

# ARTEFACT

A modern glass skyscraper at night, illuminated with vibrant blue and purple lights. The building's facade is covered in digital data visualizations, including line graphs and bar charts, suggesting a focus on technology and finance. The sky is dark, and the overall atmosphere is futuristic and high-tech.

How can Causal AI unlock growth  
for Financial Services CMOs?

# ARTEFACT

AI IS ABOUT PEOPLE

WE ACCELERATE DATA AND AI ADOPTION  
TO POSITIVELY IMPACT  
PEOPLE AND ORGANIZATIONS.



**27**  
COUNTRIES

**2500**  
EMPLOYEES

**+1000**  
CLIENTS

Artefact is a global leader in consulting services, specialized in data transformation and data & digital marketing, from strategy to the deployment of AI solutions. We are offering a unique combination of innovation (Art) and data science (Fact).



DATA / AI STRATEGY & TRANSFORMATION | DATA / AI FOR CUSTOMER | DATA / AI FOR OPERATIONS | DATA / AI FOR SUPPORT FUNCTIONS | IT & AI TECHNOLOGIES

# What if Financial Services could measure Trust and create Value through Causal AI?

---

## Executive Summary

- 01 Marketing Measurement is a Strategic Imperative for FSI**  
A. Financial Marketing Landscape is changing  
B. The Structural Complexity of FSI Marketing  
C. Industry Outlook 06
- 02 Causal Methodologies for Measurement in FSI**  
A. Observational Methodologies  
B. Quasi-Experimental Methodologies  
C. Multi-Touch Attribution (MTA) 13
- 03 Marketing Maturity and Organisational Success**  
A. The ROI of Maturity  
B. Measurement Without Adoption Is Worthless  
C. Core Recommendations for Institutional Implementation 31

## Conclusion: The Path to Marketing Value Optimisation

## References



# Executive Summary

Most Financial Service Institutions (FSI) players still cannot answer a question that any CFO will eventually ask: what did all those investments actually return? The combination of long sales cycles, brand-driven baselines, macroeconomic confounding, and some of the stringent data privacy and fairness regulations in any industry makes marketing measurement in FSI structurally harder than in almost any other sector. The result is a measurement gap that misallocates budget, undervalues mid-funnel and brand investments, and leaves marketing perpetually on the defensive when facing finance.

The authors of this paper argue that closing that gap requires a shift from correlation-based reporting to causal inference from dashboards that describe what happened to designs that prove what marketing effectively caused to happen. This is called the pursuit of a single causal truth: a defensible, evidence-based link between marketing investment and the outcomes that matter to FSIs Net Cash Flow, Net Interest Margin, AUM growth, and Customer Lifetime Value.

Drawing on Artefact's work with FSIs across payments, retail banking, asset management, and retail insurance, the authors present three interconnected measurement methodologies: Observational Models (Bayesian MMM & Structured Causal Models), Quasi-Experimental designs, and Multi-Touch Attribution and the unified Marketing Return of Investment (mROI) framework that integrates them. The authors also examine the organisational maturity required to deploy them and offer practical recommendations for institutions at every stage.

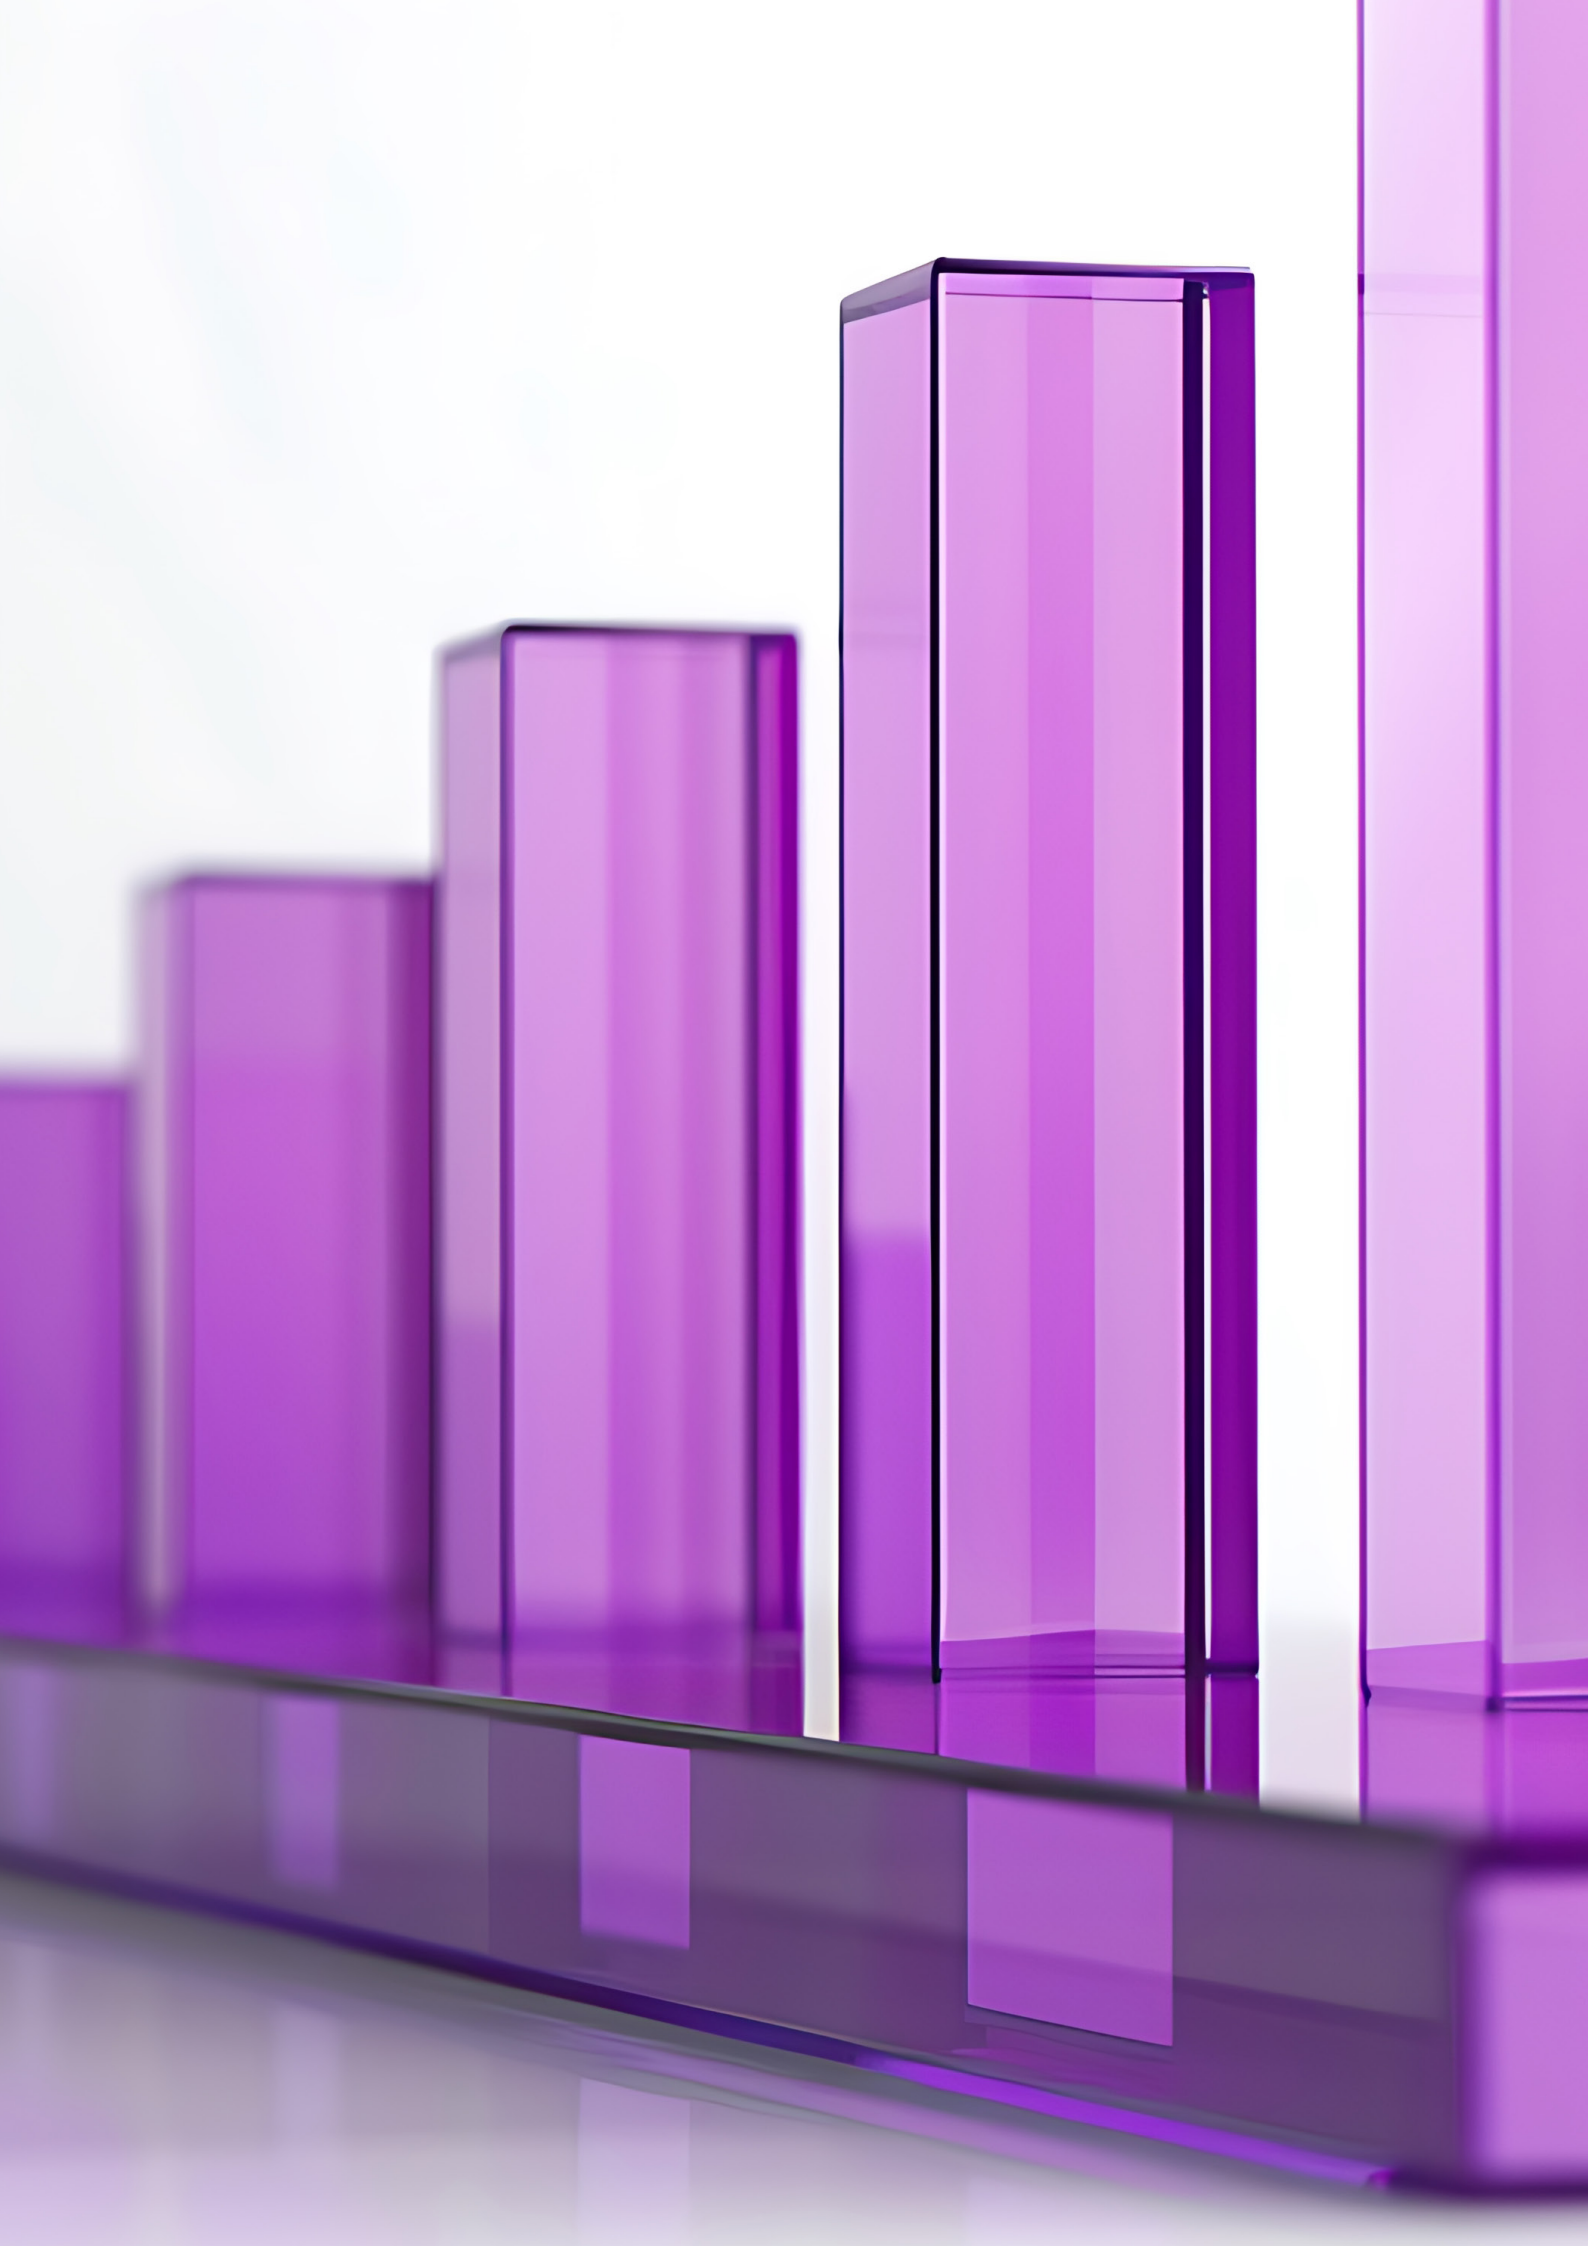
The firms that master causal measurement will not just defend their marketing budgets more effectively. They will allocate them better and build a durable competitive advantage in the process.



Financial Services institutions spend billions marketing trust they cannot measure. Advances in causal inference have fundamentally changed what is now provable and the institutions that recognise this will allocate capital more effectively than those that do not.”

Sid Mohan, Senior Director Data Science & AI  
ARTEFACT





# 01. Marketing Measurement is a Strategic Imperative for FSI.

## A. Financial Marketing Landscape is changing

Financial Marketing is being transformed from being a mere brand stewardship to a core value-generating competency. Traditionally, there has been a strong reliance on referral-based growth, proven track records and closed circles among asset and wealth managers. In today's world, influenced by intense competition and a shift in investor psychology as well as demographics, there needs to be a significant update in the role of Marketing. Marketing is no longer a peripheral support function; it is an «expanded mission» where data quality and digital agility are non-negotiable.

There is a clear pressure to demonstrate marketing ROI across financial services. According to Nielsen's 2024 Annual Marketing Report that whilst 84% of marketers in financial services express confidence in their measurement capabilities, only 38% evaluate the holistic value of marketing across both offline and digital channels – a gap that is particularly pronounced in an industry where measurement complexity is highest [1]. For instance, in asset management specifically, research conducted by Broadridge, surveying 20 senior marketing executives representing FSI institutions with AUM ranging from \$15 billion to \$2.4 trillion, showed that the pressure to demonstrate Marketing ROI has reached a critical threshold [2]. These leaders identify three primary structural challenges in driving acquisition and revenue generation:

01

**ROI Tracking and Proof:** Measuring and proving marketing's Return on Investment remains significantly more complex than tracking direct sales outcomes.

Financial marketers must upgrade their measurement capabilities to assess how capital is allocated and how value is proven and optimised in a diverse and sceptical market.

02

**Resource Constraints:** Limitations in marketing budgets and staffing levels continue to hinder the execution of sophisticated digital strategies at scale.

