Marketing Science Institute Working Paper Series 2026

Report No. 26-106

Open-Source Media and Marketing Mix Modeling: Overview, Challenges, Opportunities

Julian Runge and Koen Pauwels

Open-Source Media and Marketing Mix Modeling: Overview, Challenges, Opportunities

Julian Runge

Northwestern University
Medill School of Journalism, Media, Integrated Marketing Communications
1845 Sheridan Rd., Evanston, IL 60208, United States
julian.runge@northwestern.edu

Koen Pauwels

Northeastern University
D'Amore-McKim School of Business
370 Huntington Ave, Boston, MA 02115, United States
kpauwels@northeastern.edu

Abstract

Privacy-enhancing technologies are reshaping online advertising, making it increasingly difficult for advertisers to track customer journeys and link marketing actions to economic outcomes. As a result, interest in probabilistic measurement techniques such as media and marketing mix modeling (m/MMM) has surged, particularly among digital-first advertisers. Many of these advertisers are small and midsize firms that rely heavily on digital channels yet lack the resources to commission bespoke proprietary models. Against this backdrop, grassroots analytics communities and large online advertising platforms such as Meta Platforms and Google have started developing open-source packages for m/MMM. The packages share core principles but differ on important aspects such as their degree of automation, ease of use, and modeling paradigms. This article provides a narrative review and practice-oriented synthesis of these developments and discusses the challenges and opportunities that come with "packaged up" m/MMM. It concludes with an outline of topics that marketing scholars can engage with to vet and extend open-source m/MMM, and to align its continued development with its users' needs.

Keywords: media mix modeling, marketing mix modeling, open source, computational packages, applied marketing analytics

Acknowledgements: The authors are grateful to Elea McDonnell Feit, Igor Skokan, William Grosso, Mark Rieke, Michael Kaminsky, Brendan Sinnott, Jordan Nafa and Gufeng Zhou for comments on earlier versions of the article.

1. Introduction

Marketing organizations have long turned to media and marketing mix modeling (m/MMM, Arora et al. 2023) as a systematic approach for quantifying how marketing investments translate into customer responses and revenue. These models have shaped not only academic understanding of consumer behavior and marketing strategy but also organizational practices for budgeting, forecasting, and evaluating the return on marketing spending (Borden 1964; Lambin 1972; Little 1975; Jagpal and Brick 1982). Despite their importance, m/MMM implementations have historically been costly, opaque, and confined to a limited set of specialized vendors. This creates a gap between the methodological advances documented in scholarly research and the tools readily available to practitioners, possibly constraining organizations' ability to align marketing resource allocation with evolving customer needs.

Recently, this gap has sprung to the forefront of digital marketing practice as privacy-enhancing technologies are making it difficult, and in many cases impossible, to directly measure and observe digital customer journeys (Johnson, Shriver, and Du 2020; Johnson, Runge, and Seufert 2022). Bottom-up approaches from user-specific data are less feasible, making top-down approaches more appealing (Pauwels and Aksehirli 2025). Many online and digital-first advertisers, however, almost exclusively relied on these deterministic measurement and attribution approaches that are increasingly being rendered obsolete. The emergence of open-source m/MMM tools offers a promising response to this challenge. Grassroots analytics communities and large online advertising platforms such as Meta Platforms and Google have started developing open-source packages that aim to fill this measurement gap and support digital advertisers in adopting m/MMM.[1]

The first foray to offer an open-source package for m/MMM was Facebook's release of Robyn in the early 2020s (Meta Platforms 2024; Runge et al. 2024; CRAN 2025). Google followed suit with Lightweight (Google 2023) and Meridian (Google 2025; Nair 2025). PyMC is a longstanding package for Bayesian modeling, first initiated in 2007. In November 2022, PyMC Labs released PyMC-Marketing that specifically focuses on marketing analytics, including PyMC-Marketing-mmm for m/MMM (PyMC Labs 2025). Robyn's entry triggered both enthusiasm and skepticism. On one hand, it represents a philosophical and practical shift: from black-box modeling to transparent and fast estimation with a high degree of automation. On the other hand, it poses methodological and managerial challenges: Is open-source m/MMM "good enough" for high-stakes budget decisions? Can firms without deep analytics teams use it responsibly and understand model assumptions and performance adequately?

Providing accurate and convincing automation for causal inference problems such as m/MMM is challenging. Some components of the open-source packages are novel and unvetted, e.g., the multi-objective optimization approach in Robyn (Runge et al. 2024). By providing a practice-oriented overview of open-source m/MMM, this article aims to enable marketing scholars to gauge strengths and weaknesses and facilitate dialogue and collaboration towards vetting and testing these packages. It is a narrative review and conceptual integration of recent developments in open-

[1] The incentive for the largest online advertising platforms to start developing these solutions is that measurement commonly is a prerequisite for investment. If digital advertisers are unable to measure their advertising on these platforms (because the deterministic attribution approaches they relied on are increasingly faulty and obsolete), they may cut their advertising spend . Another explanation is that the platforms could embrace methodological approaches in their open-source packages that evaluate advertising on their platform more favorably. In our review and testing of the packages, we could not find evidence of this latter explanation.

source m/MMM and bridges marketing practice and scholarship to guide and improve applied marketing measurement to meet advertisers' measurement needs (Deighton, Mela, and Moorman 2021). A main goal of the paper is to facilitate rigor in m/MMM while it is becoming increasingly accessible to practitioners that want to assess the incremental effect of marketing activities and optimize future marketing decisions. We also raise normative questions for scholars about the role of transparency, model vetting, and automation in probabilistic marketing measurement.

The article proceeds with conceptual and practical background on m/MMM before it provides a practice-oriented overview of the different packages. We then discuss the challenges and opportunities that come with "packaged up" m/MMM and conclude with an outlook on future directions in open-source m/MMM.

2. Conceptual Background

2.1. The origins of m/MMM and its renaissance

Media mix modeling (mMM), a narrower subdomain of the broader marketing mix modeling (MMM), focuses specifically on media mix allocations, whereas MMM encompasses all "4Ps" of the marketing mix (Arora et al. 2023).[2] Almost as old as the term marketing mix itself that was coined by Neil Borden in 1949 (Borden 1964), MMM became a predominant analytical approach to propose strategic refinements to a company's marketing actions in pricing, promotion, product and distribution (Vidale and Wolfe 1957; Little and Lodish 1969; Lambin 1972; Little 1975; Jagpal and Brick 1982; Mela, Gupta, and Lehmann 1997; Jedidi, Mela, and Gupta 1999; Ding et al. 2020).

[2] We follow Arora et al. (2023) and adopt m/MMM as an umbrella term comprising both media and marketing mix modeling, so both mMM and MMM.

In the historic marketing data landscape where individual consumers could not be deterministically tracked across media and company properties, m/MMM (so, the combined practice of mMM and MMM) leverages variation in sales and marketing mix data over time to estimate the effects of a firm's marketing strategy choices on relevant outcomes. Advanced m/MMM integrates short-term and long-term effects and models competitor actions and dynamic synergy starting from an interactive market system framework (Naik and Raman 2003; Pauwels 2004; Cain 2022). To reduce bias and increase precision, models can also be calibrated against experimental estimates (Vidale and Wolfe 1957; Little and Lodish 1969), albeit it is unclear how widely adopted this practice is (Runge, Geinitz, and Ejdemyr 2020).

Recent accelerations in the rollout of privacy-enhancing technologies (Johnson, Shriver, and Du 2020; Johnson, Runge, and Seufert 2022) are fundamentally changing the data landscape in online advertising, challenging deterministic measurement approaches and bringing traditional probabilistic techniques en vogue again (Arora et al. 2023). While traditional large advertisers and brands are accustomed to, and tend to have deep expertise in, m/MMM, many digital-first and interactive advertisers have no experience with these models and often lack the resources to build the necessary expertise and systems internally (Gordon et al. 2021; Runge et al. 2024). These advertisers often heavily relied on digital advertising (Lee et al. 2025) and its increasingly dated measurement methods such as deterministic decision rules or digital attribution requiring cross-platform tracking, e.g., via third-party cookies (Li and Kannan 2014; Berman 2018). The property of m/MMM to not require user-level data is a foremost reason for its growing popularity among these smaller and midsize interactive advertisers (Arora et al. 2023).

