Multi-Agent Causal Reasoning Framework: Optimizing Advertising Incrementality via Daily Budget Allocation Under a Fixed Lifetime Budget

Paul Lo, Erte Pan, Evelyn Wu, Jing Pan, Jingxi Yu, and Simon Hollands

Snap Inc., USA

Abstract. We study a practical approach for improving advertising incrementality by reallocating spend across days within a fixed lifetime budget throughout the campaign period. Our key observation is that the marginal gains in *incremental unique reach* (additional distinct users reached due to ads) vary systematically over time (e.g., weekly seasonality) and diminish as spend increases; therefore, a non-uniform daily budget split can increase total incremental unique reach and improve both incrementality outcomes and the statistical power of measurement. We propose the *Multi-Agent Causal Reasoning Framework*, which plans a multi-day daily-budget schedule under a fixed lifetime constraint using learned reach response models trained on past campaigns. The system comprises four agents:

- *Lifetime Budget Splitter (LBS), the Planning Agent:* generates candidate daily budget schedules.
- *Daily Reach Estimator (DRE), the Estimation Agent:* predicts daily incremental reach given campaign/advertiser context and candidate budgets.
- *Overall Reach Aggregator (ORA), the Governing Agent:* combines per-day reach predictions and rank schedules.
- *Outcome Simulation Module (OSM), the Simulation Agent:* simulates end-to-end incrementality performance improvements under explicit reach-to-outcome mapping assumptions.

In offline evaluation, the Daily Reach Estimator achieves $R^2 = 0.78$ with mean absolute percentage error (MAPE) of 18.8% on a held-out validation set. Under reach-to-CPiC mapping assumptions, a feasibility simulation suggests 63–88% CPiC reductions as reach scales, and a case study estimates a ∼60% CPiC improvement.

Keywords: Budget allocation · Reach response modeling · Incrementality · Multi-agent planning · Causal inference

1 Introduction

Most advertising campaigns operate under a fixed lifetime budget, with spend paced evenly across days. Advertisers, however, increasingly care about incrementality, the causal lift attributable to ads. The effectiveness of each additional

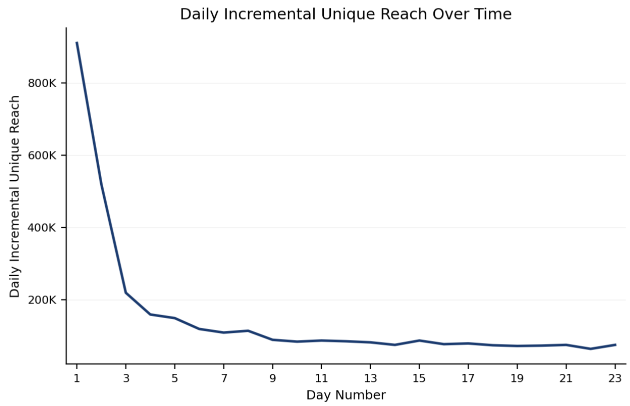

Fig. 1: Daily incremental unique reach decreases rapidly over time

dollar can vary substantially across days due to weekday seasonality, competition, user availability, and audience saturation. A pacing policy that splits spend uniformly is simple and stable, but it can leave meaningful incremental value on the table when marginal returns differ over time.

This paper studies improving incrementality by reallocating spend across days while keeping the lifetime budget fixed. We use incremental unique reach as a surrogate objective. Our key observations are:

– **Systematic time variation.** Weekdays and weekends can differ in user intent, inventory availability, and competitive pressure, leading to predictable differences in incremental opportunity across the week.
– **Diminishing returns.** For a given day, incremental reach typically saturates as spend increases (Fig. 1). Once a large fraction of the eligible audience has already been reached, additional spend increasingly goes to users already exposed to advertisement, reducing marginal incremental reach.

These properties imply that a deliberate non uniform daily budget schedule can increase total incremental reach over the period relative to a uniform split.

Beyond improving results, budget reallocation can also improve the reliability of incrementality measurement. Incrementality is typically measured through randomized experiments, and lift is easier to detect when the effect is larger and more users are exposed to ads over the campaign. This benefit is largest for advertisers with limited budgets or short campaigns, where concentrating spend on high-reach days produces more useful treatment exposure and tightens lift estimates.