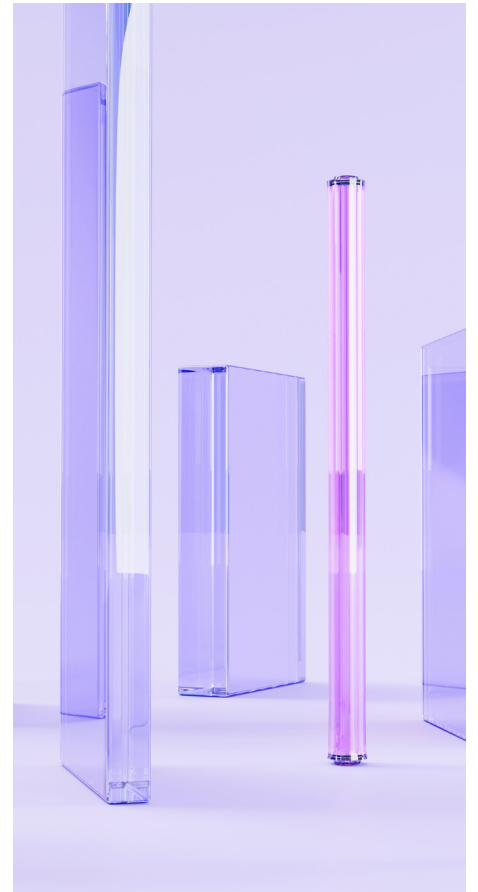
For the modern financial marketer, the challenge is to move beyond the surface-level metrics of clicks and conversions toward a «single causal truth» that aligns marketing investment with business-critical outcomes like Net Interest Margin (NIM), Assets Under Management (AUM) growth, and Customer Lifetime Value (CLV).

03

**Technology Gaps:** Existing infrastructure often fails to integrate the disparate data sets required to manage modern, multi-channel customer journeys effectively.

As Matt Schiffman, Principal of Distribution Insight at Broadridge Financial Solutions, emphasised: «The digital delivery world will not wait.» Several financial service institutions are already pivoting, as evidenced by the 2024 digital ad spend in the US alone, which reached above \$30 billion, an 18.2% increase from the previous year [2].

The industry is at a point where traditional methods of quantifying marketing effectiveness are inadequate to drive value in a landscape characterised by privacy-first regulations, brand value, interest-rate volatility, demographic changes, and “democratisation” of complex investment options.



## B. The Structural Complexity of FSI Marketing

### a. Role of Marketing in FSI

There is significant discussion in retail Financial Services on the role of marketing whether it is to simply acquire new customers, or if it also drives sustainable, longer-term growth. Because achieving «primacy» in an account holder's financial life is critical, expanding the number of products and services a customer uses is now a primary indicator of loyalty and institutional growth. Consequently, evaluating marketing performance based solely on short-term conversions or click-through rates fails to capture the true organisational impact of marketing. To align marketing efforts with enterprise-wide KPIs like Net Cash Flow, institutions are prioritising metrics such as Customer Lifetime Value (CLV), engagement quality, and retention and churn rates. To successfully link marketing dollars to these broader goals, marketers must map the immediate actions customers take directly to long-term relational outcomes.

The complexity of measuring this organisational impact lies in determining whether marketing actually caused the financial inflows or merely took credit for them. This requires a shift away from traditional attribution models which only show what a customer «touched» before converting toward measuring true incrementality.



### b. Importance of Brand

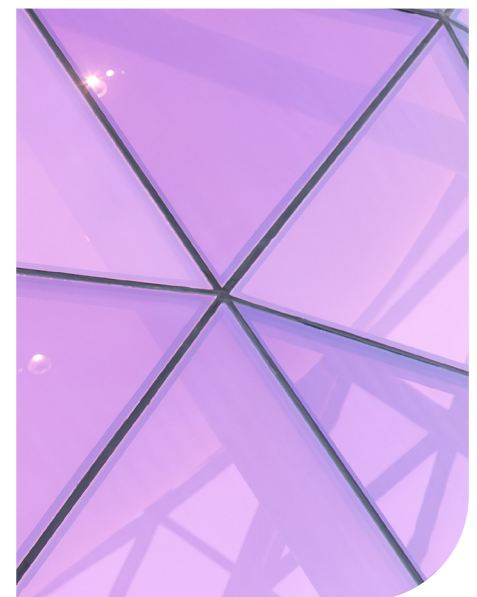
FSI works on trust. In financial services, where trust and reputation are paramount, an institution's brand credibility and thought leadership are foundational causal drivers for converting interest into long-term client relationships and capital allocation. Because a consumer's path through their financial life is increasingly unpredictable, maintaining a strong, visible brand image is essential for institutions to establish long-term relationships, loyalty, and trust.

Given this, any and all marketing investments have an impact on brand metrics that will translate into long-term evolution of the cash flow baseline. Incorporation of brand metrics is critical for any value measurement on the total value of marketing for FSI.



### c. External Market Dynamics

FSI is specifically prone to volatile market conditions and macroeconomic changes. Macroeconomic variables such as interest rates, inflation, and foreign exchange (FX) rates are critical components of financial services marketing measurement. Economic headwinds require a cohesive marketing strategy to ensure retention and reduced outflows, while tailwinds need to be capitalised upon. Failing to control for these macroeconomic shifts means marketers risk misattributing their impact whether tailwinds or headwinds to media campaigns. In retail banking, the macroeconomic environment dictates an institution's «interest income effect» and «balance sheet effect». While rising interest rates might increase immediate earnings, they simultaneously increase the «opportunity cost» for depositors, which can trigger the exact outflows and liquidity threats. Empirical analysis indicates that a 100-basis-point (bps) increase in short-term interest rates increases the Net Interest Margin (NIM) by approximately 13 bps for an average U.S. bank. However, this same rate increase can reduce the Economic Value of Equity (EVE) a measure of a bank's net worth by 8% to 18% (Abdymomunov, Gerlach & Sakurai, Federal Reserve Bank of Boston, 2023) [5].



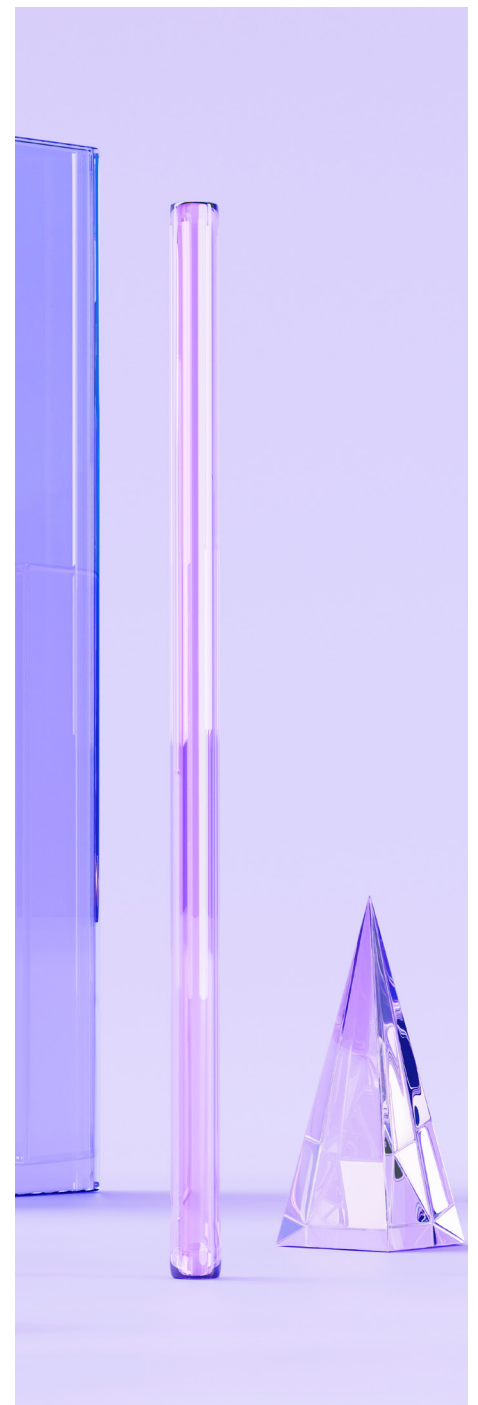
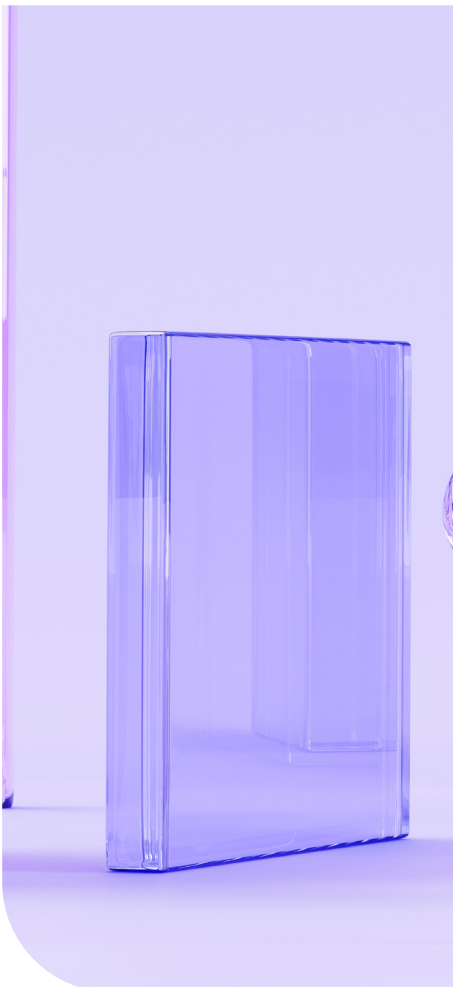
## d. Demographic Change

FSI institutions that fail to engage consumers under 45 risk ceding a once-in-generation opportunity. The US is going to witness a record wealth transfer of \$84 trillion to \$124 trillion from the Baby Boomers and Silent Generation to the Millennials and GenZ by 2048 [7]. The transfer is estimated to be over four times the size of the entire U.S. economy, (Cerulli Associates, 2024). To equip themselves for this demographic shift toward Gen Z and Gen Y (Millennials), FSI marketing teams are fundamentally changing how they engage, measure, and target consumers. A good example here is set by neobanks, which have completely embraced an elevated digital experience. The results speak for themselves: neobanks now account for close to 40% of new account openings globally, and two in five consumers expect a neobank to become their primary bank within three years – a structural shift led by digital-first players such as Revolut, Bunq, and N26 [8].

**Shifting to Algorithmic Advice and Digital Engagement:** Financial institutions recognise that younger audiences do not interact with banks the same way previous generations did. Firms are leveraging full-stack digital technologies and AI to deliver algorithmic advice and experiences that specifically appeal to younger investors. Because younger members primarily interact with digital products rather than visiting physical branches, traditional feedback loops are breaking down; for instance, younger users are far less likely to participate in surveys or leave comments than previous generations.

**Meeting Gen Z and Millennials on Social Media and Video:** To successfully market to individuals under 45, FSIs are abandoning broad traditional media in favour of highly targeted digital channels. Research shows that Gen Y and Gen Z are significantly more likely to prefer social media platforms like YouTube and TikTok for financial guidance and marketing compared to older generations. Additionally, data confirms that younger audiences (aged 18-30) have a much stronger positive response rate to Social Media and Pay-Per-Click (PPC) ads, whereas older demographics (46 and older) still prefer traditional email.

The format of communication is also changing. Video conferencing has become a top preferred channel for high-net-worth Gen Y investors to conduct portfolio reviews and receive alerts about investment opportunities. Recognising that over five million high-net-worth investors now use social media to make financial decisions, modern banks are increasingly replacing uncoordinated paper advertisements with dynamic social media strategies.



## C. Industry Outlook

Each FSI vertical carries distinct measurement challenges. What follows examines the specific dynamics – behavioural, regulatory, and structural – that determine what a rigorous measurement framework must account for in each segment.

### a. Payments Industry: Neural Mechanisms and Behavioural Economics

As described before, the payments industry is characterised by high-frequency data, where one needs to unpack the deep psychological and structural drivers of consumer behaviour. The core issue in payment marketing lies in separating the influence of a marketing promotion from the underlying “payment method effect”.

#### Neural Activation on Spending

According to Prof. Drazen Prelec and University of Utah Prof. Sachin Banker, coauthors of “Neural mechanisms of credit card spending,” consumers spend more when using credit cards than cash, based on an assumption that credit cards reduce the pain of payment and increase “motivation to spend”. Research utilising fMRI technology, which looks at brain activity during purchase, showed that consumers actively «step on the gas» by sensitising the dopaminergic reward networks in the brain, specifically the striatum [4]. This neural sensitisation drives greater purchasing by creating an anticipation of pleasure, making consumers more likely to make impulse buys and give larger tips. For payment providers, the causal question is what is the true impact of a marketing campaign for a rewards program?

A Double Machine Learning (DML) study found that higher availability of credit softly shifted consumers towards higher-value, discretionary luxury purchases instead of uniformly increasing spending across necessities [9].

#### Two-sided market

The market for payment cards is inherently two-sided, characterised by network externalities where consumers benefit from merchant acceptance, and merchants benefit from consumer adoption. Structural models have shown that merchants are significantly more sensitive to increases in usage costs than consumers. When credit card fees rise, merchants quickly reduce acceptance, which causally decreases the utility of the card for the consumer, even if the consumer’s own costs remain unchanged. Conversely, merchants who accept all means of payment can generate additional revenues of up to 3% of total sales [10].



## b. Retail Banking: Interest Rate Sensitivity and Compliance by Design

### Susceptible to Interest rate fluctuations

implies the need for an inextricable link to the macroeconomic environment to effectively measure the value of marketing. The profitability of a retail bank depends on the management of its loan portfolio and overall asset composition, which determines exposure to credit and interest rate risks. There are two primary mechanisms through which interest rates impact the banking sector: the «interest income effect» and the «balance sheet effect» [5]. The former implies a change in a customer’s purchasing power and spending behaviour due to a change in interest rates, while the latter refers to a change in household net worth (assets - liabilities), which influences spending behaviour [5]. This means that while rising interest rates increase a bank’s Net Interest Margin or earnings, they can also lead to a decrease in the Economic Value of Equity or net worth. The 2023 deposit migration crisis illustrated why separating marketing signal from macro noise matters.

As the Fed’s rate hiking cycle peaked, deposit outflows accelerated across incumbent retail banks – driven primarily by interest rate differentials, not marketing failures. Marketing drives value in this space when it can impact the «interest income effect» and the «balance sheet effect».

### The Macro-Causal Framework

In a world where interest rate increases can increase the Net Interest Margin (NIM) but also decrease Economic Value of Equity, optimal allocation of marketing investments to retain deposits, effectively identifying the «deposit beta» or price sensitivity of different customer segments, is essential. Any measurement methodology framework needs to factor this in.

### Regulatory Constraints

The banking sector operates under a higher level of scrutiny than traditional retail, where marketing models must not only be effective but also fair and explainable. The use of precision targeting and machine learning can unintentionally create «disparate impact,» where models disadvantage certain groups even if the intent is neutral. Variables such as credit scores, income, and homeownership often correlate with protected classes, making it difficult to be confident in the legality of model outputs without causal oversight.

To address this, leading institutions are adopting «compliance by design,» embedding regulatory guardrails into the data architecture and metadata. This involves integrating compliance partners at the campaign scoping stage and extending Model Risk Management (MRM) policies—traditionally used for credit risk—to marketing models. [11]

Régulatory Framework	Marketing Impact	Compliance Requirement
SEC Marketing Rule (206(4)-1)	Governs promotion of investment services; prohibits cherry-picking results.	Hypothetical performance data must include clear disclosures and context.
FINRA Rule 2210	Mandates that retail communications be fair, balanced, and not misleading	All performance claims must include prominent disclosures about risk.
GDPR / CCPA	Restricts use of personal data; grants “right to be forgotten”.	Explicit consent required for data collection; must delete data upon request
Fair Lending Standards	Prohibits discrimination in financial products.	Marketing models must be monitored for “disparate impact” on protected groups.

## c. Asset Management: Sophistications in Measurement

Asset management represents the most sophisticated application of causal measurement and inference in financial services, characterised by long-duration sales cycles, heavy reliance on intermediaries (financial advisors), and the compounding value of brand authority. At the end of 2024, U.S. professionally managed assets reached USD 73.7 trillion, with retail channels (USD 36.6 trillion) nearly equal to institutional channels (USD 37.1 trillion) [6]. As fee compression continues to squeeze margins on passive investments—expected to shrink by 15 basis points by 2030—asset managers must pivot from product distribution to delivering «differentiated alpha» [12][13]. This shift demands a measurement framework that can prove the causal link between marketing efforts and AUM growth.