2.2. Incrementality and choice of counterfactual in marketing measurement

Firms want to measure the incremental impact of their media spending on outcomes (e.g., sales) to guide their future budget and allocation decisions. Incremental outcomes are outcomes the firm would not have gotten without the ads, i.e. the counterfactual. This counterfactual could be regions where the ads did not run (geo-experiments), consumers who were not exposed to the ad (RCTs), or time periods when the firm spent less, more, or no money on the ad strategy in question (m/MMM). Because the counterfactual should be as similar as possible to the actual situation, the choice in practice comes down to which comparison is most feasible and cost effective. RCTs require a randomized hold-out sample of unexposed consumers, and cluster level-randomized experiments of unexposed clusters, while m/MMM requires variation of ad spending in past periods and control variables to account for other ways these past periods differed from each other. Modeling approaches that integrate variation over time and from experiments can be particularly powerful (Vidale and Wolfe 1957; Little and Lodish 1969; Wiesel, Pauwels, and Arts 2011; Hanssens and Pauwels 2016; Valenti et al. 2024).[3]

2.3. The related marketing analytics process

For effective data-driven decision-making in real-world marketing operations, m/MMM needs to be embedded in an analytics process that informs relevant actions, is understood and agreed upon by key stakeholders, and updated regularly. Marketing scholars may be less familiar with this important process, which is scarcely discussed in academic papers. We draw on Runge et al. (2024) to conceptualize such a marketing analytics process, as compatible with the open-source

[3] In m/MMM, this is often called calibration. Observational models are "calibrated" against results from experiments that can be expected to provide unbiased (albeit often only short-term) estimates of marketing effectiveness. See Runge, Patter, and Skokan (2023) for a hands-on guide.

implementations of m/MMM. Figure 1 depicts an iterative five-step process that starts with the firm *acting* (1) in the marketplace with media/marketing mix policy $\pi(t)$, running calibration experiments as planned, and making use of suited attribution models. The media/marketing mix policy can be expressed as specific levels of spend allocated to available media channels, tactical options within channels, other marketing actions, planned calibration experiments, and adopted attribution models, if any.

The firm then uses the observational, experimental, and attribution data (in and up to period t) to *estimate* (2.1) models of the causal effects of marketing actions on demand with different specifications. We frame this as the estimation of a set of specified models $M_i(t)$, some or all of which may include calibration against experimental or attribution results via priors (in Meridian; Google 2025) or multi-objective estimation (in Robyn; Facebook Experimental 2024, 2024a). The packages encourage users to consider a set of model diagnostics to assess statistical fit and business sensibility. We call this step inference as modelers are asked to ensure that the specified model(s) produce results that reasonably *infer* (2.2) the causal effects of the firm's media/marketing mix policy. If the diagnostics do not look reasonable, a new set of models $M_{i+1}(t)$ is specified and *estimated* (2.1), to then look at diagnostics and *infer* (2.2) again (motivating our choice of a two-way arrow in Figure 1).

In the third step, the firm needs to select a consensus model $m(t)$ that it believes to accurately model and reflect the incremental effects of its advertising on the considered outcomes, taking into account both statistical and business considerations. This model serves to prescribe a new media/marketing mix policy for the next period $t + 1$. In the packages, the prescription step (step 4) tends to be called budget optimization (in Meridian and PyMC-Marketing-mmm – Google 2025 and PyMC Labs 2025) or budget allocation (Robyn – Meta Platforms 2024). Then, the modeling chain repeats.

For each step, the figure lists its key output to be passed on to the next step, the key (human) skills needed to succeed with the step, and the key risk introduced by automation that needs to be mitigated through supervision by human expertise. The cadence for iteration can be monthly, quarterly, and sometimes weekly or half-yearly. Iteration cycles can be shorter for direct response-focused interactive advertisers and longer when longer-term efforts such as brand building are an important part of the firm's marketing strategy.

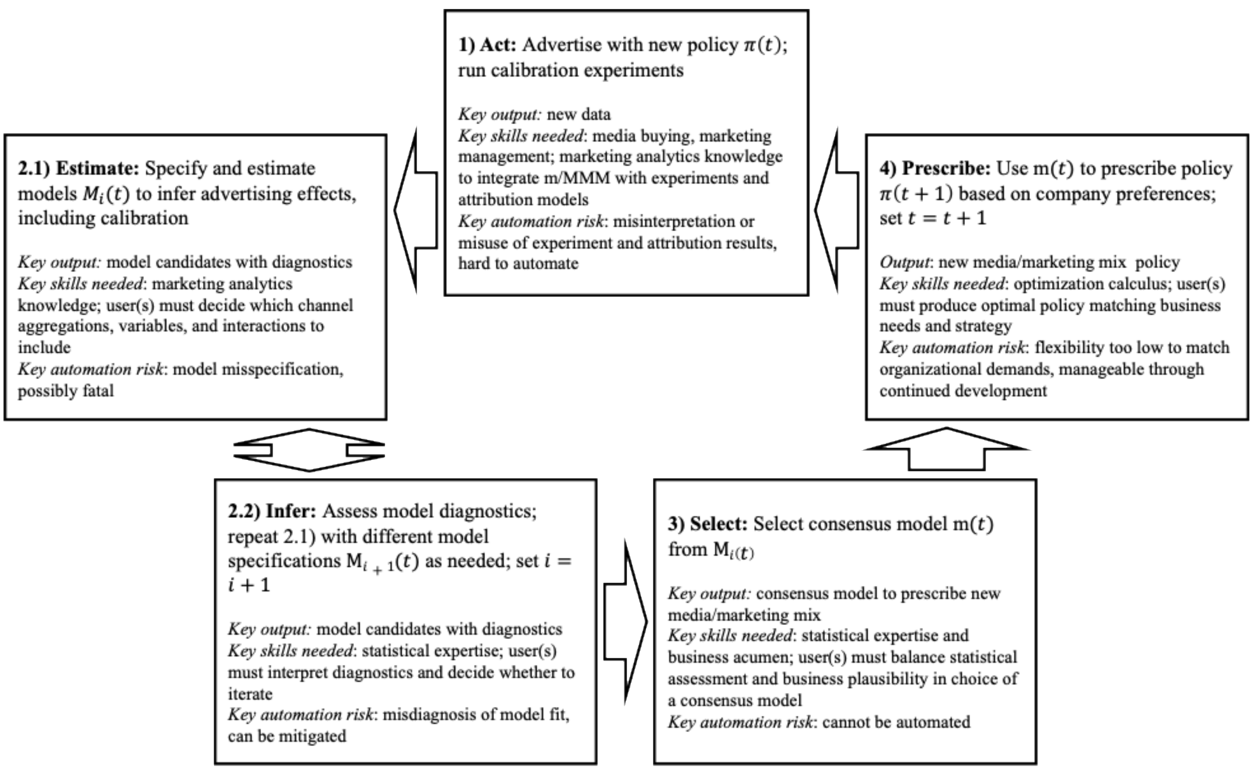

Figure 1: A conceptualization of the analytics process that open-source m/MMM packages can be embedded in, adapted from Runge et al. (2024). The process involves the steps (1) act, (2) estimate and infer (where iterative model tuning, calibration, validation and backtesting play an important role, hence the bidirectional arrow), (3) select, (4) prescribe – and act again in the next period. For each step, we note key output, key skills needed on the user(s)' part, and the key risk originating from automation and "packaged up" m/MMM.

