested. They are not interchangeable; rather, they are quite distinct, with different strengths and limitations. The challenge facing the industry is how to integrate the best of both approaches.

At the end of the day, marketers want to reach the right person at the right time with the right message. And they need the tools and systems to do that dynamically and in real time – in campaign. But that's a lot harder than it looks.

Right now, marketers appear to be using both tools – marketing mix and attribution models. But results rarely agree and marketers don't know which results to follow. The need to unify mix and attribution models is keen, but the science is just emerging.

There may be an elegant, holistic solution down the road, but the best practice we have seen to date is layering attribution beneath marketing mix (Figure 1, next page).

This technique takes the best aspects of marketing mix models and attribution and links them to business outcomes and measures. The layering approach includes four distinct levels. At the top are the brands' KPIs – the metrics that guide the business – sales and sales-related behavioural data along with key brand health metrics. This layer doesn't require advanced analytics, just good measurement. Importantly, this layer frames the brand's reality and modelling must align with this reality to be useful. And that's a hint. As we use analytics to dig beneath the surface and reveal the causal factors underneath, they must align with the reality above.

FIGURE 1: LAYERING MARKETING MIX AND ATTRIBUTION ANALYTICS

LAYER	FOUNDATION	METRICS	PARAMETERS	RESULTS
BRAND KPIs	Measurement via sample or customer database	Sales, market share, penetration, brand awareness	Monthly, weekly, national, regional, sales territories	Measures business financial performance
MARKETING MIX MODELS	Econometric models	Contribution to sales from: baseline, promotion, advertising by medium	Data: weekly/market level or store level Results: annually or quarterly	Measures marketing factor sales lift and ROI. Ties back to sales performance
GRANULAR MIX MODELS	Econometric models	Contribution to sales from: specific media vehicles, creative executions	Data: daily/zip based Results: quarterly, bimonthly	Measures granular marketing factor sales lift and ROI. Ties back to marketing mix model elasticities
ATTRIBUTION MODELS	Data science	Response to specific ad exposures by: vehicles, ad, target audience	Data: event level Results: near real time	Optimises short-term performance of one or more related media types

Marketing mix modelling provides the second layer. It reveals the drivers of the brand's marketplace results. The model replicates the brand's reality by correctly identifying the contribution of each of the underlying causal factors. This contribution and cost of each element of the marketing mix yields its ROI.

The fact that a marketing mix model accounts for all of the major causal factors and explains most of the variance in the brand's sales is its great strength. It gives the organisation, particularly finance, confidence. The other great strength of modelling at this level is that it evaluates all of the brand's marketing expenditures on the same basis.

Although marketers can be confident that mix models accurately replicate the brand's reality, and they can comparably assess most of the marketing mix, they are hampered by the limited ability to drill deeply into the factors that caused the outcome. Marketing mix is broad and comprehensive, but shallow. Generally, only broad media types are evaluated. Drilling down to see the relative contribution of specific media vehicles or ad executions is typically beyond reach.

However, traditional mix modelling techniques are able to dig more deeply. And this is the third layer of the analytic stack – more granular mix models. This

layer is uncommon, but quite effective. This approach recognises that the biggest limitation in marketing mix models is not the analytic technique employed; it's the data that fuels the models.

Most often, mix models are built on market/week data or store/week data. But we have seen models built at a far more granular level – days and finer levels of geography such as zip code-defined trading areas. Recent work by IRI has shown how modelling at finer levels can produce a strong estimate of the contribution of cinema or out-of-home advertising, two media whose effects are highly concentrated geographically, as compared to television or radio. More granularity in mix models also provides enough degrees of freedom and real variation to tease apart the effects of different executions in a campaign. This is not something marketing mix models have been called upon to do. In fact, even though most modellers and marketers acknowledge the importance of measuring creative executions (about 75% of the lift generated by advertising, according to the recent ARF Ground Truth Experiments), marketing mix models most often do not measure the effect of individual creative executions.

Although definitely an improvement, this third analytic layer – more granular marketing mix models – is still not the ultimate analytic

solution for many marketers. These models can't isolate the effect of the right message to the right consumer in the right moment – the advertisers' mantra. They also fail to recognise the role each medium or ad plays. They simply consider the presence of each element and its impact in the outcome variable, not the specific role it plays in the mix. And finally, the lag between measurement, analysis and activation is too great to have maximum effect on the campaign.