This paper introduces the *Multi-Agent Causal Reasoning Framework*, an automated planning system that selects a multi-day daily-budget schedule under a fixed lifetime budget. The agent uses learned reach response models to allocate spend across days in order to increase total unique reach and, indirectly, improve incrementality measurement and cost efficiency for advertisers on Snap.

2 Preliminaries

2.1 Terminology

- **Incrementality (Lift) and Conversion Lift Study:** *Incrementality*, or *Lift*, is the causal increase in conversions attributable to advertising. It is measured via a *Conversion Lift Study*, a randomized experiment in which a *holdout group* (a randomly selected subset of eligible users) is withheld from ad serving for the campaign duration to provide a counterfactual baseline [1]. We use *Incrementality* and *Lift* interchangeably.
- **Cost Per Incremental Conversion (CPiC):** the average cost of driving one incremental conversion, defined as total campaign spend divided by the total number of incremental conversions generated by the treatment group relative to the holdout group. A lower CPiC indicates higher cost efficiency.
- **Incremental Unique Reach:** the number of additional distinct users reached by the campaign on day $N+1$ who were not previously reached during days 1 to N. This metric captures the campaign's ability to expand its cumulative audience by reaching new users over time.
- **Pacing:** the mechanism by which an ad platform distributes a campaign's budget over time, adjusting bid rates and delivery speed to smooth spend across the day or campaign period rather than exhausting it all at once.
- **Statistically Significant Positive (SSP) lift:** a binary outcome indicator equal to 1 when the estimated lift is positive and statistically significant ($p < 0.05$, $\hat{L} > 0$). Used as a secondary evaluation metric in Section 6.2.

2.2 Incrementality and Reach as a Surrogate Objective

Incrementality captures the causal impact of advertising on an outcome by comparing the observed result with the counterfactual outcome in which ads are not shown. Platforms typically measure incrementality through randomized experiments, in which eligible users are randomly assigned either to a treatment group that can be exposed to ads or to a holdout group that is prevented from seeing them. The holdout group provides the counterfactual baseline used to estimate the causal impact of advertising. In this paper, the primary surrogate objective is incremental unique reach, defined as the additional number of distinct users reached at least once due to advertising. We focus on reach because it responds predictably to budget, is typically easier to forecast than incrementality directly, is less sparse than conversions (i.e., many more users are reached than convert, making reach a more statistically stable signal), and often responds more smoothly to budget changes.

2.3 Connection to Measurement Power

Scheduling can also affect the quality of incrementality measurement. Experimental precision improves when treatment generates a larger effect and when the campaign accumulates more informative exposure. If certain days yield higher reach gains per dollar, concentrating spend on those days can increase the signal observed during the campaign, improving power to detect lift and reducing uncertainty, assuming the per exposure causal lift remains stable under the reallocated schedule. This is especially relevant for short campaign durations and moderate budgets where lift experiments are frequently underpowered.

2.4 Practical Constraints and the Need for Modular Reasoning

Real systems impose constraints beyond the lifetime budget (e.g., smooth pacing, feasibility constraints, and guardrails for delivery stability). The objective is also multi-day and non-linear because reach deduplicates across days and marginal returns saturate. These factors motivate modular reasoning: separating planning, causal estimation, and aggregation lets each component respect its own constraints, reduces error propagation, and allows targeted validation and iteration without reworking the entire system.

3 Related Work

We summarize prior work on budget planning and optimization in online advertising, organized by modeling methodology, optimization paradigm, and the level at which control is applied.

3.1 Time Series and Structural Models for Budget Allocation

Media mix modeling (MMM) and budget optimization often use time series structure to capture saturation and carryover. Bayesian MMM with adstock (a model of advertising carryover effects where past spend decays over time) and non-linear response curves share information across channels or regions while optimizing sales or revenue under budget constraints [2, 3]. These methods fit longer horizon, multi channel planning rather than day level pacing tied to experimental incrementality. Dynamic formulations for social media budget allocation model bidding and outcomes over time and allocate spend under reach and duration constraints [4], but they typically target observational outcomes and rely on strong structural assumptions.