The aim of this conceptualization of the m/MMM analytics process is to connect the discussed open-source packages to marketing analytics practice and provide a basis for structured analysis and comparison of the packages. Case in point, Appendix A illustrates how the analytics process in Figure 1 maps onto the computational workflow of Robyn, including code examples. We also refer the interested reader to Runge et al. (2024) that provides a more detailed exposition, including best practice recommendations and six mini case studies, in Section 4 of that paper.

3. Open-source m/MMM packages

3.1. A practice-oriented overview of the packages

There currently are four prominent open-source m/MMM packages: Meridian (Google 2025), LightweightMMM (Google 2023), Robyn (Facebook Experimental 2024, 2024a), and PyMC-Marketing-mmm (PyMC Labs 2025). In our further discussion, we focus on Meridian, Robyn, and PyMC-Marketing-mmm as Meridian has replaced LightweightMMM, which is no longer actively supported (Google 2023). Table 1 compares the packages on key computational elements as needed to support the analytics process depicted in Figure 1.

Meridian (Google 2025), Google's latest causal inference-based m/MMM framework (Nair 2025) and PyMC-Marketing-mmm (PyMC Labs 2025) use a Bayesian modeling approach that allows users to incorporate prior knowledge, e.g., experiment results. One of the strengths of hierarchical Bayesian models is their ability to leverage variation across geographies or other relevant hierarchies, making them a powerful tool for marketers running campaigns across multiple regions and wanting to account for related heterogeneity. The packages also include built-in support for return-on-investment (ROI) estimation, trend and seasonality adjustments, adstock and saturation[4], and budget optimization (Dekimpe and Hanssens 1995; 1999). Given their relatively higher level of customizability, Meridian and especially PyMC-Marketing-mmm are best suited for advanced users and marketing scientists.

Robyn (Facebook Experimental 2024, 2024a) is the first-to-be-released and arguably the most automated of the solutions. Unlike Google's and PyMC-Marketing-mmm's Bayesian frameworks, Robyn is based on Ridge regression (Hilt and Seegrist 1977) with hyperparameter tuning via evolutionary algorithms, requiring less advanced statistical expertise when using the package. Robyn is particularly useful for companies that need fast and automated media mix analysis and budget allocation recommendations. Its built-in carryover and diminishing returns modeling allows for realistic estimations of long-term campaign effects. Like Meridian, it does not allow for flexible estimation of each period's effect. Since it automates many aspects of m/MMM, Robyn is ideal for practitioners and marketers who need quick insights with minimal manual intervention. It also offers model calibration against experiment results via multi-objective optimization that is relatively easier to use compared to Meridian's and PyMC-Marketing-mmm's approach via priors. This approach is, however, also much less understood and vetted (a point we will return to later).

[4] In m/MMM, saturation describes how you model diminishing returns to additional spend in a channel and adstock is a transformation of media spend or other input variables like impressions, views, or clicks that addresses that (1) yesterday's advertising still has some effect today (carryover) and (2) the effect of past advertising gets smaller as time passes (decay).

PyMC-Marketing (PyMC Labs 2025) offers a fully Bayesian and highly customizable m/MMM solution with custom priors and hierarchical models.[5] Unlike Robyn, which is optimized for automation and ease of use, PyMC-Marketing-mmm is geared toward advanced practitioners who want custom-built MMM models to fit specific business needs. While Robyn and Meridian assume static effects for the estimation time period, PyMC-Marketing-mmm allows to specify a time-varying media baseline (PyMC Labs 2025a). Its advanced customizability makes it most suited for researchers and professionals with strong Bayesian modeling expertise.

At the budget allocation/optimization step (step 4 per Figure 1), Meridian offers the most comprehensive set of options, including: (i) working from a set budget to maximize ROI or another incremental KPI; (ii) working from a flexible budget, targeting a set minimum ROI; (iii) working from a flexible budget, targeting a set marginal ROI (Google 2025a). Robyn currently supports the first two options, and PyMC-Marketing-mmm requires customization to support budget allocation/optimization.

Table 2 compares the three packages on key usability considerations. Robyn is the most automated solution, allowing users without deeper statistics or marketing analytics knowledge to run an end-to-end m/MMM modeling chain within less than a day. A notable strength are the concise but comprehensive diagnostic one-pagers it outputs for model candidates. However, the risk of model misspecification from automation is relatively high. For practitioners needing fast and automated insights and who have more basic data science knowledge and coding skills, it is currently the best choice. Meridian provides the most "packaged up" Bayesian solution. Users need a robust

[5] Hierarchical (multi-level) models estimate the same effect for many groups (countries, regions, campaigns, cohorts, etc.) jointly instead of separately. Each group gets its own parameter, but those parameters are regularized toward a common distribution, so groups with little data "borrow strength" from groups with more data. That regularization is often termed partial or hierarchical pooling.

command of data science coding to adopt it successfully and can do so within several days based on the authors' experience. Some modeling choices such as default uniform priors facilitate broad adoption but introduce risk of misuse in the authors' opinion. For advanced users comfortable with Bayesian statistics, PyMC-Marketing-mmm provides the most flexibility. To leverage it for a full m/MMM modeling chain, users will need to customize it and be able to write custom code. Completion of a first full modeling run can take up to two weeks based on the authors' experience. To effectively do so, users need a good handle of Bayesian modeling, especially priors. Appendix B provides an adoption guide for advertisers who are new to m/MMM, helping them select the appropriate package based on their organizational contingencies and use cases.

Computational element	Robyn	Meridian	PyMC-Marketing-mmm
2a. Estimate			
Channel / media aggregations	Manual	Manual	Manual
Dynamic effects, saturation, seasonality	Supports geometric and Weibull adstock; models saturation using S-curves; incorporates seasonality via time-series decomposition	Implements adstock with geometric decay; saturation using Hill functions; and seasonality through time-varying intercepts	Provides customizable adstock and saturation functions; seasonality modeling is user-defined
Multi-collinearity	Regularization through Ridge regression	Utilizes Bayesian priors to mitigate collinearity	Manages via Bayesian priors and model structure
Reach & frequency data	Supports reach and frequency data	Integrates reach and frequency data	Requires manual integration
Calibration against experiment results	Enables calibration using lift test results	Incorporates experimental results through ROI priors	Supports inclusion of experimental data via priors
Calibration against business plausibility	Allows manual calibration and inclusion of business constraints	Integrates business knowledge through ROI priors	Facilitates incorporation of expert knowledge via priors
Long-term effects (native baseline models)	Models long-term effects using adstock and saturation curves	Accounts for long-term effects via time-varying intercepts	Requires explicit modeling of long-term effects
Dynamic coefficients	Does not support time-varying coefficients	Does not explicitly support time-varying coefficients	Allows time-varying media baseline (PyMC Labs 2025a)
Channel interaction effects	Limited support; requires manual feature engineering	Limited support; requires manual feature engineering	Limited support; requires manual base code adaptation
Hierarchical modeling	Does not natively support hierarchical structures	Implements hierarchical geo-level modeling	Fully supports hierarchical modeling structures
2b. Infer + 3. Select			
Concise diagnostic output	Provides model diagnostics and convergence plots in one-pagers	Offers comprehensive diagnostics and uncertainty quantification	Requires custom diagnostic plotting and analysis
Refresh / updating function	Supports model retraining for continuous updates	Enables continuous modeling with new data	Allows for updating through model retraining
4. Prescribe			
Budget allocation / optimization	Includes built-in budget allocator with multi-objective optimization	Offers budget optimization tools with scenario analysis	Requires custom implementation for budget optimization
Different scenarios	Facilitates scenario planning with built-in tools	Supports scenario analysis through model simulations	Allows flexible scenario modeling via user-defined simulations

Table 1: Comparison of the leading open-source m/MMM packages on key computational elements that underpin the analytics process in Figure 1 (see first column for a mapping). None of the packages currently support implementing the prescribed media/marketing mix policy, running experiments, or attribution modeling, so step 1 "Act" is missing from this table. Step 2a receives comparatively most attention as risk from automation is highest per Figure 1, so scholarly study and scrutiny of what the packages offer here is important.