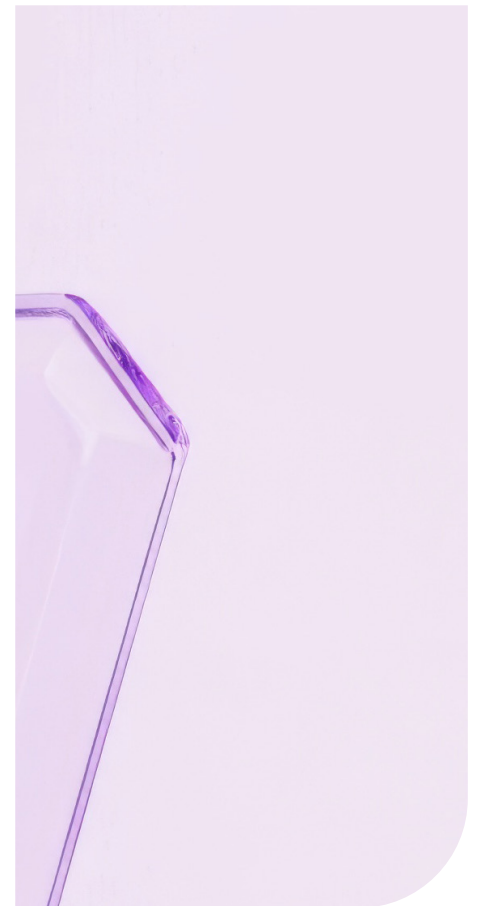
### The Financial Advisor Dynamic

The relationship between asset managers and financial advisors is the primary engine of growth, yet it is often the most difficult to measure. Today, advisors are increasingly engaging with asset managers through digital channels, even though external wholesalers are still central to closing sales.

An external wholesaler is essentially a territory-based sales representative employed by an asset management firm (such as a mutual fund or ETF provider). Their primary job is to build relationships with financial intermediaries like the retail financial advisors and wealth managers to educate them on the firm's specific investment products and persuade them to allocate their clients' capital into those funds. While «internal wholesalers» generally operate from a desk to provide data, scheduling, and inside sales support, external wholesalers have traditionally been the «boots on the ground.»

They travel their territory to conduct face-to-face meetings, host educational seminars, and build personal rapport to secure sales. This traditional reliance on face-to-face relationship building is exactly why the sources note that the wholesaler's engagement model is currently having to adapt to a new hybrid reality.

A major challenge in asset management marketing is «fragmented data,» where advisor interactions are trapped in silos, leading to uncoordinated engagement where both marketing and sales target the same advisor simultaneously. Marketing measurement in asset management must account for the fact that brand awareness grows slowly and compounds; while lower-funnel metrics like webcast attendance are easier to connect to conversion, prioritising them at the expense of brand building can hamstring long-term growth opportunities.



### Growing Audience of Retail Investors

Two primary groups of investors are prime audiences for asset managers:

- Institutional Investors, who move large blocks of shares that influence market trends. Reaching them necessitates marketing to have detailed financial insights, robust wealth data, and demonstrations of complex regulatory navigation.
- Retail Investors, who are narrowing the gap with improved access to research, are non-professional individuals who invest their own money to achieve personal financial goals. For instance, a significant portion of American households participate in the stock market through retirement savings accounts like IRAs and 401(k)s

As asset managers increasingly build direct-to-consumer engines to engage retail investors and bypass traditional intermediaries, they face a highly complex marketing measurement landscape. While the sales cycle for retail and high-net-worth investors is typically shorter (3 to 9 months) than for institutional clients, proving the exact return on investment (ROI) involves several major hurdles:

01

### The Slow Compounding of Brand Awareness:

Brand authority and awareness grow slowly and compound over time, meaning the true value of upper-funnel marketing is real but not always immediately measurable. While it is easier to connect lower-funnel engagement metrics such as webcast attendance, email nurturing, and thought leadership downloads directly to conversions, prioritising only these easily measurable bottom-funnel tactics can ultimately hamstring long-term growth.

02

### Strict Regulatory Compliance:

Retail financial marketing is heavily scrutinised. Measurement and marketing teams must operate under «compliance by design» to avoid steep penalties. For instance, FINRA Rule 2210 dictates that retail communications must be fair, balanced, and free of exaggerated claims, while the SEC Marketing Rule strictly governs how gross, net, and hypothetical performance data can be advertised to prevent misleading retail investors

03

### Signal Loss and Privacy Hurdles:

Financial marketers are increasingly constrained by stringent data privacy regulations (like GDPR and CCPA) and the disappearance of third-party cookies, making it nearly impossible to track individual user journeys deterministically. To maintain compliance while still accurately measuring marketing impact, the industry is pivoting away from deterministic tracking (which relies on specific, identifiable user data). Instead, the focus is shifting towards probabilistic measurement and causal inference. These advanced algorithmic approaches allow marketers to estimate the likelihood of marketing efforts influencing consumer behaviour without compromising user privacy.



## d. Retail Insurance: Complex Risk Factors and Price Elasticity

Marketing measurement in retail insurance is highly complicated because insurers must constantly balance customer acquisition with complex risk factors, claims costs, and strict fairness regulations. The core issue in insurance marketing lies in determining the optimal price for policies and accurately assessing their true price elasticity.

Insurers must untangle the impact of their specific marketing efforts such as offering a new rewards programme from broader macroeconomic market conditions that causally impact the overall demand for policies.

### Claims Experience and Fair Decisioning

In retail insurance, a core driver of long-term profitability and customer retention is the claims experience itself. Insurers are increasingly relying on advanced causal models to assess the causal impact of a customer's claims experience on their likelihood to renew their policy. Furthermore, these models help evaluate the difficult financial and operational trade-off between improving the customer experience by auto-approving claims versus the costs associated with investigating fraud.

Finally, much like in retail banking, insurers must operate with «compliance by design» to ensure that their policy approval processes especially for sensitive areas like Life & Health policies remain fair and unbiased across all scenarios.



# 02. Causal Methodologies for Measurement in FSI.

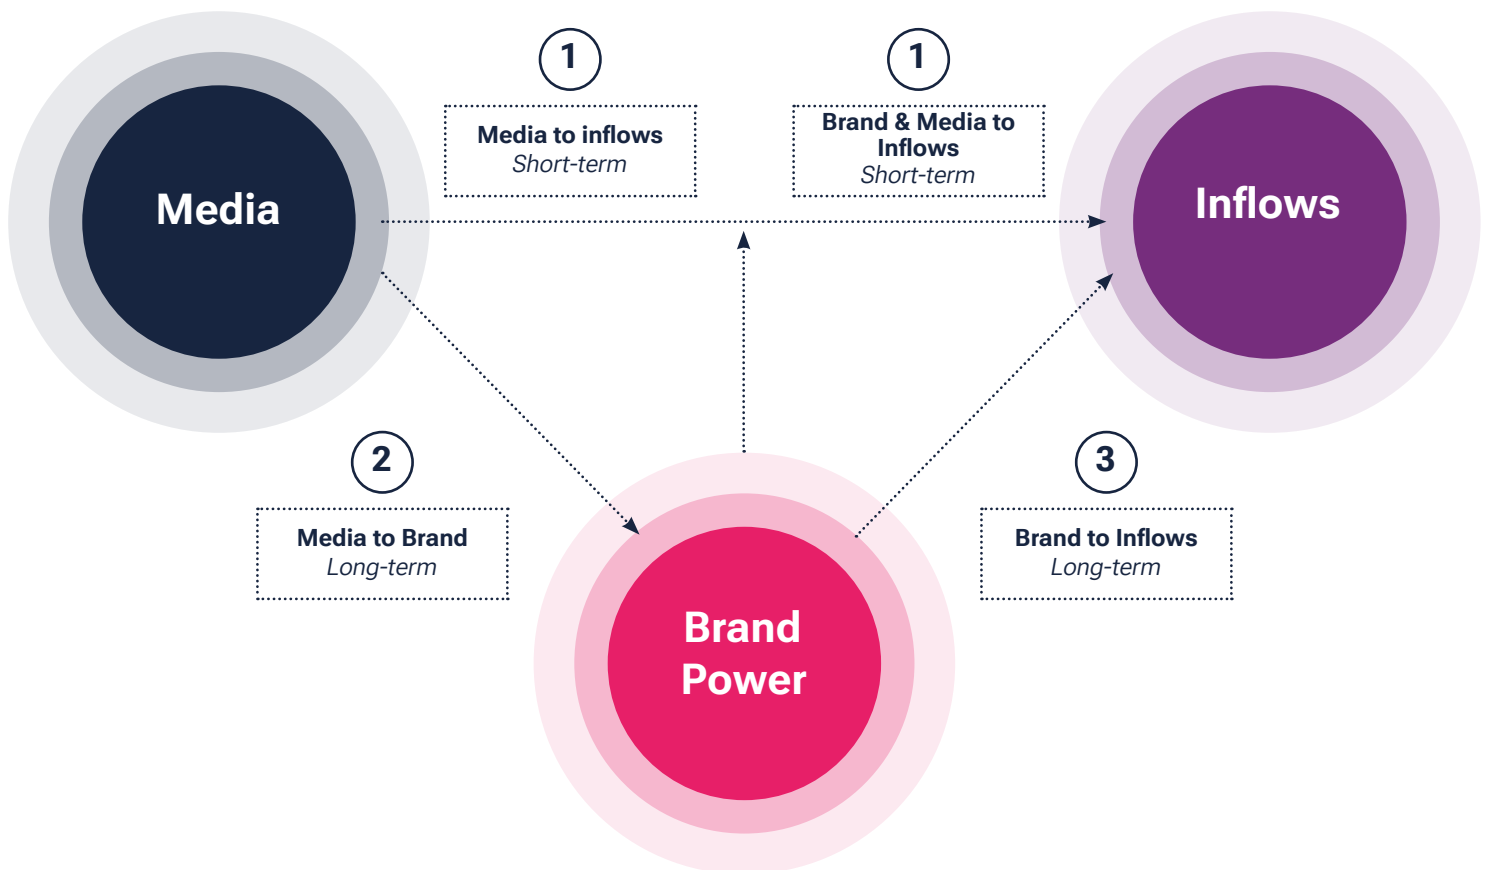
To establish a defensible return on investment, we believe Financial Service institutions should leverage three overarching methodologies. These techniques are particularly valuable in scenarios where randomised experiments are unfeasible due to cost or regulatory constraints.

## A. Observational Methodologies.

Observational Methodologies in marketing measurement are commonly and usually referred to as Marketing Mix Models (MMM).

There are two approaches in this space - Bayesian Regression and Causal Directed Acyclic Graphs - with the latter

being a more expanded and complex mathematical illustration of the former.



In order to capture the compounding value of brand, the full impact of upper and mid-funnel media as well as the long sales cycles (when involving intermediaries like financial advisors), it is vital to construct full funnel MMMs. Full MMMs capture:

- Short term effect of media on inflows / net cash flow
- Long-term effect of brand on inflows / net cash flow
- Short term of brand on inflows/ net cash flow
- Long-term effect of media on inflows / net cash flow via brand

## Bayesian Regression Design for Marketing Mix Modelling

This is a foundational statistical approach used within Marketing Mix Modelling (MMM) that uses aggregate, time-series data to quantify how much each marketing investment, macroeconomic condition, and competitive force contributes to a core business outcome such as net cash flow, Assets Under Management (AUM) growth, or new account openings.

Traditional MMMs, however, were built for a simpler world: single-channel campaigns where sales conversions were the primary KPI and effects played out over weeks, not quarters. Financial services breaks these assumptions across every dimension. Brand trust compounds over years. Customers consult financial advisors before acting. Interest rate cycles distort baseline performance in ways that dwarf the signal from any single campaign. Marketing measurement in FSI demands an architecture that can distinguish what marketing does right now from what it builds over time and prove the value of both.

## The Two Phenomena That Make FSI Measurement Hard and How the Model Solves Them

Two fundamental properties of advertising make financial services measurement uniquely challenging, and both are embedded in the Bayesian framework Artefact employs.

The first is **carryover**. Advertising does not switch off the moment it stops running. In his seminal 1979 analysis of television advertising, Simon Broadbent introduced the concept of «adstock» the lingering, decaying effect of a media impression on consumer behaviour over time. A brand campaign for a wealth management product running in Q1 may still be shaping deposit decisions in Q2 and Q3, even if no further spend occurs. In financial services, where consumers deliberate for months before switching providers or opening a new product, these carryover effects are especially pronounced. Failing to account for them systematically undervalues brand investment and misattributes performance to the most recent campaign, rather than the sustained brand presence that created the conditions for conversion.

The second is **saturation** – the principle that each additional dollar invested in a media channel produces progressively smaller incremental returns. Beyond a certain threshold, audiences have been reached, frequency caps have been hit, and further spend yields diminishing value. In their landmark 2017 paper, Jin, Wang, Sun, Chan, and Koehler at Google Research formalised both phenomena into a unified probabilistic framework, using the Hill function to model saturation curves and geometric decay to model adstock – providing the mathematical architecture that underlies modern Bayesian MMM at institutional scale [14].

## Why Bayesian? The Case for Probabilistic Budgeting

What distinguishes Bayesian regression from classical regression methods is not precision alone it is defensibility. A Bayesian model does not return a single point estimate for a channel's return on investment. It returns a full probability distribution, explicitly quantifying the range of plausible outcomes and the model's degree of certainty in each estimate.

For a CMO presenting budget allocations to a CFO, a risk committee, or an institutional board, this distinction is material. Rather than asserting «brand media drove 18% of inflows this quarter,» a Bayesian framework allows a marketer to state, with statistical rigor: «there is a 90% probability that brand media drove between 14% and 23% of inflows, conditional on the prevailing interest rate environment.» In financial services, where governance standards are high and scrutiny of marketing spend is increasing, this kind of probabilistic accountability transforms marketing from a cost center into a measurable, defensible driver of enterprise value.

Bayesian methods also allow institutions to encode prior knowledge directly into the model – a critical advantage in FSI, where usable historical data is often limited by regulatory changes, product launches, or structural market shifts. Rather than beginning each measurement cycle from scratch, Bayesian inference enables teams to carry forward calibrated beliefs from previous campaigns and update them systematically as new evidence accumulates, improving both model stability and the speed at which insights can be actioned [15].

## A Three-Model Architecture for Full-Funnel FSI Measurement

Standard single-equation MMMs cannot capture the full causal value of marketing in financial services. Binet and Field's landmark IPA research (2013), drawing on nearly 1,000 advertising effectiveness case studies across 30 years, established that brand investment and sales activation operate on fundamentally different time horizons – and that the most significant long-term business effects of marketing manifest not in immediate conversion spikes, but in the sustained elevation of the performance baseline over time [16]. Marketing builds brand equity; brand equity sustains the baseline; the baseline is what makes future campaigns more effective. A model that captures only the short-term effect will systematically undervalue brand, misallocate budget, and erode competitive positioning.

To measure the full value chain from media impression to short-term conversion to long-term baseline growth, Artefact deploys a three-

model system that separates and quantifies these effects in sequence before consolidating them into a single, integrated view of total marketing value.

### Model 1 - Short-Term Causal Effects

The first model measures the direct, near-term impact of all marketing and media activity on core business KPIs gross inflows, new account openings, or net cash flow typically over a horizon of up to two months. The model simultaneously controls for macroeconomic variables (interest rate movements, inflation indices, FX rates), pricing dynamics, and loyalty or reward incentive programs. This is essential in FSI: without explicitly modelling these forces, a rising interest rate environment or a competitor's pricing change can be incorrectly attributed to a marketing campaign and vice versa.

For institutions operating across multiple markets – national retail banks, multi-regional asset managers, or global

insurance groups – Artefact deploys a Hierarchical Bayesian architecture within this model. Formalised at the geo level by Sun, Wang, Jin, Chan, and Koehler (2017), hierarchical models allow a measurement framework to share statistical learnings across geographies while still respecting local market differences [17]. A retail bank's deposit acquisition campaign may perform very differently in New York than in Boston – but what the model learns about media saturation rates in one market can inform priors in another, dramatically improving estimate precision in data-sparse regions without imposing a uniform national coefficient.

The output of Model 1 is a defensible decomposition of short-term revenue: the percentage contribution of each marketing channel, each macroeconomic driver, and each pricing lever presented as probability distributions, not point estimates.

What is the impact of my spend on short-term (ST) impact?