3.2 Reinforcement Learning and Bandit Based Budget Planning

Budget allocation is often framed as sequential decision making. Alibaba's planner uses reinforcement learning to split budget across stages and coordinate with auto bidding, highlighting the value of learned stage-wise policies [5]. Amazon formulates cross-campaign allocation as a combinatorial bandit with knapsack constraints and Bayesian sharing [6]. These methods enable exploration but can be operationally complex and sensitive to delayed, nonstationary feedback.

Reach and frequency constraints have also been studied in related work. For example, Yahoo formulates guaranteed delivery as a reach allocation problem under impression supply and frequency requirements, solvable via quadratic programming [7]. These formulations highlight explicit constraint handling that complements our focus on daily budget scheduling.

3.3 Budget Allocation Versus Bidding Level Control

Prior work applies control at different layers of the ads stack. Some methods tune bids over time to meet reach goals or minimize spend [4, 8]. Others allocate budgets across channels or stages and rely on downstream bidding and pacing [3, 5, 6, 9]. Our work focuses on a high level daily budget schedule under a fixed lifetime budget, complementary to downstream bidding and frequency capping [10].

3.4 Positioning of Our Approach

Compared with the above literature, our setting emphasizes incrementality and experiment grounded signals. We plan day level schedules under a fixed lifetime budget while optimizing a reach based surrogate objective that is learned from historical incrementality experiments. This distinguishes our approach from observational response models and pure online reinforcement learning. Finally, our modular multi agent architecture separates schedule generation, causal reach prediction, multi day aggregation, and outcome translation, enabling component level validation and safer iteration under production constraints.

4 Problem Statement

Given an experiment running for N days with a fixed lifetime budget B, we seek a daily budget allocation vector $(b_1, \ldots, b_N)$ that maximizes the expected lifetime unique reach.

Formally, we formulate this task as a constrained optimization problem:

$$\begin{aligned} \text{Maximize} \quad & h(b_1, \ldots, b_N \mid c, t) \\ \text{Subject to} \quad & \sum_{i=1}^{N} b_i = B, \\ & b_i \geq 0, \ \forall i \in \{1, \ldots, N\}. \end{aligned}$$

Here, b_i represents the daily budget allocated to Day i, c represents specific campaign information (e.g., bidding strategy, optimization goal), and t denotes targeting information (e.g., device OS, geographic region). To reduce the search space, we reformulate the standard budget inequality constraint ($\leq B$) into an equality constraint ($= B$). This equality follows from the assumption that incremental reach is non-decreasing in spend: since additional spend cannot reduce reach, the budget constraint binds at the optimum.

This monotonicity property holds in the non-saturated regime (i.e., when the campaign has not yet reached a large fraction of its target audience) and is supported by two lines of evidence: (i) at the theoretical level, each additional dollar of daily spend continues to surface previously unreached users as long as the target audience is not exhausted, producing a non-decreasing reach curve with diminishing marginal returns—a pattern consistent with published platform research [3, 4]; and (ii) at the empirical level, internal budget-reallocation simulations confirm that shifting daily budgets toward high-reach days within a fixed lifetime budget consistently yields higher total unique reach than flat pacing.

4.1 Challenge 1: Resource Constrained Optimization

The primary challenge lies in optimally allocating a strictly finite resource, i.e., the fixed lifetime budget B, across multiple discrete time steps (days) to maximize a specific outcome. This resource allocation task is complex due to the non-linear, dynamic nature of advertising delivery.

First, the relationship between daily budget and reach is non-linear with diminishing returns. As spend rises, marginal new reach falls because impressions increasingly go to already reached users. Over-allocating a single day wastes budget on redundant impressions, while under-allocating can forfeit high-quality reach.

Second, the value of the resource fluctuates over time. The cost and availability of unique reach vary daily due to temporal seasonality (such as weekday versus weekend traffic patterns) and dynamic competitor auction density. Because the marketplace is constantly shifting, an expenditure of $1,000 today will not reliably yield the exact same reach as $1,000 tomorrow.

Finally, the allocation strategy must account for interacting platform constraints. Modifying daily budget limits directly alters the behavior of lower-level bidding execution. Pacing algorithms rely on dynamic "target remaining spend" goals to pace delivery; thus, changing the daily budget parameter alters the underlying bidding behavior in complex ways.

Because of these interacting constraints, fluctuating marketplace dynamics, and the constant threat of daily audience saturation, identifying the optimal budget vector is a highly non-trivial resource allocation problem that requires systematically balancing daily diminishing returns against the strict lifetime budget limit.