	Robyn	Meridian	PyMC-Marketing-mmm
Ease of use	High; users need to know basic R or Python and consult the documentation to adopt; full modeling run possible within a day (after data preparation is completed)	Medium; users need good command of Python and a robust understanding of Bayesian modeling; full modeling run possible within several days	Comparatively low; users need good command of Python, a robust understanding of Bayesian modeling; full modeling run possible within one to two weeks
Degree of automation	High; after data aggregation, users can run an end-to-end modeling chain even without deeper understanding of regression and statistics	Medium; human effort needed repeatedly along the modeling chain; chosen defaults make more advanced modeling elements relatively easy to handle for inexperienced users	Comparatively low; users need to be able to write custom code to run a full modeling chain
Risk from automation	High; risk of model misspecification is relatively high; users need to carefully consult documentation to correctly read diagnostic one-pagers	High; default uniform priors make Bayesian modeling relatively straightforward but increase risk of improper use and model misspecification	Low; users without robust understanding of statistics and coding are unlikely to adopt successfully
Notable strengths	Concise but comprehensive diagnostic one-pagers; end-to-end automation with minimal human effort needed; the overall most “packaged up” m/MMM solution	Most “packaged up” Bayesian m/MMM solution; makes a Bayesian framework broadly accessible, risk of misuse is somewhat elevated as a result	Highly customizable and extendable; can provide a robust basis for advanced modeling solutions; part of the broader, large and active PyMC open-source community
Target users	Small and midsize advertisers with limited proprietary modeling experience and marketing analytics capabilities	Small and midsize advertisers with more advanced marketing analytics experience; advertisers with existing m/MMM experience but new to Bayesian modeling	Advertisers of all sizes with substantial marketing analytics capabilities and initial proprietary modeling experience

Table 2: Comparison of the leading open-source m/MMM packages on key usability considerations; appendix B provides an adoption guide, substantiating how target users were selected in this table.

3.2. The most “packaged up” solution: Robyn

Meta Platforms’ m/MMM package Robyn was the first package to be released and, as we assessed through use of all packages, offers the highest degree of automation and shortest, and easiest path to estimating a full model among the available packages. As of November 27, 2025, it is being developed by 34 contributors, starred on Github more than 1,400 times, forked 416 times (Meta

Platforms 2024),[6] and has seen more than 121,000 downloads (CRAN 2025a). The package is available in both R and Python and distributed under the MIT open-source license which permits private and commercial use free of charge.[7]

Robyn is designed as a "semi-automated" MMM solution – once the data is prepared, it automates model building end-to-end: transforming variables, running hyperparameter optimization for adstock and saturation, and outputting a set of Pareto-optimal models and diagnostics. It also automates common adjustments like detecting seasonality and provides built-in model selection support. This high level of automation means analysts do not need to manually tune lag or diminishing return parameters, but the package identifies those.

This level of automation is novel in m/MMM specifically and in causal inference more generally. Some modeling components in the package further are unvetted. Two of these computational features warrant specific discussion.[8] The first is that Robyn currently works merely from time series of outcome and media data, without consideration of additional cross-sectional dimensions such as geographic clusters (Aurier and Broz-Giroux 2014). While this approach facilitates data preparation, easy model estimation, and broad accessibility, it can limit customer and advertising insights. Meridian and PyMC-Marketing-mmm accept panel data with cross-sectional hierarchies such as geo clusters. The second component is the multi-objective optimization that is available in Robyn's model estimation step. Robyn allows users to set two objectives in addition to (statistical) minimization of NRMSE: avoidance of extreme results (producing more business-plausible

[6] PyMC-Marketing has 83 contributors, is starred 987 and forked 341 times, Meridian (with 13 contributors) 1200 and 209 times respectively, per their Github pages as of November 27, 2025. Note that PyMC-Marketing comprises other functionality beyond m/MMM, which complicates this comparison, especially when it comes to the number of contributors.
[7] The authors recommend using the R version of Robyn – it was more stable than the Python version when they used both in the first half of 2025.
[8] For a comprehensive introduction to Robyn's architecture, see Runge et al. (2024).

results) and minimization of mean absolute percentage error (MAPE) against estimates from experiments or other trusted sources such as established attribution models (Vidale and Wolfe 1957; Little and Lodish 1969; Berman 2018; Gordon et al. 2019). This approach is different from the other discussed packages that use Bayesian estimation and priors to calibrate models against estimates from experiments.

This multi-objective optimization essentially performs model selection by generating a Pareto frontier of models that trade off statistical fit (NRMSE), complexity (Decomp.RSSD), and calibration (MAPE.LIFT).[9] The package then recommends the models closest to and on the Pareto frontier to the user for review in the inference step (2.2) per Figure 1. The user might select the model that makes the most business sense, e.g., avoiding a model that unrealistically attributes zero effect to a major channel. While providing users with a model for decision-making, this approach does not output a single definitive model but a set of candidates, reflecting ambiguity in model specification. This approach in Robyn epitomizes a key challenge in "packaging up" mMM: finding the right trade-off between automation and easy accessibility on the one hand and customizability and rigor on the other – which we will discuss in greater detail now.

4. Challenges and Opportunities in Open-Source Marketing Mix Modeling

In this section, we comment on challenges and opportunities offered by the developments in the open-source m/MMM space. We thereby select challenge-opportunity pairs of particular relevance to marketing scholars and analysts to ensure the rigor and reliability of the developed methodologies. Table 3 provides an overview.

[9] See Runge et al. (2024) for details.

Challenge	Opportunity
Finding the right trade-off between automation and easy accessibility versus customizability and rigor	Provide robust customer insight and marketing strategy analytics in a privacy-first era; democratize access to m/MMM, especially for small and midsize advertisers
Integrating m/MMM with other marketing decision support tools such as attribution models and testing frameworks	Fostering best practices, e.g., in calibration and integration with a suite of measurement tools, through built-in workflows, detailed documentation and educational offerings
Vetting and testing novel approaches prior to broad adoption by the open-source community, e.g., Robyn's multi-objective optimization	Harnessing the power of open-source development, incorporating methodologies with cross-disciplinary origin in m/MMM
Ensuring rigorous use, adequately quantifying and communicating uncertainty	Facilitating good organizational behaviors and outcomes through package design

Table 3: Challenges and opportunities in open-source m/MMM; practice-academy collaboration can help materialize the opportunity paired to each challenge.

First, the trade-off between accessibility/automation and rigor/customizability has long permeated the world of marketing effectiveness (Lautman and Pauwels 2009). Automation can generate large time savings, but overly simplistic models can be imprecise or find insufficient buy-in with stakeholders. Different packages prioritize customer needs differently: Robyn is arguably the most automated solution, easiest to adopt, and hence democratizes access to m/MMM most broadly. On the other hand, this easy accessibility may lead to misuse and ill-informed decisions because users' understanding of the underlying methodologies is too limited. PyMC-Marketing offers the most advanced and flexible solutions but requires much more modeling proficiency from its users. Meridian finds a middle ground. In conjunction, the different packages serve a multitude of different customer needs along a spectrum of sophistication. The opportunity here – that is already realizing itself – is much broader access to and adoption of m/MMM across small and midsize advertisers, especially interactive digital-first ones.

Second, a key challenge for marketers is often the orchestration and integration of different measurement methods. Information from attribution models, m/MMM, experiments, and other sources needs to be integrated and aligned to optimally support marketing decision experiments (Vidale and Wolfe 1957; Little and Lodish 1969). A related challenge is how to integrate open-source m/MMM with these other decision support tools. The opportunity here is to foster best practices in calibration and integration with a suite of measurement tools through built-in workflows in the packages, detailed documentation, and educational offerings (Runge, Patter, and Skokan 2023; Runge et al. 2024).

