

GALILEO

Product Brief

Independent marketing mix model audits with trust-scored recommendations. Upload your spend and revenue data. Receive a credibility-certified report in minutes — with honest confidence ranges that tell you when to act and when to collect more data first.

Anjan Goswami | SmartInfer Inc. | May 2026

The Problem: Overconfident Recommendations

Most companies allocate marketing budgets based on numbers reported by the platforms selling them the media. Google reports Google's value. Meta reports Meta's value. Independent measurement studies consistently show these self-reported numbers overstate true incremental contribution, often by 40–80%.

This creates a capital allocation problem with four compounding risks:

- **Over-attribution from platforms.** Budget decisions based on biased inputs systematically over-invest in channels that look good on paper but may not deliver proportional incremental revenue.
- **Under-investment in demand generation.** Upper-funnel channels (TV, brand video, awareness) drive downstream search and consideration, but this effect is invisible in platform-attributed models.
- **False confidence from dashboards.** Platforms report precise numbers without acknowledging uncertainty. Decisions made on false precision carry hidden risk.
- **Undetected cross-channel effects.** Channels amplify or undermine each other. When these effects are unmeasured, credit flows to the wrong channels.

What Galileo Does

Galileo is a full marketing mix modeling system that adds trust and transparency at every stage. Upload weekly spend and revenue data, and receive a credibility-certified audit report with channel contributions, cross-channel interactions, budget optimization, and trust-scored recommendations — in minutes, not weeks.

Trust-Scored Recommendations

Every recommendation carries a computed confidence score — not a static label. When the model is strong and the scenario is well-supported, Galileo gives specific dollar-impact estimates. When confidence is low, it says so explicitly and suppresses unreliable numbers. This means Galileo will refuse to give you a precise dollar estimate when it knows the estimate is not reliable. That is a feature, not a limitation.

Multiple Estimation Methods

Galileo supports multiple estimation approaches — from fast statistical methods that deliver results in minutes, to full Bayesian techniques for deeper uncertainty quantification. The system selects the appropriate method based on data quality and client needs.

Structured Client Interview

An optional conversational interview captures what the marketing team already knows: lift study results, natural experiments, seasonality patterns, and confounding factors. This knowledge is validated against statistical plausibility checks and incorporated as model priors — so the model reflects domain expertise, not just the data.

Cross-Channel Interaction Detection

Most budget models treat channels as independent. In reality, TV drives branded search, social drives website visits, email primes retargeting. Galileo detects these interactions and quantifies their annual revenue impact. This is often the most valuable part of the audit.

Credibility Certificate

Every report includes an automated credibility assessment: 40+ quality gates covering data completeness, model diagnostics, and optimization consistency. This is what turns a model output into an audit.

What You Receive

Deliverable	What It Tells You
Channel contributions	How much revenue each channel drives (with confidence ranges)
Saturation analysis	Which channels have diminishing returns; where spend is wasted
Interaction map	Which channel combinations amplify or cannibalize each other
Budget optimization	How to reallocate the same budget for more revenue
What-if scenarios	Trust-scored impact of cutting or scaling specific channels
Pattern discovery	Hidden spending regimes, efficiency breakpoints, seasonal effects
Credibility certificate	Model quality assessment with pass/flag/fail on each dimension
Interactive report	Web-based dashboard with chat interface for follow-up questions
PDF report	Presentation-ready document for meetings and stakeholder review

How Trust Scoring Works

Every output — every reallocation recommendation, every counterfactual scenario, every chat answer — carries a computed confidence score derived from measurable inputs:

- **Model fit quality:** How well the model explains observed revenue
- **Confidence interval coverage:** Whether uncertainty estimates are calibrated
- **Spend shift magnitude:** How far the scenario departs from observed data
- **Data density:** How much spend history exists for the channels being changed
- **Credibility verdict:** The result of 40+ automated quality gates

Class	Score	What It Means
HIGH	≥ 0.70	Act on these recommendations. Model is strong, scenario is well-supported.
CAUTION	0.40–0.70	Directionally useful. Verify before committing budget.
LOW	< 0.40	Do not act. Dollar amounts suppressed. Collect more data first.

Illustrative Example

A DTC brand with \$18M annual ad spend across 6 channels received the following audit findings from a demo run:

- One upper-funnel channel was overweighted by ~3 percentage points with diminishing returns at current spend levels

- An owned channel had significant headroom — only ~35% saturated with strong marginal returns
- Cross-channel synergies were detected, including upper-funnel channels driving branded search
- Budget-neutral reallocation was estimated at 8–12% revenue uplift

Illustrative results from synthetic data. Real-world results vary depending on data quality, channel mix, and model assumptions.

How It Works

Step	What Happens	Time
1. Upload	Provide weekly spend + revenue CSV (26+ weeks)	Minutes
2. Interview	Optional conversational interview captures domain knowledge	5–10 min
3. Analysis	Automated analysis pipeline runs	< 1 min
4. Review	Interactive report, chat interface, PDF, and credibility certificate	Immediate

When It's Most Valuable

Strong fit: Brands spending \$5M–\$200M annually across multiple paid channels (upper + lower funnel). 26+ weeks of weekly data. Teams that have run lift studies or have domain knowledge to inform the model. Budget cycles where reallocation decisions are on the table.

Less useful: Single-channel advertisers. Fewer than 26 weeks of data. Constant spend across channels. Questions about creative optimization or audience targeting.

Commercial Positioning

	Platform	Open-Source	Galileo	Consultancy
Cost	Free (bundled)	Free + labor	Moderate (SaaS)	\$50–200K
Time	Real-time	Days–weeks	Minutes	8–16 weeks
Independence	No	Yes	Yes	Yes
Trust scoring	None	None	Every output scored	Consultant judgment
Credibility audit	None	None	40+ gates	Consultant judgment
Expertise needed	None	Data scientist	Analyst-level	Econometrician

Data Governance

- **Your data stays in your environment.** All data is stored within the Galileo instance. No external database, no third-party analytics, no data leaving your network.
- **Access isolation.** Each user can only access their own audit jobs. No cross-user data visibility.
- **Time-limited accounts with automated cleanup.** User accounts have explicit expiry dates. Expired data is removed through an auditable process with multiple automated safety checks.
- **No credentials in outputs.** Reports, PDFs, and exports contain only analytical results.

What Galileo Does Not Do

- **Not a replacement for experiments.** Lift tests and geo-holdout studies provide causal evidence that observational models cannot match.
- **Not individual-level attribution.** This is an aggregate time-series model. It answers “how much revenue did TV drive?” not “which users converted because of a TV ad.”
- **Not creative optimization.** Operates at channel-budget level, not ad-set or creative level.
- **Does not guarantee correctness.** All observational models are subject to omitted variable bias. The credibility certificate flags known risk factors but cannot rule out unknown confounders.

Engagement Structure

Element	Detail
Format	Fixed-fee pilot audit
Timeline	Results delivered within minutes of data upload
Commitment	Single engagement — no long-term contract required
Deliverable	Interactive report + credibility certificate + downloadable PDF
Follow-up	Optional quarterly re-runs as spend data accumulates

Getting Started

1. Prepare a weekly CSV with date, channel spend columns, and revenue
2. Upload through the web interface or API
3. Optionally complete a short conversational interview (5–10 minutes)
4. Receive interactive report + downloadable PDF + credibility certificate

For a pilot audit with your data: anjan@smartinfer.ai