■ Spend ST contribution to KPI  
■ KPI Baseline

## Model 2 - Baseline Decomposition and Long-Term Brand Effects

In every marketing mix model, there is a residual: the «baseline» the portion of business performance that cannot be explained by near-term marketing or observable external variables. In standard MMM practice, this baseline is treated as a fixed intercept, essentially a black box. Artefact's framework treats it as the most important signal in the data.

The theoretical foundation for this approach originates with Nerlove and Arrow's canonical 1962 model of advertising as an investment in goodwill – a depreciating stock of brand equity that accumulates from sustained marketing activity and decays in its absence [18]. The key insight is that brand equity does not manifest as a short-term conversion spike; it manifests as a sustained elevation in the business baseline over time. An institution with strong brand health will see its baseline inflows hold steady – or grow – even in the absence of active campaigns. An institution with eroding brand health will watch its baseline deteriorate regardless of how effectively it optimises performance media.

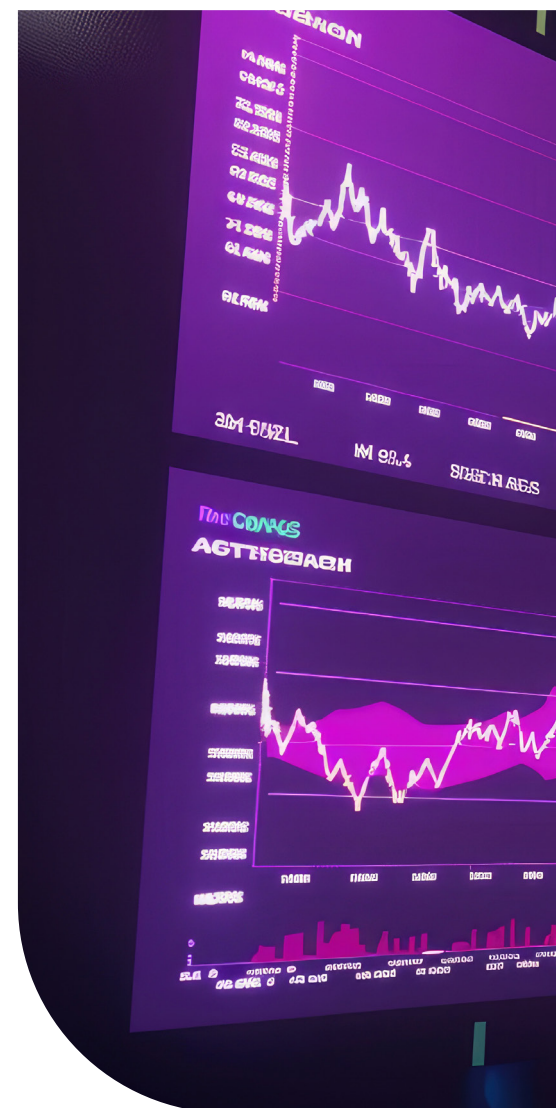
There is, however, a critical technical challenge between this conceptual logic and its practical implementation. The baseline produced by Model 1 is not a smooth, interpretable time series. It is an aggregate residual capturing everything the short-term model could not explain which means it contains both the genuine long-run brand signal Artefact seeks to measure and a layer of short-term noise: model residuals, idiosyncratic weekly fluctuations, and transient events with no informational content about brand health. If this raw baseline were regressed directly against brand power metrics, that noise would contaminate the relationship, producing unreliable estimates.

A spike in the raw baseline driven by a competitor's temporary price promotion or a one-off market event would be incorrectly attributed to brand strength and brand's true contribution would be masked by volatility it did not cause.

To isolate the true long-run signal, Artefact applies a **Gaussian transformation** to the raw baseline before entering Model 2. Gaussian processes are a class of non-parametric probabilistic models particularly well-suited to this task: they decompose a noisy time series into its constituent components using kernel functions that encode specific assumptions about the smoothness and temporal scale of each component [19]. By selecting a kernel that captures slow-moving, smooth trends – such as a squared exponential or Matérn kernel – while treating shorter-cycle fluctuations as noise, the transformation produces a smoothed, long-run baseline series that reflects only the gradual, brand-driven movements in the underlying business trajectory.

The intuition is direct: brand power does not move on a weekly timescale. It shifts over months and years compounding with sustained marketing investment and eroding in its absence. The Gaussian transformation aligns the temporal frequency of the dependent variable with the temporal frequency of the driver being measured. What gets regressed against brand power metrics is therefore the signal that brand actually moves, not the noise layered on top of it. This is the methodological step that makes the brand attribution in Model 2 credible and it is what separates a rigorous full-funnel measurement framework from one that conflates brand equity with residual variance.

Model 2 then takes this Gaussian-transformed baseline as its dependent variable and regresses it against brand power metrics (awareness, consideration, net promoter scores), macroeconomic trend components, and structural pricing shifts. The output is a decomposition of the long-term baseline: what share of the institution's sustained performance trajectory is attributable to brand health, and how that brand equity is trending across time. For a CDAO, this unlocks a previously unmeasurable dimension of marketing's enterprise value one that is particularly consequential in financial services, where brand trust is a primary driver of customer primacy and product cross-sell.



**Model 3 – Media’s Indirect Effect on Brand Power**

The third model closes the causal loop between near-term media investment and long-term brand equity. It estimates how changes in media spend – both performance and brand formats – causally drive movements in brand power metrics, treating brand equity as an outcome variable rather than a fixed input.

This model enables Artefact to trace the complete causal pathway of any marketing investment: its immediate effect on inflows (measured in Model 1),

its contribution to brand equity over time (Model 3), and the downstream effect of that brand equity on the long-term performance baseline (Model 2).

Together, the three models produce a total marketing value estimate that is both comprehensive – capturing short-term activation and long-term brand compounding – and causally grounded, satisfying the evidential standards required by institutional governance and regulatory audit.

What is the contribution of my spend drivers to brand metrics?



- Spend LT contribution to brand
- Brand Baseline

What is the impact of my spend on short-term (ST) effect?



- Spend ST contribution to KPI
- Spend LT multiplier to KPI
- External Factors to KPI

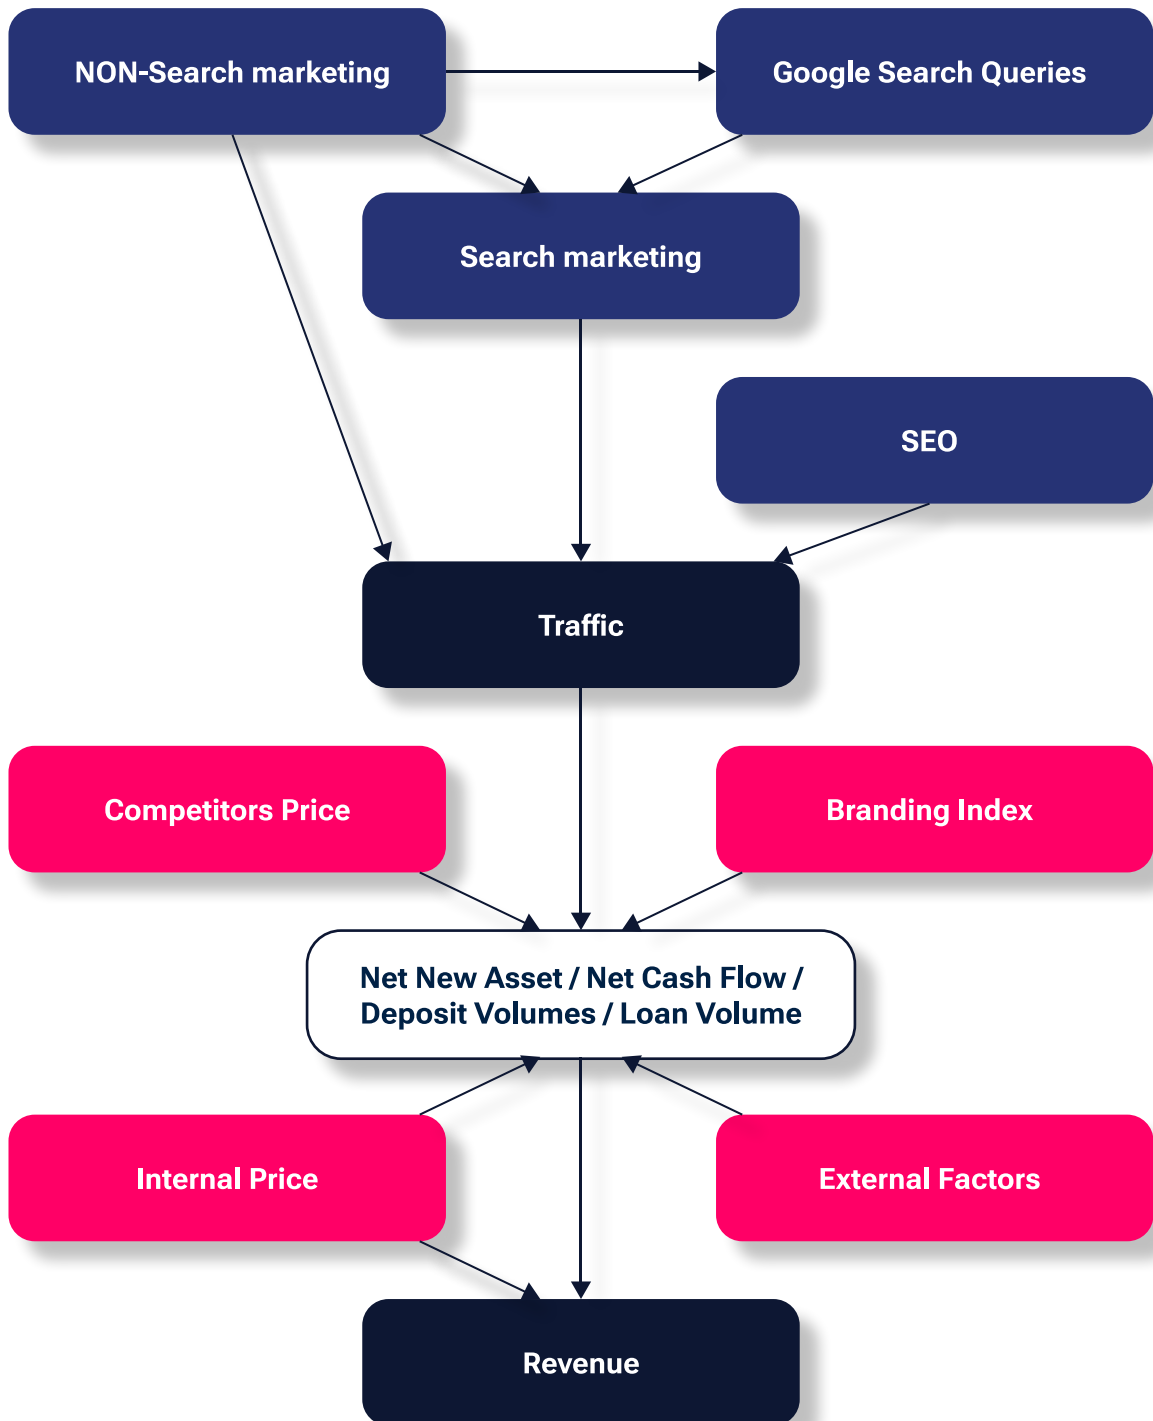


## An Alternative Architecture: Structural Causal Models and the DAG-Based Approach

The three-model Bayesian regression architecture described above is effective and operationally tractable. But it is sequential: each model is estimated independently, with outputs passed forward as inputs to the next stage.

This means that uncertainty does not fully propagate across the chain – an estimate of brand power in Model 3 carries no formal memory of the confidence interval it inherited from Model 1.

For institutions where the evidentiary bar is highest – those subject to regulatory model risk management, or those making nine-figure budget decisions – a fully unified causal framework may be warranted. That framework is the Structural Causal Model (SCM).



### From Regression to Causation: The SCM Difference

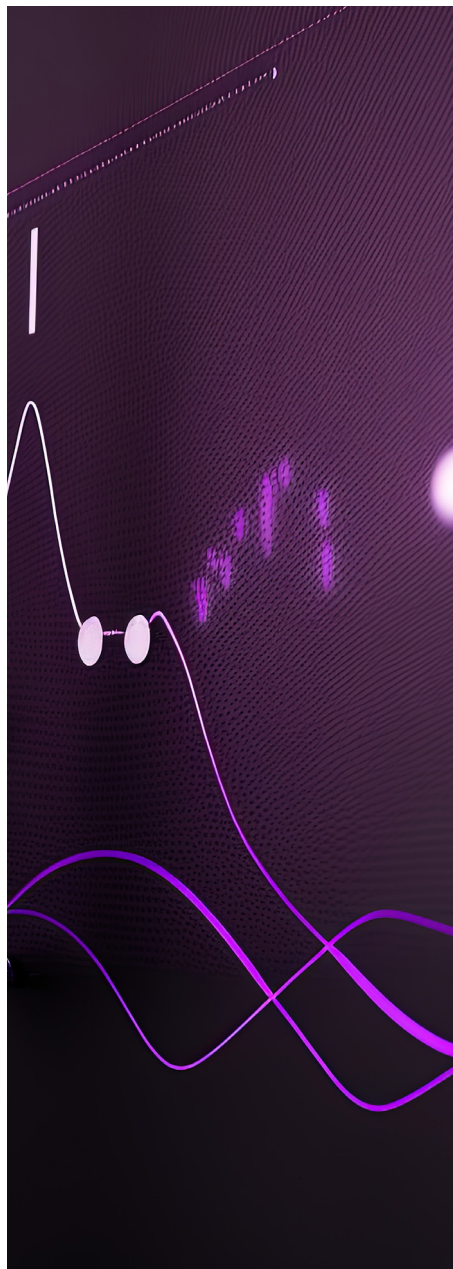
An SCM, formalised within Judea Pearl's causal hierarchy, moves beyond statistical association to encode the actual data-generating process of a system [20]. Rather than fitting three separate regressions and linking them procedurally, an SCM represents all variables – media spend, brand power, macroeconomic conditions, business outcomes – as nodes in a Directed Acyclic Graph (DAG), connected by directed edges that assert causal direction. Each node is governed by a structural equation that specifies exactly how its value is determined by its causal parents and a noise term. Crucially, this formalism supports interventional and counterfactual reasoning in a way that regression chains cannot: rather than estimating correlations from observational data, an SCM can answer questions like «what would net inflows have been in Q2 had we allocated 20% more to brand media in Q1, holding all else equal?» – a question that requires the mathematical machinery of Pearl's do-calculus to answer rigorously.

The practical implementation of this framework at scale is enabled by DoWhy, the open-source causal inference library developed at Microsoft Research by Amit Sharma and Emre Kiciman (2020) [21]. DoWhy's Graphical Causal Model (GCM) extension allows practitioners to specify a DAG, assign causal mechanisms to each node, fit the model to observational data, and then query it for direct effects, indirect effects, counterfactuals, and root cause attributions – all within a unified probabilistic framework.

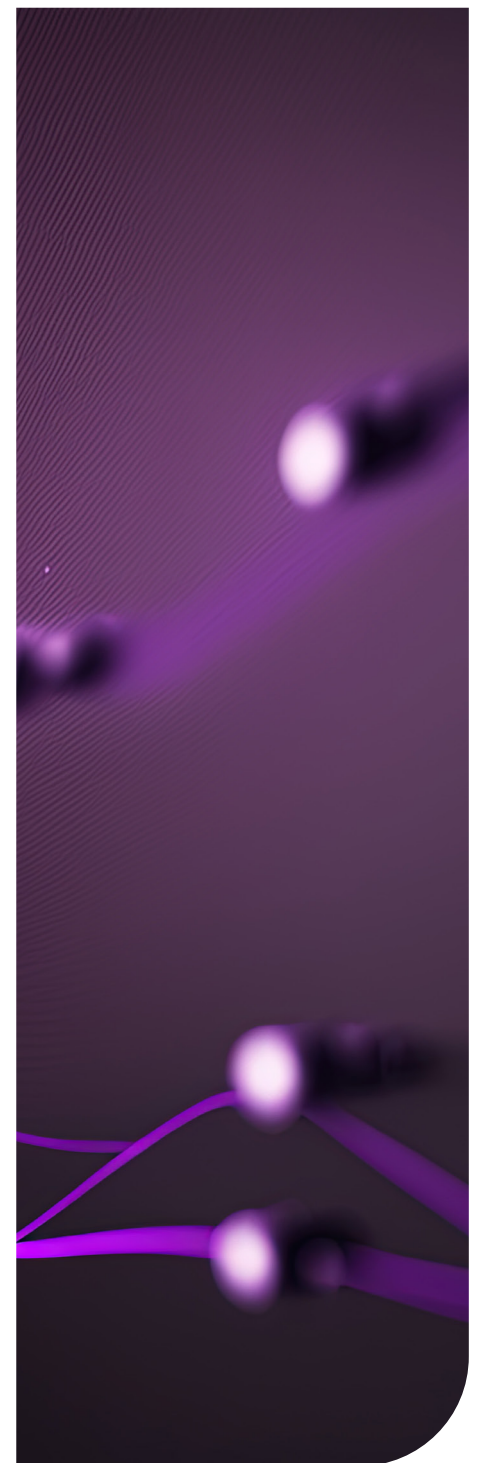


### The Marketing DAG for FSI

In the SCM formulation of the full-funnel FSI measurement problem, the causal graph takes the following structure. Media spend (M) operates through two pathways simultaneously: a direct short-term path to business outcomes (Y), and an indirect long-term path mediated through brand power (B) and the long-run baseline (L). Macroeconomic variables (E) enter as observed confounders, directly influencing both baseline dynamics and business outcomes. Brand power is treated as a latent variable – it is not directly observed in the data, but is inferred through its effects on the baseline and, in turn, on outcomes.