Third, a challenge is to vet and test new computational approaches and pathways that some open-source communities have embarked on. E.g., Robyn applies optimization across three distinct objectives, balancing statistical error against the data, statistical error against trusted results from other sources, and business error (by formalizing a preference for non-extreme outcomes in the estimation step). This approach is novel and interesting as business plausibility is a criterion often used in model evaluation as much as statistical fit (Manchanda, Rossi, and Chintagunta 2004; Moorman and Day 2016). On the other hand, this approach is unvetted and a nontrivial degree of the resulting models may not approximate a reasonable data generating process but something more arbitrary. Thus, while Robyn's functionality provides flexibility to users in an otherwise highly automated package (a strength), it entails ambiguity in uncertainty specification and methods that have not been rigorously studied. A risk here is that Robyn evokes a sense of confidence in the results that is not actually backed by the data and deep statistical understanding on the side of the package's users. The opportunity however may be equally large: Marketers can harness the power of open-source development to incorporate new methodologies with cross-disciplinary origin in m/MMM.

Fourth, m/MMMs are modeling complex and dynamic market processes. Representing uncertainty appropriately and ensuring rigorous use of open-source m/MMM is challenging. Two key issues are often at play here: (1) Is there enough variation to estimate the associations between media spend and outcomes? (2) Can the estimates be interpreted as causal? As to the first issue: because the variation in marketing allocations over time can be insufficient to generate well-identified estimates (Manchanda, Rossi, and Chintagunta 2004; Arora et al. 2023), it is important to consider intentionally amplifying it. This can be achieved by regularly implementing well-thought-out tactical changes and recording these changes in the mMM data and taking action on mMM results, i.e., changing media spend allocations. A further concern is the strategic timing of advertising that tends to intensify around high-demand holidays. Parameters and data addressing systematic demand shifts should hence always be included in the estimation. Speaking to the second issue, the estimation of causally interpretable estimates is a critical concern in m/MMM. Observational studies of individual behavior are particularly prone to overestimating the causal impact of advertising due to "activity bias" and "user-induced endogeneity" (Lewis, Rao, and Reiley 2011). Customer-initiated ads (De Haan, Wiesel, and Pauwels 2016) are triggered by visiting a particular website (page), making it challenging to attribute the resulting purchase to the ad versus the consumer self-selection. In aggregate models, a similar bias arises from the strong correlation between many on- and/or offline activities, complicating the establishment of cause-effect relationships and potentially leading to systematic misestimation of advertising effects.

To mitigate both types of risks, we encourage validation of m/MMM results with experimental methods (Vidale and Wolfe 1957; Little and Lodish 1969; Hanssens and Pauwels 2016), such as field experiments (e.g., Wiesel, Pauwels, and Arts 2011; Valenti et al. 2024), advertising experiments (Eastlack and Rao 1989; Blake, Nosko, and Tadelis 2015; Lewis and Rao 2015; Gordon et al. 2019; Johnson 2023), or lower-level attribution results that the organization has found to reflect incremental effects (Feit al. 2013; Li and Kannan 2014; Berman 2018). For instance, last click-attributed outcomes may approximate the causal incremental effects of campaigns, particularly when baseline effects are minimal or absent, as might occur with new or unknown products and brands with little or no organic discovery. The opportunity here is to facilitate good organizational behaviors and outcomes through smart intentional package design (Runge et al. 2024). A close dialogue between marketing academics and open-source m/MMM communities can be a fruitful pathway to achieve a convergence between long-established and -documented best practices and open-source development.

5. Outlook

In the past, the technology industry has witnessed a series of transformative open-source "wars," where community-driven, transparent alternatives have challenged, and in some domains displaced, long-standing proprietary solutions. These clashes – ranging from operating systems (Linux vs. Windows), web browsers (Firefox vs. Internet Explorer/Chrome), to infrastructure tooling (Kubernetes vs. proprietary orchestrators) – have shaped the technological foundations of today's digital commerce, as many domains in computer science and engineering have seen open source triumph as a de facto standard.[10] In contrast, marketing analytics – particularly m/MMM – is often still the domain of closed, consultant-driven models. As a result, many companies continue to base marketing and media mix decisions on multitouch attribution (MTA), including last-click methods shown as inferior by several authors (e.g., Li and Kannan 2014; Blake, Nosko, and Tadelis 2015; De Haan, Wiesel, and Pauwels 2016; Berman 2018).

[10] Special thanks to William Grosso for a comment to this avail during friendly review.

By making state-of-the-art methods freely available, open-source initiatives lower barriers to adoption, enhance transparency, and enable reproducibility. They enable small and midsize advertisers to adopt models that have thus far often been beyond their reach and can accelerate the diffusion of innovations in modeling approaches. This democratization of m/MMM has the potential to improve how organizations learn from data, design more relevant marketing strategies, and respond more quickly to shifts in customer preferences. It is all the more important given increasing demands for privacy-enhancing technologies by platforms and regulators that make deterministic attribution approaches – that many digital-first advertisers relied on – increasingly obsolete (Johnson, Shriver, and Du 2020; Johnson, Runge, and Seufert 2022).

Note that this review is limited to publicly available documentation and supported features as of November 27, 2025. We focus on core modeling features relevant for typical in-market users. As these packages are under continuous development, new capabilities may emerge rapidly. At this point, it is hard to anticipate whether open-source m/MMM will become a lasting and integral part of applied marketing measurement. Each package is likely to find its specific use cases as shown in Table 2 and discussed in the text. E.g., Robyn may serve small and midsize businesses looking for ease of use and automation and to cross-validate and triangulate other modeling approaches. Meridian may serve somewhat more sophisticated advertisers with interest in Bayesian modeling. PyMC-Marketing and the wider PyMC umbrella package may be a backbone for specialized and tailored in-house models with custom extensions.

Marketing scholars can contribute to the open-source development efforts directly or collaborate with practitioners to help shape the future of applied marketing measurement and define rigorous and valid best practices generally, and for different verticals and organizational contingencies specifically. We believe that a few areas are worthy of particular attention from scholars:

- Given recent developments in large language models (LLMs) and LLM-based chatbots, LLMs trained on a corpus including open-source m/MMM code and documentation can be expected to be able to apply this knowledge and to estimate and output models as prompted by the user. We believe this to be an area of particular interest to scholars who could train related models and design m/MMM agents. This would also facilitate testing and vetting of the ongoing developments in open-source m/MMM via automation.
- Relatedly, methods from machine learning and artificial intelligence provide interesting angles for innovation in m/MMM (Proserpio et al. 2020; Ding et al. 2020), e.g., to include unstructured, text, image, audio, and video data in models directly (Mulc et al. 2025).
- Third, extended functionality to estimate dynamic coefficients, e.g., via Gaussian processes and with long-term decomposition (Dew, Ansari, and Li 2020; Dew, Padilla, and Shchetkina 2024; PyMC Labs 2025a), would be useful. Related models would be particularly useful to understand and model long-term trends and brand effects in a firm's advertising (Kireyev, Pauwels and Gupta 2016).
- Fourth, m/MMM in general needs a deeper integration with experimentation and other measurement sources (Vidale and Wolfe 1957; Little and Lodish 1969; Zantedeschi, Feit, and Bradlow 2017; Gordon, Moakler, and Zettelmeyer 2023), including setting up experiments for channels with high model uncertainty, to test diminishing returns and to calibrate model estimates. Open-source models, and related artificial agents per the first point in this list, can lead the way here. Reinforcement learning can extend the recommended Model-Experiment-Model-Experiment (MEME) iteration applied in practice prize papers (Wiesel, Pauwels, and Arts 2011; Hanssens and Pauwels 2016; Valenti et al. 2024).

- It would be useful if the packages provided tools for sensitivity analysis to evaluate how much unobserved confounding or violations of key assumptions would change recommendations. In the absence of theoretical guarantees, such tools can provide confidence when using observational data to guide decision making.
- Relatedly, more comprehensive and explicit validation, e.g., via stability checks and parameter recovery on simulated data, built into the packages could help provide more robust modeling results.
- Finally, flexible long-term effects, competitive interactions and dynamic synergy are absent from the currently available open-source packages but can yield important insight in dynamic marketing environments (Dekimpe and Hanssens 1999; Jedidi, Mela, and Gupta 1999; Zantedeschi, Feit, and Bradlow 2017; Pauwels 2018; Cain 2022).