4.2 Challenge 2: Causal Understanding and the Surrogate Metric

The ultimate objective of the advertiser is not merely to reach users, but to maximize causal lift (e.g., incremental conversions). However, directly optimizing a bidding system or budget planner for incrementality is difficult due to measurement noise, delayed conversion windows, and the lack of full observability into counterfactual user behavior.

This requires a deep causal understanding of the advertising delivery mechanism to identify a reliable surrogate metric. We establish *Unique Reach* as this surrogate on the basis of a large-scale, propensity-score-matched meta-analysis across Conversion Lift studies, which consistently shows that lower reach is associated with degraded lift detection (Section 6.2). Expanding unique reach combats audience saturation (where delivering additional impressions to already-reached users yields near-zero incremental value) and improves the statistical power of the experiment.

Simulations calibrated to a real advertiser campaign (Section 6.1) confirm this pathway: CPiC falls by 63–88% as reach scales, while per-exposure lift remains stable—consistent with maximizing reach directly satisfying the causal objective.

To operationalize this insight, we frame lifetime budget planning as a causal decision problem: choose daily budgets to maximize the surrogate (Unique Reach) that reliably drives lift under real-world measurement constraints.

5 Multi-Agent Causal Reasoning Framework

We redefine the system as a Multi-Agent System where agency is fundamentally grounded in the comprehension of causal pathways:

$$\textit{Daily Budget} \rightarrow \textit{Unique Reach} \rightarrow \textit{Incrementality (Lift)}.$$

In this framework, specialized agents interact through coordination process to resolve the trade-offs between daily constraints and lifetime objectives.

5.1 The Causal Agency Architecture

Our system comprises a collective of four autonomous agents:

- **Lifetime Budget Splitter, the Planning Agent:** acts as the strategic lead. It leverages the causal assumption that non-uniform budget splits can unlock latent incrementality. It proposes eligible daily budget schedules based on historical priors.
- **Daily Reach Estimator, the Estimation Agent:** predicts the expected incremental reach for each candidate budget, providing the evidence needed to rank schedules.
- **Overall Reach Aggregator, the Governing Agent:** aggregates daily predicted reaches and selects the schedule that maximizes total estimated incremental unique reach across the campaign period.
- **Outcome Simulation Module, the Simulation Agent:** runs counterfactual simulations to quantify how the chosen budget schedule affects CPiC and statistical power under explicit reach-to-outcome mapping assumptions.

Figure 2 illustrates the overall architecture and the interaction between agents across the Planning and Execution layers.

- **Planning layer (agent):** chooses a daily budget schedule under a fixed lifetime budget.