This graph makes explicit something the three-model regression chain only implies: that brand power is not merely a covariate to be included in a regression, but a mediating causal mechanism with its own temporal dynamics. The SCM forces the analyst to commit to that structure before seeing the data – which is both its greatest methodological strength and the source of its higher implementation burden.



### Embedding the Gaussian Baseline within the SCM

The most methodologically significant contribution of the SCM formulation is how it handles the long-run baseline and specifically, how it subsumes the Gaussian transformation described in the previous section into the causal structure itself.

In the sequential three-model approach, the Gaussian smoothing of the Model 1 baseline is a pre-processing step: a separate procedure applied to the raw residual before it enters Model 2. The SCM approach makes this unnecessary by modeling brand power directly as a latent variable with a Gaussian Process prior. Drawing on the framework developed by:

Witty, Ha, Gretton, and Dhir (2020) for causal inference with structured latent confounders, the brand equity node  $B(t)$  is specified as a GP-governed process: its evolution over time is assumed to be smooth and slowly varying, governed by a kernel function (such as a squared exponential kernel) that encodes the temporal scale at which brand equity changes [21]. Formally:

$$B(t) \sim GP(\mu(t), K(t, t'))$$

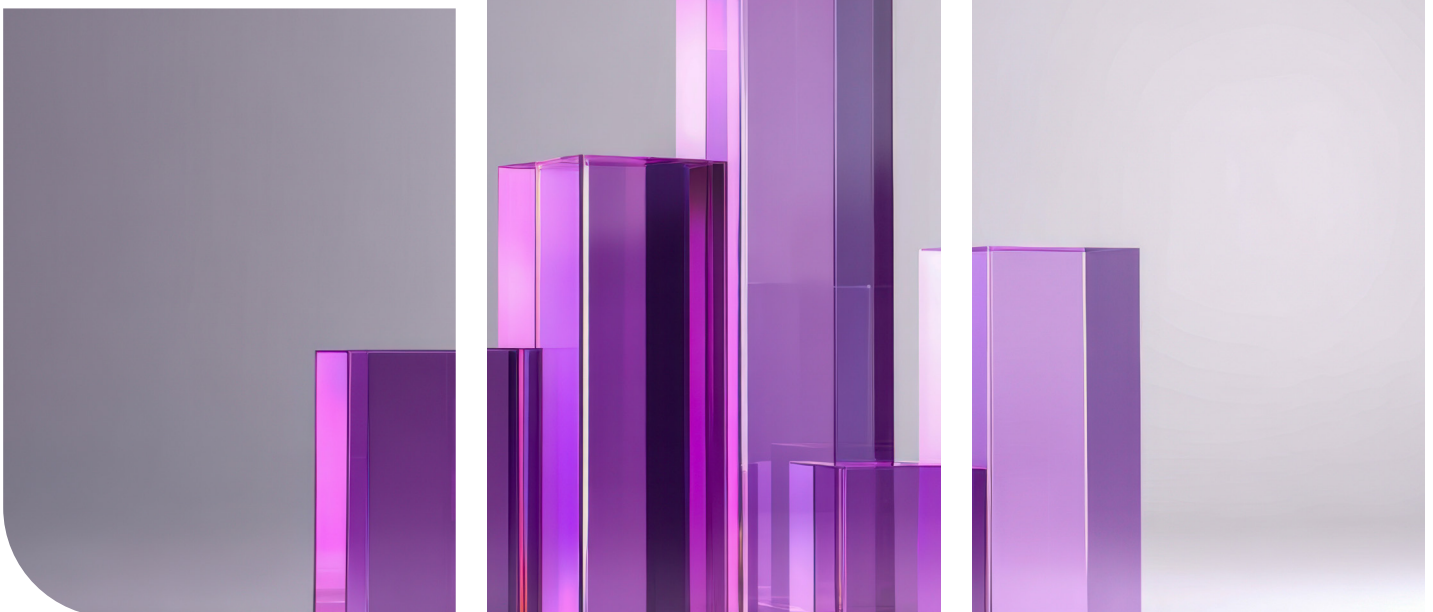
where  $K(t, t')$  is a kernel that assigns high covariance between nearby time points and decaying covariance across longer intervals mathematically enforcing the intuition that brand power does not oscillate week-to-week but compounds or erodes over months and years.

The Gaussian process prior ensures that when the SCM is fitted to observational data, the inferred brand equity trajectory will inherently possess the slow-moving, noise-filtered character that the GP transformation imposes post-hoc in the regression framework.

This has a further advantage that is particularly relevant for CDAO audiences: the brand equity trajectory is now a fully Bayesian latent variable, with a posterior distribution that propagates uncertainty about brand health forward into every downstream estimate.

The consequence is elegant. What the three-model approach achieves through an explicit Gaussian smoothing step applied to a residual series, the SCM achieves through the structural specification of brand equity's temporal dynamics. The smooth, long-run signal is not extracted it is defined as a property of brand power within the model.

Rather than a point estimate of brand's contribution to the baseline, the SCM delivers a probability distribution quantifying not just the most likely brand trajectory, but the full range of plausible brand dynamics consistent with the observed business outcomes.



## The Practical Trade-Off

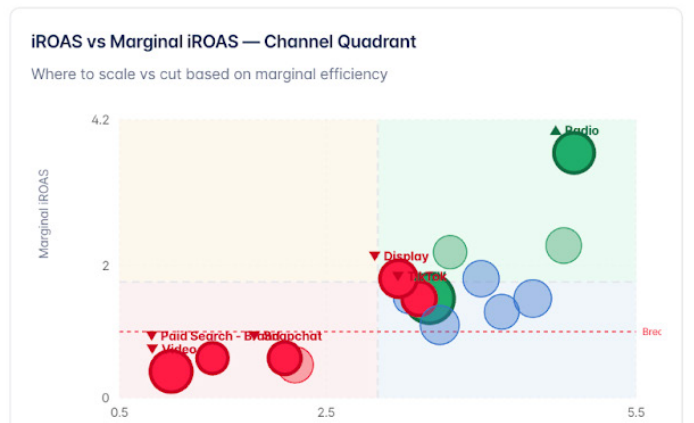
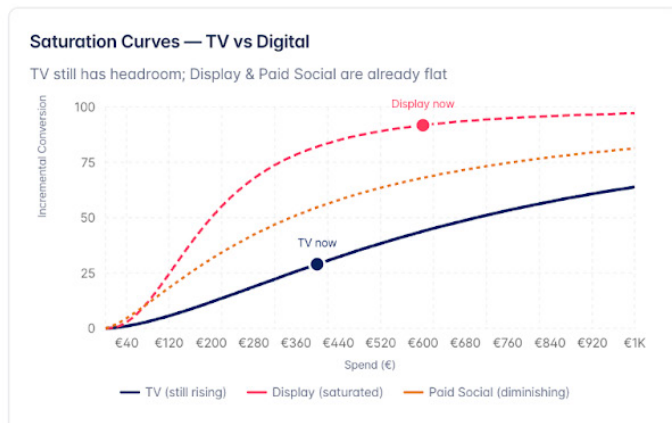
The SCM approach is not a straightforward upgrade from the three-model framework – it is a different philosophical and operational commitment. Its advantages are substantial: unified uncertainty propagation, formal counterfactual reasoning, regulatory explainability through explicit causal graphs, and the elimination of procedural coupling between model stages. Peters, Janzing, and Schölkopf (2017) demonstrate that this kind of structural formalism is also more robust to distributional shifts – an important property in financial services, where market regime changes can invalidate the assumptions of models fitted on historical data [23].

Its challenges are equally real. The DAG must be specified a priori, which requires domain expertise and stakeholder alignment on the assumed causal structure a non-trivial process in organisations where marketing and finance teams may hold competing theories of how media drives outcomes. Fitting GP-augmented SCMs is computationally intensive and requires more data than sequential regression. And the outputs, while more rigorous, are harder to communicate to a non-technical board audience than a straightforward channel contribution chart.

Artefact’s recommendation for FSI institutions is pragmatic: begin with the three-model Bayesian architecture as the operational measurement baseline, and evolve toward the SCM formulation as data maturity, institutional alignment, and governance requirements demand it. The two frameworks are not competing alternatives they are sequential stages of measurement sophistication, appropriate for different points on the organizational maturity curve described in Section 3.



### Why? — MMM Evidence



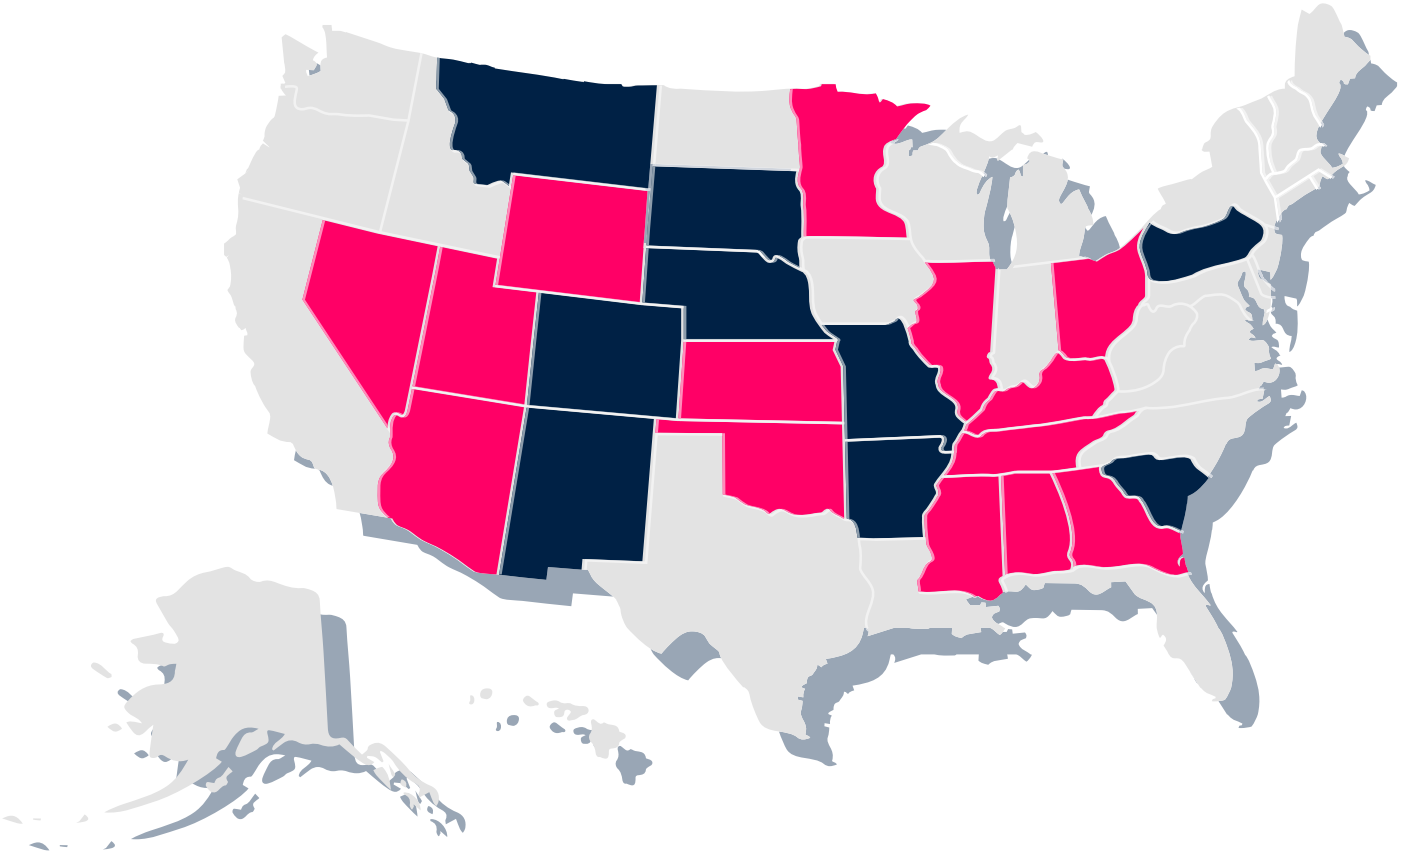
## B. Observational Methodologies.

Quasi-experimental methodologies are techniques leveraged in real-world situations where it is impractical to randomly assign participants or subjects into groups. These methodologies allow for estimating the causal impact of a marketing intervention in the absence of random individual assignment.

There are three main steps when employing quasi-experimental designs:

### 1. Match-Market:

Identify like-for-like geos or regions when treatments or interventions in the form of specific marketing campaigns can be applied



● Treatment 1

● Treatment 2

● Holdout

### 2. Power Analysis to estimate Budget & Duration Threshold:

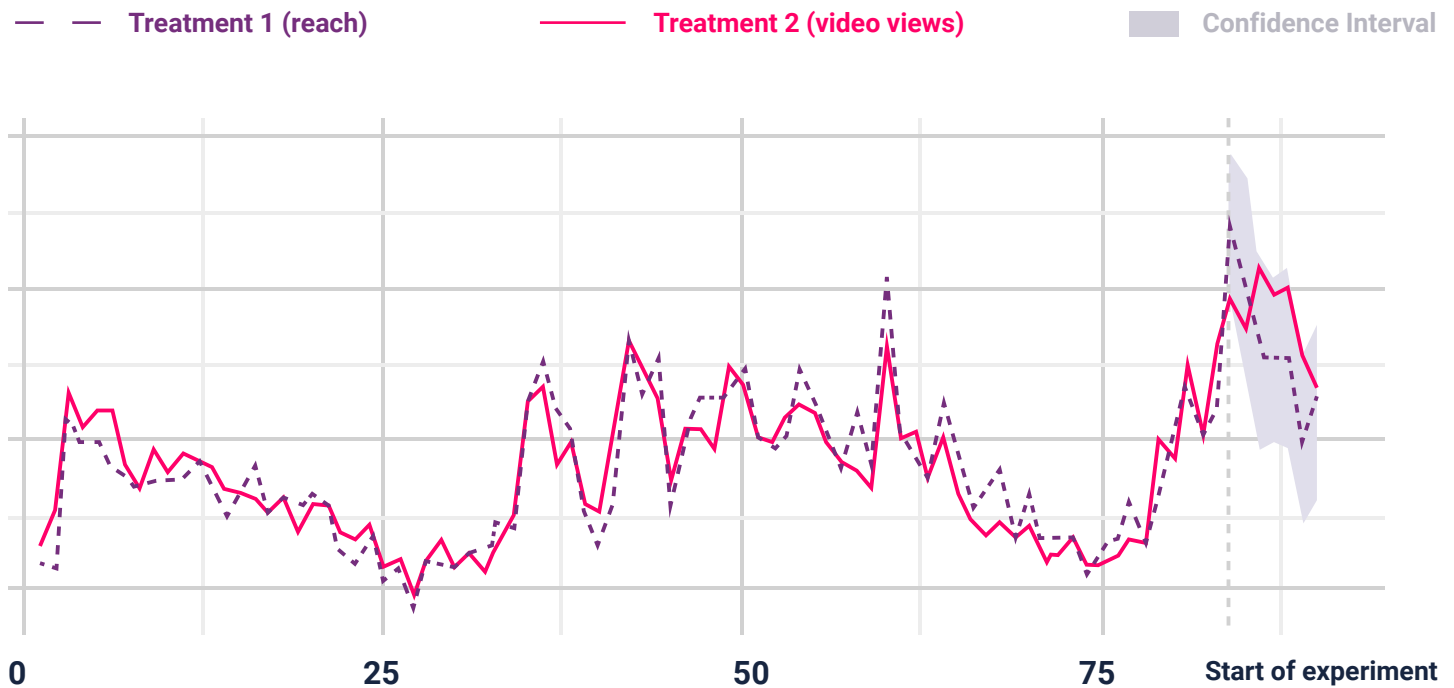
The Power analysis is a critical step in designing a geographic tests. It evaluates whether the experiment design, given its parameters, can reliably detect a meaningful effect size. Rather than relying solely on theoretical variance assumptions, we use a simulation-based

power analysis grounded in real-world Gross Inflows/ Net Cash Flow data. This approach provides an empirical check on whether the design is likely to produce statistically significant results, and helps quantify the test's sensitivity under realistic conditions.