Attribution provides the fourth and bottom layer of the analytical stack. Attribution modelling was born in the digital world to evaluate event-level data in near real time to drive the next ad placement decision in a continuous improvement framework. It potentially checks all three boxes: right message to the right consumer in the right moment. Potentially.

There have been improvements in attribution science recently, but challenges remain. In its earliest incarnation, attribution was entirely arbitrary. Credit for sales generation was assigned a priori – last click, first click, weighted sequence of clicks, with no validation and no rationale.

More recently, the best attribution models are based on sound statistical modelling with parameters fitted to real data. That said, there is no consensus yet on which techniques are most suitable. The

techniques in use range from time-series models to game theory approaches and there is not always the transparency we would expect from providers. This may be changing – we expect this will settle out under increased scrutiny in the near term. Iterative, self-adjusting algorithms and machine-learning artificial intelligence will ultimately provide better accuracy and drive more sensitive forecasts.

The next challenge in attribution is busting the media silos between digital and traditional media, mainly television. Because attribution operates solely in the digital media silo, the models are unaware of the effect of traditional media, the rest of the marketing mix and the brand itself. This isolation results in gross overstatements of the impact of digital media.

For example, if a brand runs a flight of banner advertising to drive retail traffic to the store, this investment will yield much more positive results if it follows a flight of television advertising that generates awareness, interest and a predisposition to the brand. Attribution tends to credit all of

not readers, can be brought into the data stacks. And linking these disparate data sources is an enormous task. One marketer told us they “spent eighteen months cleaning and matching the data – matching person-level digital data to household-level television data”. The process “introduces error at every step. It’s very difficult – and the data will never be apples to apples.”

Further, the impact of non-media marketing factors can be extremely important but tends to be outside the realm of attribution modelling. Trade promotion, a critical component of consumer packaged goods marketing, can be imputed based on geography, but the lack of clean store data – not knowing which stores were shopped – makes that a noisy estimate.

Since attribution operates at such a granular, often household level, modellers face the challenge of identifying all of the digital devices the household uses. In the digisphere, devices are observed as independent agents through cookies or IDs. But it’s more complicated and device graphs are required to help you observe

Net, attribution modelling offers the ultimate in granularity and timeliness, promising clear and immediate guidance for the continuous improvement of campaign performance. But right now, there is a lot of coverage bias, imputation, and estimation that goes into that process. It will get better, but we are not sure when. A modeller at a major advertising company told us: “We’re four years into this journey – and we have a long way to go.” Yet they remain optimistic. “We are working through it; we have promising things in the works.”

THE WAY FORWARD

Attribution is a powerful addition to our analytical toolkit for improving advertising ROI. It can add a deep layer to marketing mix models for advertisers who want a more granular and timely read on the performance of their entire marketing mix. In an activation setting, for an agency or media provider, it can drive better ad choice and placement decisions in a continuous improvement process.

But issues and challenges raised in this paper all need to be exposed to more daylight. We have seen a lot of hand-waving, or arrogant announcements that the world has changed and these are the new answers. Science doesn’t work that way. Not even marketing science, or data science. We need experimentation and validation. We will make the fastest progress if this is done publicly and the shared learning accelerates progress.

Marketing mix modelling was a remarkable breakthrough. Twenty-five years ago, we could finally measure the ROI of marketing and figure out if the new campaign generated more sales than the old one, or if the money spent in television performed better than the investment in magazines. But it’s time for the next step – from measuring ROI to managing ROI. That’s the promise of attribution modelling. Like any development of this magnitude in its early days, a lot of work lies ahead.

“There may be an elegant, holistic solution down the road, but the best practice we have seen to date is layering attribution beneath marketing mix”

the incremental sales to the display ads. You can imagine similar types of misattribution arising from promotional activity or price reductions or any other marketing activity to which the attribution analysis is blind.

Set-top box data offers a solution for television, but the available data is incomplete and biased to some degree by the lack of full coverage. The other traditional media have proven to be more difficult to integrate and, as a result, modellers generally rely on some form of imputation when including them. Spotify and Pandora bring in digital radio data, but not terrestrial radio, and magazine subscribers,

which of the phones, tablets and computers belong to a given household. Only then can you understand the full string of exposures impacting the household’s brand choices.

And that brings us to the final issue with attribution. There are no people in attribution modelling. This is a fact most modellers, especially Big Data scientists, simply shrug off. There are only devices and households, which is tolerated from an analytic and measurement standpoint, but which puts the advertising and consumer insights industry back in the 1950s. It’s almost inconceivable that in 2016, it’s acceptable to target households not humans.