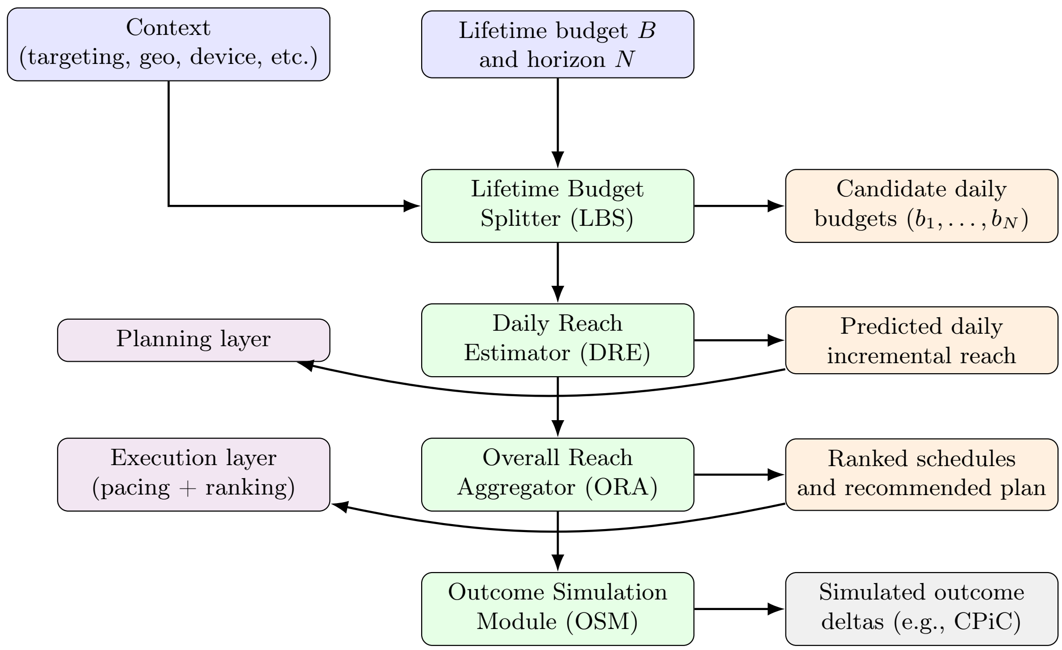

Fig. 2: Flow chart for the Causal Agency Architecture.

– **Execution layer (delivery stack):** the existing pacing algorithm and ranking models execute delivery given each day budget.

The proposed architecture is more than a relabeling of a predict-then-optimize pipeline. The decomposition separates four different reasoning roles: candidate generation under budget constraints, counterfactual reach estimation, global schedule ranking under delivery guardrails, and outcome simulation for incrementality interpretation. This separation is important in the production setting because each component has a different state/action space, validation contract, and safety constraint. It also allows the causal estimator and delivery guardrails to be audited independently from the search procedure. Moreover, this layered design mirrors existing advertiser workflows, where daily budgets are set upstream of the bidding and pacing systems. A flat single-component pipeline would entangle budget planning with the considerable complexity of real-time auction execution. The overall algorithm is summarized in Algorithm 1.

Algorithm 1: Budget Allocation Algorithm

```
Input: Number of days N; total lifetime budget B; campaign context
       (c, t); feasible split set 𝓑 = {(b_1, ..., b_N) | b_i ≥ 0, Σ_{i=1}^N b_i = B};
   DailyReachEstimator(b_i, i, c, t) → r_i;
OverallReachAggregator(r) → R.
Output: Best split b* = (b*_1, ..., b*_N) and maximal estimated reach R*.
Generate all possible daily budget splits: 𝓑 ←
 LifetimeBudgetSplitter(B, N).
R* ← −∞;   b* ← ∅.
foreach b = (b_1, ..., b_N) ∈ 𝓑 do
    r ← DailyReachEstimator(b)          // Get r_i for each day i
    R ← OverallReachAggregator(r)             // Combine daily
     reaches
    if R > R* then
        R* ← R;   b* ← b.
return (b*, R*).
```

5.2 Causal Reasoning and Surrogate Logic

Our agents operate on a specific Causal Pathway:

- **Direct Intervention:** the agents treat the daily budget as a controllable variable.
- **Surrogate Reasoning:** because direct incrementality signals are sparse and delayed, the agents adopt Unique Reach as a *causal surrogate* for Lift (empirically justified in Section 6.2). Two key assumptions underpin this choice: reach is monotonically non-decreasing in daily spend, and per-exposure causal lift remains approximately stable under budget reallocation.
- **Causal Explainability:** each recommended budget shift (e.g., increasing Monday's budget) is backed by a logical argument: "Increasing b_{mon} yields higher marginal reach than b_{wed} due to observed weekday seasonality, thereby improving overall unique user reached."

5.3 Data Processing and Features

Training data are derived from historical incrementality experiments conducted by a large-scale advertiser across multiple markets, and include both prediction targets and model inputs.

Targets. The primary target is daily incremental unique reach, defined as the count of distinct users reached on day d who had not been reached on any prior day within the same experiment (i.e., users whose first-exposure date equals d). Secondary targets include total daily unique reach and lifetime unique reach accumulated through day d.

Inputs. Model inputs comprise daily spend, daily budget, lifetime budget, experiment length, day index, and campaign targeting metadata including country, bid strategy, age group, gender, device/OS, and weekend indicator flags.

Dataset characteristics. The dataset spans over 100 experiments across multiple markets, yielding thousands of raw observations at the experiment-day level. After excluding days with zero spend or zero reach—which correspond to inactive or paused campaign periods—the majority of observations are retained for training. Lifetime budgets span several orders of magnitude with high variance, reflecting the highly skewed budget distribution typical of real advertising campaigns. Experiment durations range from 7 to 28 days. The dataset is dominated by purchase conversion events.