We hope that the overview provided in this paper is a useful starting point for interested scholars and practitioners and can encourage collaboration, study, and contribution.

Funding

This research received no external funding.

Conflict of Interest

The authors declare no conflicts of interest.

References

Arora, N., Berman, R., Feit, E., Hanssens, D., Li, A., Lovett, M., Mela, C., Wilbur, K., & Lynch, J. (2023). Charting the Future of Marketing Mix Modeling Best Practices. *Marketing Science Institute* Blue Ribbon Panel Report. [Last accessed on February 10, 2025 at https://www.msi.org/presentation/msi-blue-ribbon-panel-report-charting-the-future-of-marketing-mix-modeling-best-practices/]

Aurier, P., & Broz-Giroux, A. (2014). Modeling advertising impact at campaign level: Empirical generalizations relative to long-term advertising profit contribution and its antecedents. *Marketing Letters*, *25*(2), 193-206.

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

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

Borden, N. H. (1964). The concept of the marketing mix. *Journal of Advertising Research*, 4(2), 2-7.

Cain, P. M. (2022). Modelling short-and long-term marketing effects in the consumer purchase journey. *International Journal of Research in Marketing*, *39*(1), 96-116.

CRAN (2025). CRAN page for Robyn. [Last accessed on November 27, 2025 at https://cran.rstudio.com/web/packages/Robyn/index.html]

CRAN (2025a). CRAN download count for Robyn. [Last accessed on November 27, 2025 at https://cranlogs.r-pkg.org/badges/grand-total/Robyn]

De Haan, E., Wiesel, T., & Pauwels, K. (2016). The effectiveness of different forms of online advertising for purchase conversion in a multiple-channel attribution framework. *International Journal of Research in Marketing*, *33*(3), 491-507.

Deighton, J. A., Mela, C. F., & Moorman, C. (2021). Marketing thinking and doing. *Journal of Marketing*, 85(1), 1-6.

Dekimpe, M. G., & Hanssens, D. M. (1995). Empirical generalizations about market evolution and stationarity. *Marketing science*, *14*(3_supplement), G109-G121.

Dekimpe, M. G., & Hanssens, D. M. (1999). Sustained spending and persistent response: A new look at long-term marketing profitability. *Journal of Marketing Research*, *36*(4), 397-412.

Dew, R., Ansari, A., & Li, Y. (2020). Modeling dynamic heterogeneity using Gaussian processes. *Journal of Marketing Research*, *57*(1), 55-77.

Dew, R., Padilla, N., & Shchetkina, A. (2024). Your MMM is Broken: Identification of Nonlinear and Time-varying Effects in Marketing Mix Models. *arXiv preprint arXiv:2408.07678*.

Ding, Y., DeSarbo, W. S., Hanssens, D. M., Jedidi, K., Lynch, J. G., & Lehmann, D. R. (2020). The past, present, and future of measurement and methods in marketing analysis. *Marketing Letters*, *31*, 175-186.

Eastlack Jr, J. O., & Rao, A. G. (1989). Advertising experiments at the Campbell soup company. *Marketing Science*, *8*(1), 57-71.

Feit, E. M., Wang, P., Bradlow, E. T., & Fader, P. S. (2013). Fusing aggregate and disaggregate data with an application to multiplatform media consumption. *Journal of Marketing Research*, *50*(3), 348-364.

Google (2023). Lightweight Bayesian marketing mix modeling. *Github repository*. [Last accessed on September 3, 2025 at https://github.com/google/lightweight_mmm]

Google (2025). Meridian. *Github repository*. [Last accessed on November 27, 2025 at https://github.com/google/meridian]

Google (2025a). Budget optimization scenarios. *Google Developers blog*. [Last accessed on February 5, 2025 at https://developers.google.com/meridian/docs/user-guide/budget-optimization-scenarios]

Gordon, B. R., Jerath, K., Katona, Z., Narayanan, S., Shin, J., & Wilbur, K. C. (2021). Inefficiencies in digital advertising markets. *Journal of Marketing*, *85*(1), 7-25.

Gordon, B. R., Moakler, R., & Zettelmeyer, F. (2023). Close enough? A large-scale exploration of non-experimental approaches to advertising measurement. *Marketing Science*, *42*(4), 768-793.

Gordon, B. R., Zettelmeyer, F., Bhargava, N., & Chapsky, D. (2019). A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook. *Marketing Science*, 38(2), 193-225.

Hanssens, D. M., & Pauwels, K. H. (2016). Demonstrating the value of marketing. *Journal of Marketing*, *80*(6), 173-190.

Hilt, D. E., & Seegrist, D. W. (1977). *Ridge, a computer program for calculating ridge regression estimates*. Department of Agriculture, Forest Service, Northeastern Forest Experiment Station.

Jagpal, H. S., & Brick, I. E. (1982). The marketing mix decision under uncertainty. *Marketing Science*, *1*(1), 79-92.

Jedidi, K., Mela, C. F., & Gupta, S. (1999). Managing advertising and promotion for long-run profitability. *Marketing Science*, *18*(1), 1-22.

Johnson, G. A. (2023). Inferno: A guide to field experiments in online display advertising. *Journal of Economics & Management Strategy*, *32*(3), 469-490.

Johnson, G., Runge, J., & Seufert, E. (2022). Privacy-centric digital advertising: implications for research. *Customer Needs and Solutions*, 9(1-2), 49-54.

Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry?. *Marketing Science*, *39*(1), 33-51.

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.

Lambin, J. J. (1972). A computer on-line marketing mix model. *Journal of Marketing Research*, *9*(2), 119-126.

Lautman, M., & Pauwels, K. (2009). What is important? Identifying metrics that matter. *Journal of Advertising Research*, *49*(3), 339-359.

Lee, S. Y., Runge, J., Yoo, D., Bart, Y., Gyurak, A., & Schneider, J. W. (2025). Social Media Advertising's Role in Small Business Resilience amid Lockdowns. *Journal of Interactive Marketing*, forthcoming.

Lewis, R. A., Rao, J. M., & Reiley, D. H. (2011). Here, there, and everywhere: correlated online behaviors can lead to overestimates of the effects of advertising. In *Proceedings of the 20th International Conference on World Wide Web* (pp. 157-166).

Lewis, R. A., & Rao, J. M. (2015). The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics*, *130*(4), 1941-1973.

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

Little, J. D., & Lodish, L. M. (1969). A media planning calculus. *Operations Research*, *17*(1), 1-35.

Little, J. D. (1975). BRANDAID: A marketing-mix model, part 1: Structure. *Operations Research*, *23*(4), 628-655.

Manchanda, P., Rossi, P. E., & Chintagunta, P. K. (2004). Response modeling with nonrandom marketing-mix variables. *Journal of Marketing Research*, 41(4), 467-478.

Mela, C. F., Gupta, S., & Lehmann, D. R. (1997). The long-term impact of promotion and advertising on consumer brand choice. *Journal of Marketing Research*, *34*(2), 248-261.

Meta Platforms (2024). Robyn. *Github*. [Last accessed on November 27, 2025 at https://facebookexperimental.github.io/Robyn/]

Moorman, C., & Day, G. S. (2016). Organizing for marketing excellence. *Journal of Marketing*, 80(6), 6-35.

Mulc, T., Anderson, M., Cubre, P., Zhang, H., Liu, I., & Kumar, S. (2025). NNN: Next-Generation Neural Networks for Marketing Mix Modeling. *arXiv preprint arXiv:2504.06212*.

Naik, P. A., & Raman, K. (2003). Understanding the impact of synergy in multimedia communications. *Journal of Marketing Research*, *40*(4), 375-388.