The Power Analysis also determines the minimum budget and duration required to measure a statistically significant effect.

### 3. Treatment effect estimation:

Measure the Average Treatment Effect (ATE) against a counterfactual leveraging the control or hold-out group DMAs as the basis



#### Difference-in-Differences (DiD)

DiD estimates the causal impact of a marketing intervention by comparing the change in outcomes over time between a group that received the treatment and a comparable group that did not. By taking two differences – one across groups and one across time – the method cancels out both pre-existing baseline differences and any macro-level trends affecting everyone simultaneously, isolating the incremental effect of the intervention itself. In FSI, where campaigns frequently roll out in waves across regions or customer cohorts, the Callaway and Sant’Anna (2021) extension of DiD is particularly valuable, as it corrects for the bias that standard regression produces when treatment timing is staggered across units [24].

Consider a retail bank that activates a new deposit acquisition campaign in three cities (treatment cities) while holding spend flat in three comparable cities (control cities). A DiD design compares the change in gross deposit inflows in the treated cities before and after the campaign launch against the same change in the control cities. The effect of a central bank rate decision or a competitor’s product launch forces that affect both groups equally is differenced out, leaving a clean estimate of the campaign’s true incremental contribution.

#### Synthetic Controls

This involves constructing artificial comparison groups for scenarios where traditional randomised A/B testing is not feasible. This method creates a weighted counterfactual out of untreated units to mimic what would have happened to the treated group if the intervention had never occurred.

Consider a retail bank, for instance, that wants to measure the true causal impact of a regional marketing campaign or a regional price increase. Because the bank cannot easily randomise these macro-level interventions across individual customers, it can use synthetic controls. By taking aggregated data from multiple untreated regions or branches that historically mirror the target region, the bank can create a tailored synthetic counterfactual to accurately estimate the incremental impact of the new branch or campaign.

### Regression Discontinuity (RDD)

RDD leverages a clear assignment threshold or natural cutoff—such as a credit score or account balance—to identify causal effects. By comparing individuals who fall just above and just below this strict threshold, marketers can essentially mimic a randomised controlled trial, as the assignment around the cutoff is as good as random.

For instance, imagine an asset management firm wants to measure the impact of a premium reward and advisory program on client consumption and engagement. The program is only offered to clients whose total assets with the bank exceed a strict eligibility threshold, such as \$500,000. By using RDD to compare clients narrowly hovering just above and just below this asset threshold, the firm can isolate the true causal effect of the premium program's availability without the bias of unmeasured wealth selection, revealing how the program alters subjective expectations and investment behaviours.

### Double Machine Learning (DML)

DML combines outcome prediction with treatment assignment modelling to isolate causal impacts in complex, high-dimensional data environments. It uses advanced machine learning to separate true causal effects from confounding variables by «orthogonalising» the estimating equations with respect to nuisance parameters, thereby mitigating both bias and variance

Consider a digital payments company, in this case, that wants to understand how assigning higher credit limits causally alters the shape of consumer spending distributions.

Because credit assignment is often correlated with confounding factors that influence spending, measuring true impact is difficult. DML allows the platform to estimate these causal effects amidst high-dimensional data, revealing that increased credit lines do not just uniformly increase spending, but disproportionately shift consumers toward higher-value, discretionary luxury purchases at higher quantiles.



## C. Multi-Touch Attribution (MTA)

While MTA is losing traction in today's privacy-first world, it remains relevant in the world of financial services, where a significant portion of marketing activities are executed leveraging Owned Media channels such as email and app pushes.

**The First-Party Data Advantage:** The primary reason MTA has been heavily criticised recently is the «privacy wall»—the deprecation of third-party cookies and strict privacy regulations (like GDPR and CCPA) that make tracking individual users across external ad networks nearly impossible. However, Owned Media campaigns rely on first-party data, which is generated directly from customer interactions like website visits, app usage, and email engagement. Because banks collect this data directly and with user consent, MTA can effectively track these touchpoints without being obstructed by privacy restrictions.

**The Unified Approach:** To ensure that MTA does not fall into the trap of over-crediting owned channels that might just be «harvesting» demand (e.g., an email claiming 100% credit for a conversion that was actually driven by a mid or upper funnel campaign), FSI marketers are integrating MTA with the MMM, and calibrating with Incrementality Tests

### The Engine behind the MTA

Attribution modelling has seen significant advancements in recent years, driven by the rise of deep learning. At its core, a customer journey is a sequence of touchpoints, making it an ideal candidate for models designed to handle sequential data. This is where Recurrent Neural Networks (RNNs) come into play, particularly a specialised variant known as Long Short-Term Memory networks (LSTMs).

LSTMs excel at capturing dependencies in sequences, allowing us to train models on customer journeys and predict whether a given series of touchpoints will lead to conversion. However, predicting conversion alone isn't enough. The real challenge in attribution lies in fairly distributing credit across the touchpoints that contributed to a successful outcome.



To address this, we can borrow a powerful concept from modern AI: the Attention mechanism, a key innovation behind transformer-based models like ChatGPT. By integrating Attention into an LSTM-based framework, we gain two critical advantages:

01

#### Conversion Prediction:

The ability to model complex customer journeys and estimate their likelihood of success.

02

#### Fair Attribution:

A data-driven way to assign weights to each touchpoint, determining its true influence on the conversion.

With enough historical data, such a model can learn to identify which campaigns drive conversions, which touchpoint combinations are most effective, and how interactions across channels impact the customer journey.

Beyond touchpoints, this approach also allows for the integration of customer metadata, enabling deeper personalisation. If you have a 360-degree view of your customers, you can feed demographic, behavioural, or contextual data into the model, enhancing accuracy and unlocking insights such as which personas are most likely to convert on a given channel.

These techniques allow marketing leaders to move beyond simplistic attribution models and gain a clearer, more actionable understanding of what truly drives results.

### Key Steps in Implementing an LSTM-based Attribution Model

01

Construct customer journeys from historical interaction data.

02

Build a custom AI model tailored to your customer data.

03

Train it to maximise its accuracy.

04

Extract the contribution of each touchpoint belonging to converting journeys, according to the model.

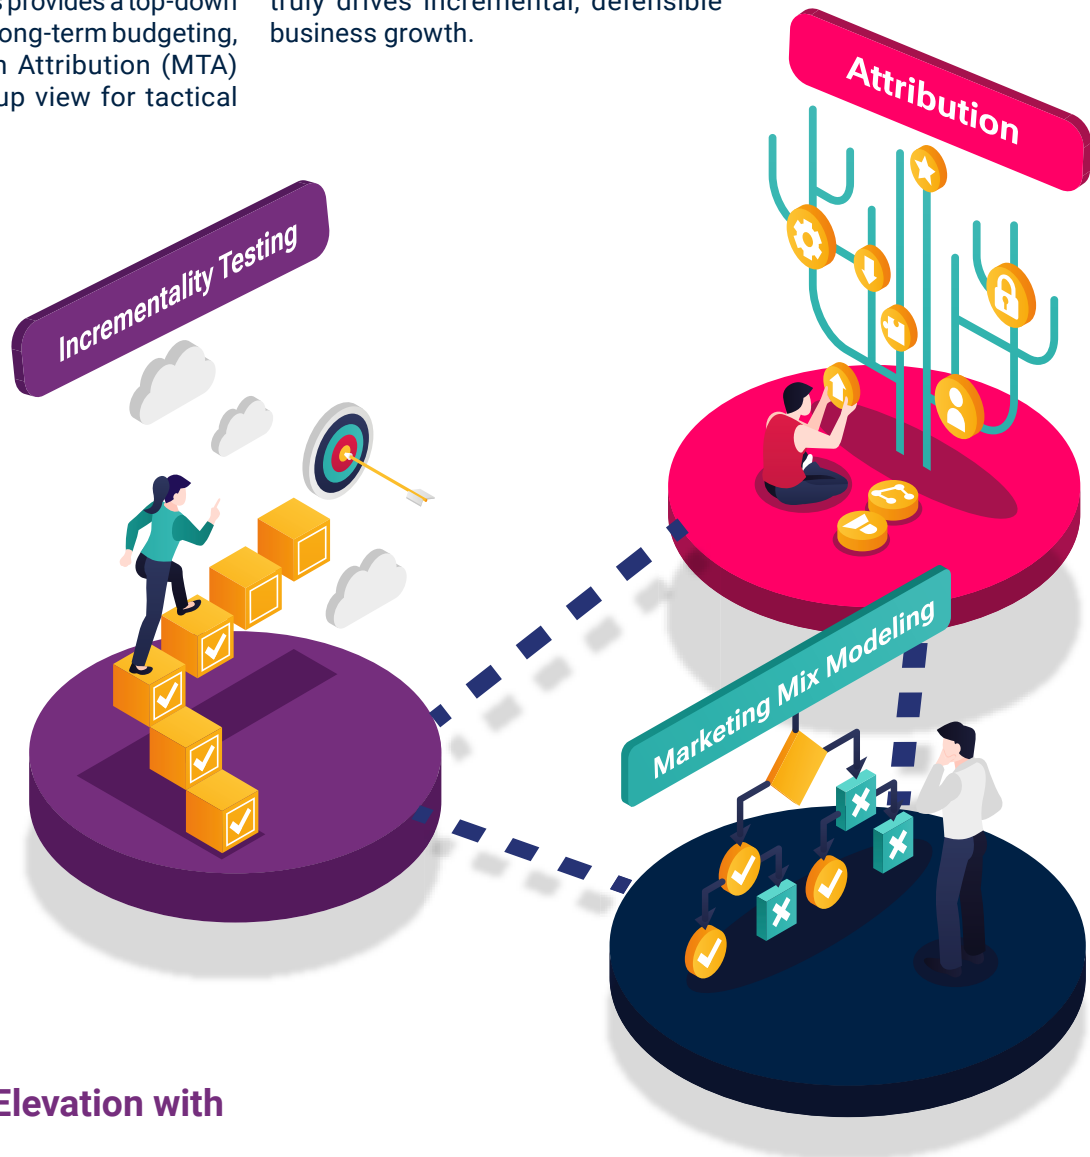
# Bringing It Together: The Unified mROI Framework

Each of the three methodologies above answers a different question. mROI is the framework that connects them into a single, defensible view of marketing value

Previously, organisations had siloed methodologies that rarely reached a consensus. MMMs provides a top-down strategic view for long-term budgeting, while Multi-Touch Attribution (MTA) offers a bottom-up view for tactical optimisation.

The modern standard is shifting toward a more unified mROI framework, which integrates these top-down strategic and behavioural approaches with in-market incrementality testing. This integration is essential because, while attribution models show where conversions appear, only unified measurement reveals what truly drives incremental, defensible business growth.

In a world of signal loss—exacerbated by the deprecation of third-party cookies and stringent data privacy regulations like GDPR and CCPA—marketers must transition from deterministic to probabilistic measurement.



## Intelligence Elevation with Agentic AI

The causal measurement methodologies described previously – Bayesian MMMs, Structural Causal Models, quasi-experimental designs, and multi-touch attribution – are technically capable of answering the questions that matter most to FSI Marketing leadership. The practical constraint from a technology perspective, however, has not been the methodology alone. It has been the complexity of running these systems

continuously, at institutional scale, and translating their outputs into decisions fast enough to influence the budget cycles and campaign activations they are designed to inform. Agentic AI directly addresses this constraint. By deploying systems of purpose-built AI agents across the measurement value chain, FSIs can operationalise their causal infrastructure in ways that were not commercially viable even two years

ago – reducing cycle times, lowering the dependency on specialist data science capacity at every step, and closing the gap between model output and business decision.

Four domains of the measurement value chain are most immediately tractable.

## Agentic Media Data Governance

The most persistent source of model failure in FSI marketing measurement is not methodological – it is data quality. Marketing data in financial services arrives from a wider array of sources than almost any other sector: CRM systems, digital advertising platforms, programmatic DSPs, branch transaction records, call centre logs, third-party aggregators, and increasingly, clean room environments operating under privacy-preserving computation protocols. Each source carries its own schema conventions, refresh cadences, coverage gaps, and failure modes. A pixel fires incorrectly. A Google Ads API call returns anomalous spend. A consent management update reduces cookie traffic by 30% in a single market. Without detection, these events silently corrupt the data that flows into attribution and MMM models, producing outputs that appear credible but are structurally unsound.

Agentic data governance addresses this systematically. Agents monitor incoming data feeds continuously, applying rule-based and configured anomaly detection to flag deviations from expected patterns – sudden drops in impression volume, implausible channel-level CPMs, mismatches between reported and billed media spend. Where data quality issues are identified, agents can quarantine affected data, trigger reconciliation workflows, and prevent corrupted observations from entering model training pipelines without human escalation. Equally important for FSIs operating under SR 11-7 and SS1/23 is the auditability that agentic governance enables: automated data lineage documentation, schema validation logs, and PII compliance checks generate the kind of model input records that independent validators and internal audit functions require [28][31]. This transforms data governance from a periodic remediation exercise into a continuous, traceable process.

In practice, consider a retail bank running a multi-channel deposit acquisition campaign across paid search, display, email, and branch marketing. An agentic governance layer continuously monitors each data feed, detects when the attribution platform is under-reporting mobile conversions following an iOS update, and flags the gap before it propagates into the quarterly MMM refresh – preserving the integrity of the model outputs that will inform the next budget reallocation decision.



## Simplification of Model Execution

The Bayesian MMMs and causal inference designs require meaningful data science expertise to configure, validate, and maintain. Posterior sampling, prior specification, convergence diagnostics, backtesting, and structural break detection are not tasks that marketing operations teams can absorb without specialist support. In most FSIs, this creates a direct bottleneck: the cadence of model refresh is constrained not by data availability but by analyst capacity. The result is quarterly – sometimes annual – model updates in an environment where media mix, consumer behaviour, and macroeconomic conditions can shift materially within weeks.

Agentic orchestration decouples model refresh cadence from analyst bandwidth. An orchestration agent can monitor for new data availability, trigger model retraining pipelines automatically, run pre-specified validation checks – holdout performance, posterior predictive accuracy, parameter stability against the prior run – and surface model health metrics to the analytics team for review. Where results fall within expected parameters, the refresh completes without manual intervention. Where anomalies are detected – a channel coefficient shifting outside its historical confidence interval, convergence failures, data coverage falling below a threshold – the agent escalates with a structured summary of the issue and the affected outputs. Analysts shift from running models to supervising them, concentrating their time on the model design decisions and strategic interpretation that genuinely require human judgement.

The commercial consequence is compressive. Shortening the model refresh cycle from four weeks to one – or from one week to twenty-four hours – does not merely make measurement faster. It changes the decisions measurement can influence. A weekly MMM refresh can inform in-flight budget reallocation. A daily attribution update can trigger campaign-level bid adjustments. The measurement stack moves from a retrospective reporting function to a live operational input.



## Agentic Decision Intelligence and Reporting

The most widespread failure in marketing analytics is not the model – it is the last mile. Gartner’s research consistently finds that a majority of organisations use analytics primarily to confirm decisions that have already been made, rather than to genuinely interrogate assumptions [27]. In FSIs, this tendency is compounded by the distance between the teams that build causal models and the leaders who control budgets: a CMO navigating a quarterly board review, a CFO approving a £50 million media plan, and a CDAO managing model risk obligations are not, in general, practitioners who will open a model output file and reason through posterior distributions. If the insight does not reach them in a form they can act on, it does not reach them at all.