Nair, H. (2025). Meridian is now available to everyone. *Google Ads & Commerce Blog*. [Last accessed on February 4, 2025 at https://blog.google/products/ads-commerce/meridian-marketing-mix-model-open-to-everyone/]

Pauwels, K. H. (2018). Modeling dynamic relations among marketing and performance metrics. *Foundations and Trends® in Marketing*, *11*(4), 215-301.

Pauwels, K., Aksehirli, Z., & Lackman, A. (2016). Like the ad or the brand? Marketing stimulates different electronic word-of-mouth content to drive online and offline performance. *International Journal of Research in Marketing*, *33*(3), 639-655.

Pauwels, K., & Aksehirli, Z. (2025). Big data analytics democratized with clean collaboration and customer privacy choice. *Journal of Business Research*, *188*, 115112.

Proserpio, D., Hauser, J. R., Liu, X., Amano, T., Burnap, A., Guo, T., ... & Yoganarasimhan, H. (2020). Soul and machine (learning). *Marketing Letters*, *31*, 393-404.

PyMC Labs (2025). PyMC-Marketing. *Github repository*. [Last accessed on November 27, 2025 at https://github.com/pymc-labs/pymc-marketing]

PyMC Labs (2025a). MMM with time-varying media baseline. *PyMC-Marketing How-To*. [Last accessed on September 3, 2025 at https://www.pymc-marketing.io/en/0.7.0/notebooks/mmm/mmm_time_varying_media_example.html]

Runge, J., Geinitz, S., & Ejdemyr, S. (2020). Experimentation and performance in advertising: An observational survey of firm practices on Facebook. *Expert Systems with Applications*, 158, 113554.

Runge, J., Patter, H., & Skokan, I. (2023). A new gold standard for digital ad measurement. *Harvard Business Review*.

Runge, J., Skokan, I., Zhou, G., & Pauwels, K. (2024). Packaging Up Media Mix Modeling: An Introduction to Robyn's Open-Source Approach. *Marketing Science Institute Working Paper Series 2024*, Report No. 24-147. [Last accessed on November 27, 2025 at https://www.msi.org/working-paper/packaging-up-media-mix-modeling-an-introduction-to-robyns-open-source-approach/]

Srinivasan, S., Vanhuele, M., & Pauwels, K. (2010). Mind-set metrics in market response models: An integrative approach. *Journal of Marketing Research*, *47*(4), 672-684.

Valenti, A., Srinivasan, S., Yildirim, G., & Pauwels, K. (2024). Direct mail to prospects and email to current customers? Modeling and field-testing multichannel marketing. *Journal of the Academy of Marketing Science*, *52*(3), 815-834.

Vidale, M. L., & Wolfe, H. (1957). An operations-research study of sales response to advertising. *Operations Research*, *5*(3), 370-381.

Wiesel, T., Pauwels, K., & Arts, J. (2011). Practice prize paper—Marketing's profit impact: Quantifying online and off-line funnel progression. *Marketing Science*, *30*(4), 604-611.

Zantedeschi, D., Feit, E. M., & Bradlow, E. T. (2017). Measuring multichannel advertising response. *Management Science*, *63*(8), 2706-2728.

Appendix A: Robyn Workflow and Code Example

This appendix demonstrates how the analytics process shown in Figure 1 can map onto the workflow of an m/MMM open-source package. It adapts the workflow and code examples from the package Robyn as presented in Appendix A of Runge et al. (2024). Runge et al. (2024) also provide a detailed walk-through of the package and list six brief case studies from companies that successfully adopted Robyn for their measurement.

The workflow provides a simplified overview of the steps involved in a complete modeling run of Robyn, on the simulated data available within the Robyn R package "dt_simulated_weekly" (Meta Platforms 2024). The workflow comprises input and output, from data loading to model estimation, budget allocation, and model refresh. In the version here, we added starred lines indicating how the workflow steps map to the steps of the m/MMM modeling chain in Figure 1.

Robyn's documentation provides additional and deeper walk-throughs and examples, e.g., at https://github.com/facebookexperimental/Robyn/blob/main/demo/demo.R (Meta Platforms 2024) and at https://drive.google.com/file/d/1LYM9g2xP5ZGhmGXlcQf--3EyXUaME6g6/view?ts=66b0ab09 for a demo video. Meridian and PyMC-Marketing offer similarly detailed documentation and guidance at https://github.com/google/meridian/blob/main/demo/Meridian_Getting_Started.ipynb and https://github.com/pymc-labs/pymc-marketing respectively.

WORKFLOW AND CODE EXAMPLE FOR ROBYN 3.11.1

****************** **MODELING STEPS *ESTIMATE* AND *INFER* PER FIGURE 1** *******************

STEP 0: SETUP ENVIRONMENT

- **Input/Steps**: Install Robyn and load into environment

- **Output**: Robyn R package is downloaded installed. This step ensures the latest stable (CRAN) or latest experimental version of Robyn is used for compatibility and new features.
- **Example Command**:

```
install.packages("Robyn")  #Install the latest stable version from CRAN
# remotes::install_github("facebookexperimental/Robyn/R") #Install the latest dev version from GitHub
library(Robyn)
```

STEP 1: LOAD DATA

- **Input**: Load in own or load the demo dataset `dt_simulated_weekly` & demo holiday `dt_prophet_holidays`
- **Output**: Data loaded into the R environment. In this step, user loads into Robyn all the necessary data for model building

STEP 2: MODEL SPECIFICATION

- **Input**: Dataset, holiday data, and model configuration parameters
- **Output**: `InputCollect` object containing all model specifications. This sets up the modeling system with all necessary variables and configurations for Robyn MMM analysis
- **Example Command:**

```
InputCollect <- robyn_inputs(
  dt_input = dt_simulated_weekly,
  dt_holidays = dt_prophet_holidays,
  date_var = "DATE", # date format must be "2020-01-01"
  dep_var = "revenue", # there should be only one dependent variable
  dep_var_type = "revenue", # "revenue" (ROI) or "conversion" (CPA)
  prophet_vars = c("trend", "season", "holiday"), # "trend","season", "weekday" & "holiday"
  prophet_country = "DE", # input country code. Check: dt_prophet_holidays
  context_vars = c("competitor_sales_B", "events"), # e.g. competitors, discount, unemployment etc
  paid_media_spends = c("tv_S", "ooh_S", "print_S", "facebook_S", "search_S"), # mandatory input
  paid_media_vars = c("tv_S", "ooh_S", "print_S", "facebook_I", "search_clicks_P"), # mandatory.
  # paid_media_vars must have same order as paid_media_spends. Use media exposure metrics like
  # impressions, GRP etc. If not applicable, use spend instead.
  organic_vars = "newsletter", # marketing activity without media spend
  # factor_vars = c("events"), # force variables in context_vars or organic_vars to be categorical
  window_start = "2016-01-01",
  window_end = "2018-12-31",
  adstock = "geometric" # geometric, weibull_cdf or weibull_pdf.
)
```

STEP 3: DEFINE AND ADD HYPERPARAMETERS

- **Input**: Definitions of model behavior parameters
- **Output**: Configured hyperparameters added to `InputCollect.` Hyperparameters like adstock and saturation define how marketing inputs decay over time and how they saturate, impacting model accuracy.
- **Example Values**: `facebook_S_thetas = c(0, 0.3)`: Range for the decay rate of Facebook spend, allowing for quick to moderate decay.
- **Note on regularization via ridge regression**: Lambda is the penalty term for regularised regression. Lambda doesn't need manual definition from the users, because it is set to the range of c(0, 1) by default in hyperparameters and will be scaled to the proper altitude with lambda_max and lambda_min_ratio.

STEP 4: MODEL CALIBRATION / ADD EXPERIMENTAL INPUT

- **Input**: Calibration data to align model outputs with other trusted outcomes (from experiments or attribution models)

- **Output**: Adjusted model parameters to reflect real-world effects, Calibration helps in adjusting the model based on ground truth data, enhancing predictive accuracy
- **Example Command**:

```
calibration_input <- data.frame(
channel name must in paid_media_vars
channel = c("facebook_S",  "tv_S", "facebook_S+search_S", "newsletter"),
liftStartDate = as.Date(c("2018-05-01", "2018-04-03", "2018-07-01", "2017-12-01")),
liftEndDate = as.Date(c("2018-06-10", "2018-06-03", "2018-07-20", "2017-12-31")),
liftAbs = c(400000, 300000, 700000, 200),
spend = c(421000, 7100, 350000, 0),
confidence = c(0.85, 0.8, 0.99, 0.95),
metric = c("revenue", "revenue", "revenue", "revenue"),
calibration_scope = c("immediate", "immediate", "immediate", "immediate")
)
InputCollect <- robyn_inputs(InputCollect = InputCollect, calibration_input = calibration_input)
```

STEP 5: BUILD INITIAL MODEL

- **Input**: InputCollect object
- **Output**: OutputModels object containing model results and diagnostics.

```
OutputModels <- robyn_run(
  InputCollect = InputCollect, # feed in all model specification
  cores = NULL, # NULL defaults to (max available - 1)
  iterations = 2000, # 2000 recommended for the dummy dataset
  trials = 5, # 5 recommended for the sample dataset
  ts validation = TRUE, # 3-way-split time series for NRMSE validation.
  add penalty factor = FALSE # Experimental feature. Use with caution.
  optimize_weights = c(1,1,1) # Set weights for multi-objective optimization across NRMSE,
Decomp.RSSD and MAPE.LIFT; defaults to c(1,1,1). New feature. Use with caution.
)
```

```
OutputCollect <- robyn_outputs(
  InputCollect, OutputModels,
  pareto fronts = "auto", # automatically pick how many pareto-fronts to fill min candidates (100)
  # min candidates = 100, # top pareto models for clustering. Default to 100
  # calibration constraint = 0.1, # range c(0.01, 0.1) & default at 0.1
  csv out = "pareto", # "pareto", "all", or NULL (for none)
  clusters = TRUE, # Set to TRUE to cluster similar models by ROAS. See ?robyn clusters
  export = create files, # this will create files locally
  plot folder = robyn directory, # path for plots exports and files creation
  plot pareto = create files # Set to FALSE to deactivate plotting and saving model one-pagers
)
```

********************* **MODELING STEP *SELECT* PER FIGURE 1** *********************

STEP 6: SELECT AND SAVE THE MODEL

- **Input**: OutputModels object

- **Output**: Selected model saved and export. This steps enables the user to choose the model that best fits their business context and save it for future use.
- **Example Command**:

```
select_model <- "1_122_7" # Pick one of the models from OutputCollect to proceed
ExportedModel <- robyn_write(InputCollect, OutputCollect, select_model, export = create_files)
print(ExportedModel)
```

********************* **MODELING STEP *PRESCRIBE* PER FIGURE 1** *********************

STEP 7: GET BUDGET ALLOCATION

- **Input**: Selected model and budget scenarios
- **Output**: Budget allocation recommendations. The budget allocator provides insights on how to allocate marketing budget effectively based on model predictions.
- **Example Command**:

`scenario = "max_response":` Optimizes for maximum return on investment.

```
AllocatorCollect1 <- robyn_allocator(
  InputCollect = InputCollect,
  OutputCollect = OutputCollect,
  select_model = select_model,
  # date_range = "all", # Default to "all"
  # total_budget = NULL, # When NULL, default is total spend in date_range
  channel_constr_low = 0.7,
  channel_constr_up = c(1.2, 1.5, 1.5, 1.5, 1.5),
  # channel_constr_multiplier = 3,
  scenario = "max_response",
  export = create_files
)
```

********************* **FURTHER RELEVANT COMPONENTS** *********************

STEP 8: MODEL REFRESH

- **Input**: JSON file of a previously exported model, new data
- **Output**: Updated model based on new data. Regular updates to the model ensure it remains accurate as new data becomes available.
- **Example Command**:

```
json_file <- "~/Desktop/Robyn_202211211853_init/RobynModel-1_100_6.json"
RobynRefresh <- robyn_refresh(
  json_file = json_file,
  dt_input = dt_simulated_weekly,
  dt_holidays = dt_prophet_holidays,
  refresh_steps = 13,
  refresh_iters = 1000, # 1k is an estimation
  refresh_trials = 1
```

STEP 9: GET MARGINAL RETURNS

- **Input**: Selected model, specific spend levels
- **Output**: Marginal returns for additional spend. Helps in understanding the incremental benefits of additional marketing spend, aiding in decision-making.
- **Example Command:**

```
Response <- robyn_response(
  InputCollect = InputCollect,
```

```
  OutputCollect = OutputCollect,
  select model = select model,
  metric name = "facebook S"
)
Response$plot
```

OPTIONAL: RECREATE PREVIOUS MODELS AND REPLICATE RESULTS

- **Input**: JSON file of a previously exported model
- **Output**: Recreated model and outputs. Ensures consistency and reproducibility by allowing users to reload and rerun models from saved states.

HYPERPARAMETERS EXPLANATION:

- **Geometric Adstock**: Theta values typically range from 0 to 0.8, reflecting the percentage of impact carried over to the next period. For example, a theta of 0.3 for TV means that 30% of the advertising effect remains the next day.
- **Weibull CDF Adstock**: Shape and scale parameters allow for a flexible decay rate. Shape influences the curve's form (S-shaped or L-shaped), while scale adjusts the inflection point of decay. Common bounds are shape (0, 2) and scale (0, 0.1).
- **Weibull PDF Adstock**: Similar to CDF but allows for lagged effects, with shape influencing the peak timing and scale adjusting the peak's position. Recommended bounds are shape (0.0001, 10) and scale (0, 0.1) for varied effects.
- **Hill Saturation Function**: Alpha and gamma parameters define the saturation curve's shape and inflection point. Typical bounds are alpha (0.5, 3) for curve shape and gamma (0.3, 1) for the saturation point.

Appendix B: Choosing the Right m/MMM Package

This appendix provides a decision aid for practitioners looking to adopt open-source mMM. We thereby focus on mMM (so, media mix modeling) as digital advertising measurement is currently a main driver of adoption among small and midsize businesses. In principle, all packages can be extended into MMM (so, marketing mix modeling) by including data on relevant marketing actions. As PyMC-Marketing-mmm offers the most flexibility and customizability, it also provides the most comprehensive basis for extension into MMM. In reality, mMM and MMM will often coexist and complement each other, e.g., with mMM informing more granular and tactical advertising decisions and laddering up into broader and more strategically oriented MMM.

Here is a step-by-step decision guide for an advertiser looking to adopt mMM for the first time via one of the packages. mMM should thereby complement, and not replace, existing and other decision support solutions like attribution models and advertising experiments. mMM can provide more strategic guidance at lower decision cadence and provide a longer term "source of truth" that incorporates and is calibrated against the other methods (Runge, Patter, and Skokan 2023).

1) Have you been advertising on more than two channels for more than a year (better: two years)?	Yes: continue to next row	No: use attribution models and ad experiments to inform your advertising operations; reconsider m/MMM when you have reached sufficient scale and data availability
2) Do you have reliable advertising spend and sales data per day or week for at least a year and at least two channels?	Yes: continue to next row	No: work to gather and prepare data first
3) Do you have in-house marketing science capabilities and proprietary modeling experience?	Yes: continue to next row	No: use Robyn
4) Are you familiar with Bayesian modeling?	Yes: continue to next row	No: use Meridian and/or Robyn
5) Do you have advanced modeling and Bayesian experience?	Yes: continue to next row	No: use Meridian

6) Do you have custom use cases and the ability to write custom code (e.g., for budget optimization and allocation)?	Yes: use PyMC-Marketing-mmm and continue to next row	No: use Meridian and continue to next row
7) Are you looking to inform marketing decisions and strategy beyond advertising?	Yes: prepare data on relevant marketing actions and include in modeling	No: revisit if this changes