Agentic decision intelligence resolves the last-mile problem by placing a natural-language interface between the causal measurement stack and the people who use it. Rather than requiring stakeholders to navigate dashboards, a decision intelligence agent proactively surfaces the outputs that matter: weekly budget allocation recommendations grounded in the latest mROI estimates, scenario analyses that quantify the revenue impact of shifting spend between channels, alerts when a campaign’s observed incrementality diverges materially from the model’s prediction, and plain-language explanations of what changed in the model and why. Stakeholders can query the agent directly – «what is the marginal return on an additional £5 million in brand investment in the UK market?» – and receive a structured response drawn from the validated causal model, with confidence intervals and flagged caveats, rather than from a generative language model producing a plausible but ungrounded answer.

The critical distinction is grounding. An agentic decision intelligence system operating over a validated MMM or SCM is not a chatbot that approximates answers from training data – it is a retrieval and interpretation layer over a model that has been built, documented, and validated to the standards the organisation has committed to. This grounding matters particularly in the FSI context, where model outputs may influence material business decisions and where SR 11-7 and SS1/23 require that automated systems producing decision-relevant outputs are interpretable, documented, and subject to ongoing monitoring [28][31].



## Agentic Media Buying and Campaign Execution

The most advanced – and, in the FSI context, most carefully bounded – application of agentic AI in the measurement value chain is the direct connection between causal model outputs and media buying execution. In the conventional workflow, the mROI framework produces budget allocation recommendations that are then manually translated into instructions for agency trading desks or in-house programmatic teams. This translation introduces latency measured in days or weeks, and interpretation variance as recommendations pass through multiple handoffs. By the time optimised allocations reach the market, the conditions that justified them may have changed.

A tightly scoped agentic execution layer can close this gap. Within pre-approved operating parameters – channel budget ranges, bid floors and ceilings, audience targeting constraints – an agent can implement the reallocation recommendations of the causal model directly, adjusting bids in DSPs, pausing underperforming line items, increasing investment in channels where the model identifies unexploited incrementality headroom, and triggering the geo-holdout tests that feed the next model refresh. Decisions that fall outside the agent's authorised parameters – significant budget shifts, new channel activations, audience targeting changes – are escalated to human approvers with a structured recommendation and the model evidence supporting it.

For most FSIs, full automation of media buying decisions is neither appropriate nor necessary. The practical architecture is human-in-the-loop: agents handle the routine, high-cadence micro-optimisations that human trading desks cannot execute at sufficient speed or scale, while strategic allocation decisions remain with accountable human decision-makers. Compliance review requirements, the relatively extended consideration cycles of financial products, and the regulatory scrutiny that attaches to automated decisions in financial services all reinforce this model. What agentic execution changes is not who makes the decisions – it is how quickly the intelligence from causal models translates into campaign behaviour, and how consistently the organisation is able to act on the evidence it has invested in generating.

**Insight Agent**  
Ask any question in natural language to generate visualisations and insights

New Chat

Hello! I'm here to help you explore and understand your Marketing Mix Modeling (MMM) data. What insights can I help you uncover? For example, you can ask me to:


- Compare channel performance:** "Which channels have the highest ROI?"
- Analyze trends over time:** "Show me the weekly revenue and spend for the last quarter."
- Decompose your revenue:** "What were the biggest drivers of revenue last month?"
- Visualize diminishing returns:** "Plot the response curve for the Display channel."

Which channels have the highest ROI

Based on the latest model results, here are the channels with the highest Return on Investment (ROI). This chart shows the overall ROI for each media channel, giving you a clear picture of which channels are performing best.

The top-performing channels are Paid Search - Brand, Meta (Think), and Meta (See). Would you like to dive deeper into these results? For example, we could explore:

- ROI trends over time:** "Show me the monthly ROI for the top 3 channels."
- Spend vs. Contribution:** "Compare the spend share to the contribution share for each channel."
- Marginal ROI:** "Which channels are most efficient at the margin?"



Channel	ROI
YouTube	5.30
Meta (Do)	5.20
Primar	4.90
TikTok	4.70
Radio	4.00
Paid Search - Generic	3.90

**Budget Optimization**  
IMM-driven budget reallocation recommendations for next month

Upload Budget

**Optimisation opportunity identified**

**Next month: Rebalance media mix towards high marginal-return channels**

For the upcoming month, shift budget from saturated digital channels (Display, Paid Search, Video) towards high-performing offline channels (TV, Radio) and scalable prospecting formats (YouTube, Meta (See)) where the next € performs best.

Apply Recommendation

**Deep Research Agent**

Based on my research, these recommendations are broadly consistent with industry benchmarks.

- TV** is widely recognised as an underutilised channel in the industry, with strong marginal returns and slower saturation curves, making it a reliable source of incremental growth.
- Radio** is frequently cited as one of the most efficient channels, delivering strong reach, particularly when combined with TV for broad coverage.
- YouTube** aligns with benchmark findings as a scalable prospecting channel that extends reach—especially among younger audiences—and complements linear TV.

Conversely, **Display** and **Paid Search** are often considered mature, demand-capture channels in the industry, where high competition and audience saturation can lead to diminishing marginal returns, supporting the case for budget reallocation.

# 03. Marketing Maturity and Organisational Success.

The financial services industry faces significant hurdles in achieving measurement maturity. This is largely due to organizational silos once measurement tools are delivered. While it is critical to invest in the technical solution, it is equally imperative to invest in change management, ensuring that budget and campaign planning processes are optimized to leverage measurement results.

When it comes to digital marketing maturity, the average maturity score for banks has actually regressed in recent years, declining from 2.3 in 2021 to 1.9 in 2023 [24]. This regression is attributed to the difficulty of integrating data from various sources, such as digital platforms and physical branches, into a cohesive system.

## A. The ROI of Maturity

Artefact research shows that high-maturity brands are 3-4 times more responsive to market changes and achieve 1.5-2.5x higher profitability growth compared to their low-maturity peers

(based on Artefact client engagements across 3 major FSI institutions, 2020–2025). A top-5 European Asset Management firm that Artefact worked with, in internalising the Causal Marketing Measurement stack,

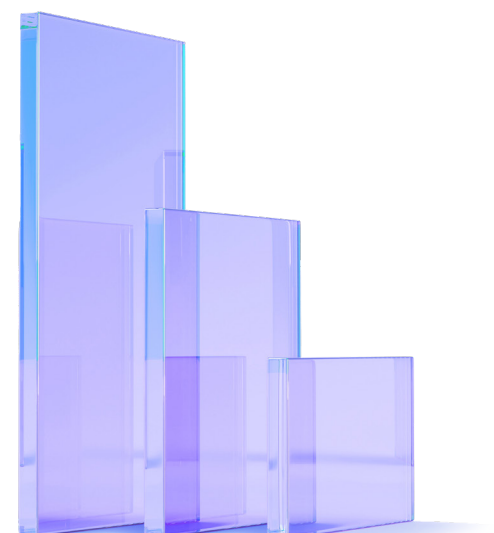
increased its Marketing Qualified Leads (MQL) by 20.3%, which was translated into a projected CLV over a five-year horizon.

## B. Measurement Without Adoption Is Worthless

Causal models, alone, do not generate value. Value is realized when the insights they produce impact decisions and teams are able to act on them – such as which channels receive higher budgets, which audiences get prioritised, and which campaigns get cut. In most financial service institutions, that translation from model output to a core decision or action is where the value is lost. The technical challenge of building a Bayesian MMM or a GP-augmented SCM model, in many aspects, is easier than the organisational challenge of persuading a CMO to reallocate budget away from divisions or channels that have always performed well on last-click metrics but demonstrably do not drive incremental revenue.

This isn't a new observation, but it continues to remain an unsolved one. McKinsey's research on large-scale organisational transformations consistently finds that approximately 70% of transformation programmes fail to meet their objectives – and that people and process account for the majority of that failure, not technology [26]. Gartner's research on analytics adoption compounds this: 65% of organisations still use data selectively to justify decisions they have already made, rather than allowing the data to genuinely inform what decision should be made in the first place [27].

For FSIs, both failure modes operate simultaneously – the transformation fails, and the data that does get produced gets reverse-engineered to fit the conclusion that was reached before the model ran.

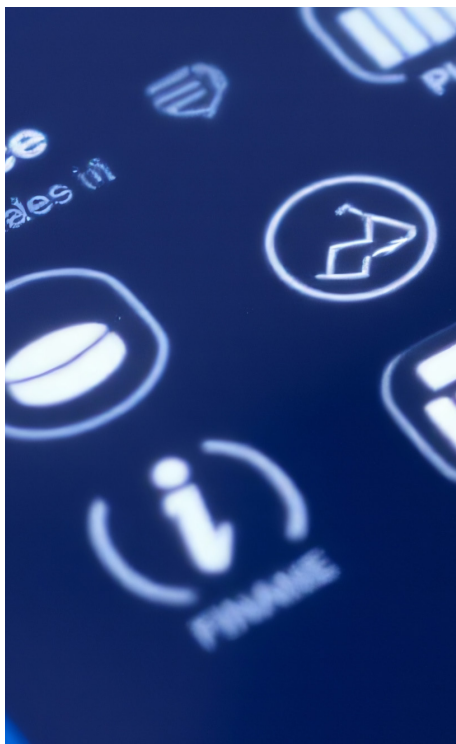


### The Three Structural Barriers to FSI

The adoption challenge in FSI is materially different from the same challenge in retail or consumer goods. Three structural dynamics create friction that does not exist in the same way outside of regulated industries:

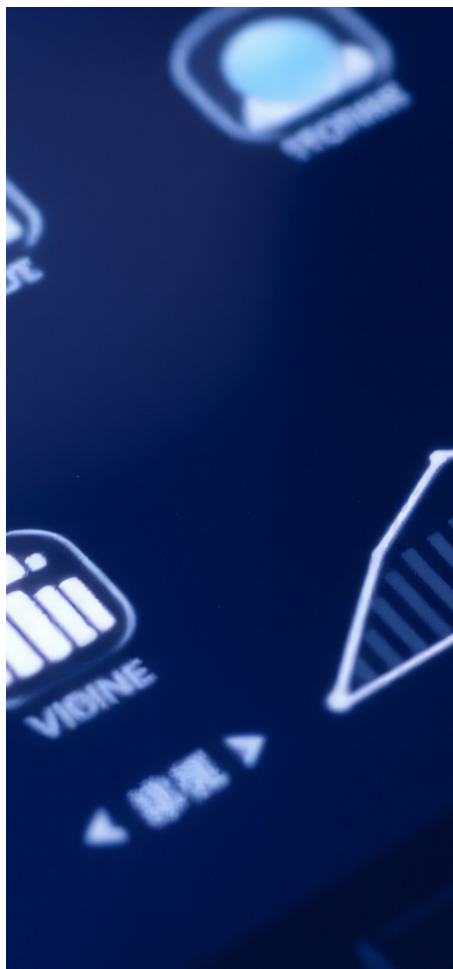
#### The marketing-finance cultural divide.

In most FSIs, marketing and finance operate on different planning cycles, accountability frameworks, and definitions of success. Marketing teams are accustomed to campaign-cycle thinking awareness or perception metrics, share of voice, click-through rates whilst finance teams operate on quarterly earnings logic, cost-of-acquisition ratios, and net present value. The result is a disconnect in the success metrics or KPIs that marketing teams are accountable for and have control over. A causal measurement design obligates these two functions into shared accountability for the same numbers: the same mROI estimate, the same incremental revenue figure and/or the same CLV projection. This is a governance change as much as a technical one, and it typically brings to the front, long-standing tensions around who owns the definition of “marketing and media effectiveness”, and who has the authority to act on it.



### The model risk management bottleneck.

FSI compliance requirements create friction at every point where a model output touches a customer-facing decision or a resource allocation with regulatory implications. Model Risk Management (MRM) teams – a function that rarely exists outside of FSI – will require validation of any causal framework before it can be used to justify resource allocation. SR 11-7 guidance from the US Federal Reserve, and PRA’s Model Risk Management Principles (SS1/23) published March 2023, requires that all models used for business decisions be subject to independent validation accompanied with detailed documentation of assumptions and continuous performance monitoring [28] [31]. For a causal MMM, this implies investing in interpretability and audit records from the outset, and not as an afterthought once the model is operationalised. Institutions that treat MRM as a final gate rather than a design input consistently see their measurement programmes stall at the point of deployment.



### The data silo problem.

The data required to run any of the causal models or experiments – media exposure and performance data, brand metrics, clean CRM records, validated KPI signals such as bank transaction data – typically sits across several different systems, owned by different teams or even external agencies, none of whom have historically had any incentive to align. Research by WBR Insights found that 54% of FSI leaders acknowledge data silos as a significant barrier to innovation, whilst the Mulesoft Connectivity Benchmark Report found that 88% of IT decision-makers in FSIs agree that data silos create operational challenges [29]. Without a deliberate data integration strategy that precedes the modelling work, marketing measurement becomes technically sound but empirically hollow.



### Operating Model for Measurement

One of the most consequential decisions an FSI institution must make when deploying marketing measurement solutions is an organisational one: where does this capability sit? Embedding it within data science or a CDAO function creates distance from the marketing planning cycles, and reduces the possibility of model outputs being acted upon in time to influence budget decisions. On the other hand, embedding it within finance, despite ensuring rigour, risks running measurement to serve cost control rather than growth.



Gartner's research on marketing Centres of Excellence (CoEs) identifies a federated model as the most effective structure for organisations attempting to scale analytics capabilities across multiple business units [30]. In the FSI context, this translates to a dedicated Measurement CoE or Accelerator that sits at the intersection of marketing, data science and finance – reporting jointly to the CMO and CDAO, with a formal interface to MRM. This structure is uncommon today, but it is the framework that differentiates institutions where causal measurement has driven material budget reallocation from those where it has produced impressive presentations that were not acted upon.

The CoE's primary mandate is not to run models. It is to ensure that model outputs enter the planning cycle at the point where decisions are still open, ensuring that model results are interpreted effectively and correctly. This means delivering mROI estimates and scenario analyses before annual budget commitments are made, not three or six months after. It means sitting in the room when channel or funnel budgets are discussed, not producing a report that circulates afterwards. And it means building the literacy – in the CMO's office, with the CFO's team, and with the brand and performance marketing teams – to read and challenge causal model outputs rather than simply receive them.

### Data Literacy as a Non-Negotiable Prerequisite

A marketing measurement programme will only operate at the speed of its least literate stakeholder. If a CMO cannot distinguish between a correlation and a causal estimate, the model will be used to confirm whatever the CMO already believed. If a CFO cannot read a confidence interval, they will either over-trust a point estimate or dismiss the entire framework as too uncertain to act on. If a marketing director cannot interpret an adstock decay or a saturation function, they will not understand why the model is recommending a budget cut to a channel that still shows positive last-touch ROI.



Investing in causal literacy is therefore not a soft priority it is a hard prerequisite for value generation. This does not mean training the CMO to build a Regression model. It means ensuring that senior stakeholders who receive causal measurement outputs understand what the model is claiming, what assumptions underlie it, and what the confidence bounds around the estimate mean in practical terms. Institutions that have built this literacy systematically through structured workshops, decision simulation exercises, and shared dashboards with layman interpretation layers consistently report faster adoption and stronger business impact from their measurement programmes than those that treat model communication as a data science documentation task.



### From Transformation to Action: Bridging to Implementation

The change management challenge described above is not a reason to delay building causal measurement capability. It is a reason to design the programme differently from the outset to treat organisational transformation as a parallel workstream, not a downstream consequence. The Core Recommendations that follow are structured precisely with this in mind. Each recommendation addresses not just what to measure or which technology to adopt, but the institutional conditions governance, data infrastructure, planning cycle integration, and compliance readiness without which even the most technically rigorous causal framework will fail to generate the business impact that justifies the investment.

# C. Core Recommendations for Institutional Implementation

The marketing-finance cultural divide, the model risk management bottleneck, and the data silo problem described above are not insurmountable. But they require deliberate institutional responses, not organic evolution.

The following recommendations are structured to address each barrier directly, whilst also providing a practical roadmap for FSIs at any stage of measurement maturity.

They are sequenced intentionally: the first three recommendations create the institutional conditions without which the remaining four will underdeliver.

01

## Build a Measurement Centre of Excellence

---

Establish a dedicated measurement function that reports jointly to the CMO and CDAO, with a formal interface to Model Risk Management. Its mandate should be defined in terms of decisions influenced, not models built – delivering mROI estimates and channel scenario analyses before budget commitments are made, not after campaigns have run. A central team of specialists that serves multiple product lines whilst maintaining methodological standards and translating outputs into planning-ready formats. The CoE should maintain a living record of decisions where causal model outputs demonstrably drove a different outcome than legacy measurement would have produced – that record is both its proof of value and its defence against budget cuts.

02

## Align Marketing and Finance Around Shared Accountability

---

Before any model is built, marketing and finance must agree on a common definition of what marketing is expected to contribute to business value and over what time horizon. In FSI, where customer lifetime value spans decades and brand trust is a structural competitive advantage, this imbalance is particularly costly. The practical mechanism for resolving it is a joint CMO/CFO scorecard that includes long-horizon metrics CLV, AUM inflows, product adoption rates alongside short-term conversion metrics such as net new account openings, and that both functions present to the board using the same numbers from the same measurement system.

03

## Treat Data Foundations as a Prerequisite

---

Audit existing marketing, CRM, and product conversion data pipelines for completeness, freshness, and cross-functional access rights before committing to a modelling programme. For institutions at low measurement maturity, this means establishing a consolidated marketing data environment with consistent definitions for media spend, customer identity, and conversion events before a single causal model is run. Causal modelling applied to fragmented, unreconciled data produces more sophisticated-looking errors, not better insights. The investment in data integration infrastructure is not glamorous, but it determines whether the measurement programme generates genuine evidence or expensive noise.

04

## Redefine How Marketing Proves Its Value

---

Replace click-through rates, impressions, and last-touch conversion counts as primary KPIs with metrics that connect directly to the income statement: incremental new accounts, CLV over a three-to-five year horizon, inflows and AUM growth, and for the most mature implementations contribution to brand equity and pricing power. These are metrics a CFO can evaluate in a capital-allocation meeting. Making this shift requires two commitments: the CMO must be willing to act on causal findings even when they are inconvenient particularly when a historically favoured channel is shown not to drive incremental revenue and the CFO and board must agree to evaluate marketing performance over a time horizon longer than a single fiscal quarter. Without the second commitment, the first is unsustainable.

05

## Invest in Always-On Experimentation

---

Build the operational infrastructure to run geo-holdout incrementality tests continuously, not as occasional one-off projects. Each experiment generates a causally identified data point that makes the next model run more accurately, creating a compounding measurement advantage. In regulated FSI environments where individual-level randomisation is constrained in credit and advisory product categories, geo-holdout designs in which matched geographic markets are randomised to different media treatments provide the most credible incrementality evidence available without triggering fair lending or fiduciary compliance concerns.

06

## Build Model Risk Management Readiness from Day One

---

Engage the Model Risk Management function as a design partner at the outset, not a compliance gate at the end. SR 11-7 guidance from the Federal Reserve requires independent validation, documented conceptual soundness, and ongoing performance monitoring for all models used in material business decisions [28]. UK institutions operate under equivalent obligations through the PRA's Supervisory Statement SS1/23 on Model Risk Management, published March 2023 [31]. An institution that builds a production-ready causal marketing model without MRM involvement will face a validation bottleneck precisely when the model is ready to influence budget allocation. The practical steps are: maintain a model inventory entry from the first version; document all assumptions and their statistical basis in language accessible to non-modelling validators; and build interpretability layers that expose parameter estimates, sensitivity ranges, and backtesting results. Institutions that invest in this infrastructure find that MRM validation improves model quality over time rather than obstructing it.

07

## Adopt Privacy-Enhancing Technologies

---

Integrate data clean room infrastructure into the measurement data architecture now, ahead of further tightening of third-party data access under GDPR, CCPA, and the ongoing deprecation of cross-site tracking. Clean rooms allow FSI institutions to combine first-party CRM data with media platform and open banking data in a privacy-compliant environment, producing richer model inputs without exposing individual customer records. Secure Multi-Party Computation extends this to encrypted joint analysis across multiple data-holding parties. Institutions that build this infrastructure now will maintain measurement continuity through the next wave of data privacy regulation; those that defer will rebuild their data environments under regulatory pressure.

08

## Leverage Agentic AI Across the Measurement Value Chain

---

Begin with the area that generates the highest immediate return and the lowest institutional risk: agentic data governance. Automating anomaly detection, schema validation, and PII compliance checks across media data feeds produces compounding model quality improvements without requiring changes to decision-making workflows or triggering additional regulatory review. Next, we can extend agentic orchestration to model execution – shortening the refresh cycle from weeks to days and freeing data science capacity for the design and interpretation work that genuinely requires human judgement. Decision intelligence agents, which translate causal model outputs into natural-language recommendations for CMO and CFO stakeholders, directly address the adoption problem. Media buying automation should be introduced last and with the most deliberate governance: agents that produce or act on outputs influencing material business decisions fall within the scope of SR 11-7 and SS1/23, and should be registered in the model inventory, subject to independent validation, and deployed within a human-in-the-loop architecture that preserves accountability for consequential decisions [28][31]. Institutions that sequence adoption in this order build agentic capability progressively, with each layer reinforcing the next.

# Conclusion:

## The Path to Marketing Value Optimisation.

Causal AI is not a new category of modelling or software. It is a new standard of evidence one that FSI marketing has needed for a long time and now has the methodological infrastructure to meet. By focusing on true cause-and-effect relationships, institutions can eliminate wasteful spending, optimise performance across the full funnel, and build a sustainable competitive advantage.

The measurement gap is closeable. The methodologies are available. The only remaining variable is whether institutions choose to act on that before their competitors do.

## About the Authors

If you would like to discuss this report, please contact the authors.



**Sid Mohan is the Senior Director Data Science** based in the Amsterdam/Utrecht office and the New York office of Artefact, and the global lead for the Marketing Measurement & Causal AI offering at Artefact. You may contact him by email at [siddharth.mohan@artefact.com](mailto:siddharth.mohan@artefact.com).



**Akhilesh Kale (AK) is the Partner** based in the New York office of Artefact, and is responsible for the Financial Services practice. You may contact him by email at [akhilesh.kale@artefact.com](mailto:akhilesh.kale@artefact.com).



**Melanie Ngo is the Consulting Director** based in the Los Angeles office of Artefact, and is the leading expert for Data and AI for Marketing. You may contact her by email at [melanie.ngo@artefact.com](mailto:melanie.ngo@artefact.com).

# References

1. Nielsen. (2024). Annual Marketing Report: Maximising ROI in a Fragmented World. <https://www.nielsen.com/insights/2024/maximizing-roi-in-a-fragmented-world-nielsen-annual-marketing-report/>
2. Broadridge Financial Solutions. (n.d.). Building a Better Marketing Engine for Asset Management. <https://www.broadridge.com/report/building-a-better-marketing-engine-for-asset-management>
3. eMarketer. (2024). US Ad Spending by Industry 2024. Insider Intelligence. <https://www.emarketer.com/content/us-ad-spending-by-industry-2024>
4. Banker, S., Dunfield, D., Huang, A., & Prelec, D. (2021). Neural mechanisms of credit card spending. *Scientific Reports*, 11, 4070. <https://doi.org/10.1038/s41598-021-83488-3>. Summary available at: <https://mitsloan.mit.edu/experts/how-credit-cards-activate-reward-center-our-brains-and-drive-spending>
5. Abdymomunov, A., Gerlach, J., & Sakurai, Y. (2023). Interest rate risk in the U.S. banking sector. Presented at the Federal Reserve Stress Testing Research Conference (FR-STRC), Federal Reserve Bank of Boston. <https://www.bostonfed.org/-/media/Documents/events/2023/FR-STRC/Interest-Rate-Risk-in-the-US-Banking-Sector.pdf> [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4395529](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4395529)
6. Cerulli Associates. (2025). Retail Channel Marketshare Gains Ground, Nearing Parity with Institutional. <https://www.cerulli.com/press-releases/retail-channel-marketshare-gains-ground-nearing-parity-with-institutional>
7. Cerulli Associates. (2024). Cerulli Anticipates \$124 Trillion in Wealth Will Transfer Through 2048. <https://www.cerulli.com/press-releases/cerulli-anticipates-124-trillion-in-wealth-will-transfer-through-2048>
8. Simon-Kucher. (2024). The New Economics of Neobanking. <https://www.simon-kucher.com/en/insights/new-economics-neobanking>
9. Distribution-valued Causal Machine Learning: Implications of Credit on Spending Patterns - arXiv, <https://arxiv.org/pdf/2509.03063>
10. Explaining the Interplay Between Merchant Acceptance and Consumer Adoption in Two-Sided Markets for Payment Methods - American Economic Association, <https://www.aeaweb.org/conference/2020/preliminary/paper/e3aH4yny> [https://www.bis.org/events/eopix\\_pw\\_prot/huynh\\_paper.pdf](https://www.bis.org/events/eopix_pw_prot/huynh_paper.pdf)
11. Complete Guide to Finance Marketing Regulations: What You Need to Know, <https://www.defianceanalytics.com/blog/complete-guide-to-finance-marketing-regulations-what-you-need-to-know>
12. Six Threats Demand a New Playbook for Banks in Wealth and Asset Management, <https://www.bain.com/insights/six-threats-demand-a-new-playbook-for-banks-in-wealth-and-asset-management/>
13. Navigating the challenges in the Asset Management industry—and leveraging their opportunities - Davies, <https://davies-group.com/consulting/blog/asset-management-challenges-opportunities/>
14. Jin, Y., Wang, Y., Sun, Y., Chan, D., & Koehler, J. (2017). Bayesian Methods for Media Mix Modelling with Carryover and Shape Effects. Google Research. <https://research.google/pubs/bayesian-methods-for-media-mix-modeling-with-carryover-and-shape-effects/>
15. Chan, D., & Perry, M. (2017). Challenges and Opportunities in Media Mix Modeling. Google Research. <https://research.google.com/pubs/archive/2d0395bc7d4d13ddedef54d744ba7748e8ba8dd1.pdf>
16. Binet, L., & Field, P. (2013). The Long and the Short of It: Balancing Short and Long-Term Marketing Strategies. Institute of Practitioners in Advertising (IPA). <https://ipa.co.uk/knowledge/publications-reports/the-long-and-the-short-of-it-balancing-short-and-long-term-marketing-strategies>

17. Sun, Y., Wang, Y., Jin, Y., Chan, D., & Koehler, J. (2017). Geo-level Bayesian Hierarchical Media Mix Modeling. Google Research. <https://research.google/pubs/pub46000/>
18. Nerlove, M., & Arrow, K.J. (1962). Optimal advertising policy under dynamic conditions. *Economica*, 29(114), 129–142. <https://www.jstor.org/stable/2551549>
19. Roberts, S., Osborne, M., Ebden, M., Reece, S., Gibson, N., & Aigrain, S. (2013). Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A*, 371(1984). <https://royalsocietypublishing.org/doi/10.1098/rsta.2011.0550>
20. Pearl, J., Glymour, M., & Jewell, N.P. (2016). *Causal Inference in Statistics: A Primer*. Wiley. <https://www.wiley.com/en-us/Causal+Inference+in+Statistics%3A+A+Primer-p-9781119186847>
21. Sharma, A., & Kiciman, E. (2020). DoWhy: An End-to-End Library for Causal Inference. arXiv preprint. <https://arxiv.org/abs/2011.04216>
22. Witty, S., Ha, J.K., Gretton, A., & Dhir, N. (2020). Causal Inference using Gaussian Processes with Structured Latent Confounders. *Proceedings of the 37th International Conference on Machine Learning (ICML)*. <https://proceedings.mlr.press/v119/witty20a/witty20a.pdf>
23. Peters, J., Janzing, D., & Schölkopf, B. (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press. <https://mitpress.mit.edu/9780262037310/elements-of-causal-inference/>
24. Callaway, B., & Sant’Anna, P.H.C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
25. Creating a Competitive Advantage Through Marketing Maturity in ..., <https://media-publications.bcg.com/Creating-a-Competitive-Advantage-Through-Marketing-Maturity-in-Banking.pdf>
26. de la Boutetière, H., Montagner, A., & Reich, A. (2018). Unlocking success in digital transformations. McKinsey & Company. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/unlocking-success-in-digital-transformations>
27. Gartner. (2024). Data & Analytics Summit: Key findings on analytics adoption. <https://www.gartner.com/en/newsroom/press-releases/2024-05-13-gartner-says-data-and-analytics-leaders-must-use-collective-intelligence-to-drive-business-value>
28. Board of Governors of the Federal Reserve System & Office of the Comptroller of the Currency. (2011). SR 11-7: Guidance on Model Risk Management. Federal Reserve. <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>
29. MuleSoft. (2024). Connectivity Benchmark Report. Salesforce. <https://www.mulesoft.com/lp/reports/connectivity-benchmark>; WBR Insights cited in: BAI Banking Strategies. (2023). Breaking down silos: The crucial role of data unification in financial institutions. <https://www.bai.org/banking-strategies/breaking-down-silos-the-crucial-role-of-data-unification-in-financial-institutions/>
30. Gartner. (2023). What makes a marketing centre of excellence? Gartner for Marketing Leaders. <https://www.gartner.com/en/marketing/insights/articles/what-makes-a-marketing-center-of-excellence>
31. Prudential Regulation Authority. (2023). SS1/23: Model risk management principles for banks. Bank of England. <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss>



### Data / AI Strategy & Transformation

- ^ Data / AI Strategy
- ^ AI Transformation
- ^ Data / AI Adoption



### Data / AI for Customer

- ^ Customer Experience
- ^ Marketing
- ^ Sales
- ^ Customer Service



### Data / AI for Operations

- ^ Office Operations
- ^ Supply Chain
- ^ Manufacturing & Interventions
- ^ Procurement
- ^ R&D



### Data / AI for Support functions

- ^ Finance
- ^ HR
- ^ Legal
- ^ Sustainability



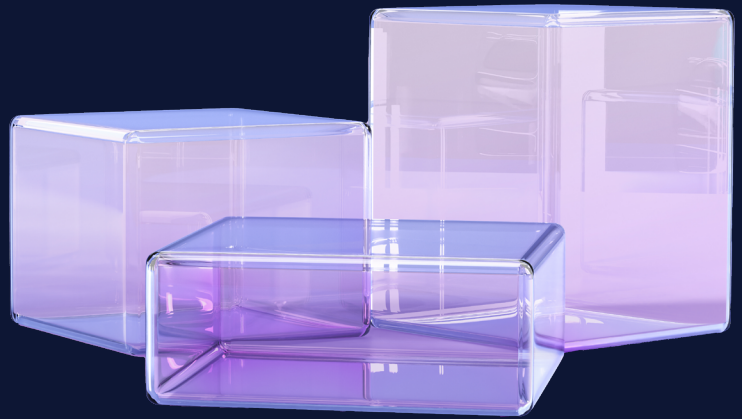
### IT & AI Technologies

- ^ IT strategy
- ^ AI for IT
- ^ AI Foundations
- ^ Analytics & BI
- ^ Data Governance & Management
- ^ Data Tech Foundations



# ARTEFACT

AI IS ABOUT PEOPLE



## CONTACT

hello@artefact.com  
+ 33 1 79 72 45 45  
artefact.com/contact-us

## ARTEFACT USA

135 W 26th Street  
New York,  
NY 10001  
United States  
hello-usa@artefact.com  
(805) 768-4730

## ARTEFACT UK

hello-uk@artefact.com  
+44 (0) 207 802 9500  
17th Floor, 5 Aldermanbury  
Square  
EC2V 7HR, London  
United Kingdom

## ARTEFACT BENELUX

marketing-nl@artefact.com  
Stationsplein 32 3511 ED,  
Utrecht  
Netherlands

## ABOUT ARTEFACT

Artefact is a leading global consulting company dedicated to accelerating the adoption of data and AI to positively impact people and organizations.

We specialize in data & AI transformation and data marketing to drive tangible business results across the entire enterprise value chain. Artefact offers the most comprehensive set of data-driven solutions per industry, built on deep data science and cutting-edge AI technologies, delivering AI projects at scale in all industry sectors.

From strategy to operations, we offer an end-to-end approach and solutions: data strategy, data quality and governance, data platforms, AI Factory, data-driven customer experience and marketing ROI. Our 2500+ employees operate in 27 countries (Europe, Asia, Americas, Middle East, Africa) and we partner with 1000+ clients, including some of the world's top 300 brands.

## ARTEFACT HEADQUARTERS

19, rue Richer  
75009 – Paris  
France

artefact.com