Due to privacy and data-sharing constraints, actual user-level records cannot be disclosed. The statistical parameters used to seed the simulation framework (Section 6) were derived from these real experiments, with noise added to the underlying user-level data prior to aggregation to protect individual privacy while preserving the distributional properties of real campaigns.

5.4 Daily Reach Estimation and Overall Reach Aggregator

The Daily Reach Estimator (DRE) agent acts as the system's predictive core, mapping campaign contexts and daily budgets to expected incremental unique reach. By accurately capturing the diminishing returns of marginal reach as spend increases, the DRE provides the essential causal modeling required for budget planning.

Because the system operates under a fixed lifetime budget, the available resources for each day are intrinsically linked, creating a zero-sum competition. The Overall Reach Aggregator (ORA) resolves this resource competition by synthesizing individual daily predictions into a unified global objective: $\mathrm{ORA}(\mathbf{r}) = \sum_{i=1}^{N} r_i$, which is valid without double-counting because each r_i is defined as the count of users first reached on day i (Section 5.3).

By leveraging these two specialized agents, we translate the theoretical framework described in the problem statement into a computationally feasible form:

$$\begin{aligned} \text{Maximize}_{\{b_i\}} \quad & \mathrm{ORA}(r_1, \ldots, r_N) \\ \text{s.t.} \quad & r_i = \mathrm{DRE}(b_i, i, c, t), \quad \forall i \\ & \sum_{i=1}^{N} b_i = B. \end{aligned}$$

For a fast prototype of the Daily Reach Estimation agent, we use gradient boosting decision trees, with mean absolute percentage error (MAPE) as the primary metric. In the first iteration, we train on observations across 14 countries and 40 experiments, achieving $R^2 = 0.78$ and MAPE = 18.8% on a 155-observation validation set. The model captures diminishing returns between daily spend and incremental daily reach.

5.5 Lifetime Budget Splitter

The splitter optimization chooses daily budgets that maximize predicted overall reach under a fixed lifetime budget. Because the objective function is unknown and may not be convex, classic closed-form methods are not directly applicable.

For prototyping, the splitter discretizes the total budget into units defined by a step size δ, which controls search granularity. Following a stars-and-bars logic [11], these budget units (stars) are partitioned into N daily allocations (bins) using $N-1$ dividers (bars), allowing the agent to systematically enumerate feasible distributions.

5.6 Guardrails

To avoid illogical allocations and preserve delivery system stability:

- **Splitter guardrails:** enforce strictly positive daily budgets and filter highly skewed splits using an entropy threshold.
- **Estimator guardrails:** cap predicted reach by an upper limit derived from eligible audience size (e.g., using recent monthly active users for the relevant geo and targeting).

6 Empirical evidence and evaluation

6.1 Feasibility evidence

We support two core claims using internal data and industry literature:

1. Increasing daily budget increases unique reach, albeit with diminishing returns.
2. Increasing unique reach improves the statistical confidence and cost efficiency of incrementality estimates.

To enable reproducible evaluation of the proposed framework, we employ a bootstrap simulation via the Outcome Simulation Module (OSM), seeded with statistical parameters derived from the real experiments described in Section 5.3. User-level signals are noised prior to parameter extraction to prevent re-identification, and the simulation regenerates synthetic campaign trajectories that preserve the distributional characteristics—budget scale, reach curves, and experiment duration—of the original data. Results should be interpreted as an evaluation of the methodology under realistic operating conditions rather than a direct replay of historical campaigns.

Using the OSM, a bootstrap simulation with 1,000 iterations evaluates how scaling reach affects statistical power and CPiC under explicit reach-to-CPiC mapping assumptions (Fig. 3). Here, *confidence* refers to the simulated probability that the z-test for lift crosses the threshold for statistical significance (i.e., simulated power at each reach level). Based on parameters from a historical experiment for an e-commerce advertiser, key findings include:

- A 10% reach increase first crosses the 80% confidence threshold.
- A 100% reach increase achieves 90% confidence, and a 200% reach increase is required to achieve 95% confidence.
- CPiC decreases substantially (by up to 88% at maximum scale) as reach increases, while the estimated incremental lift effect size remains stable at approximately 0.12%.

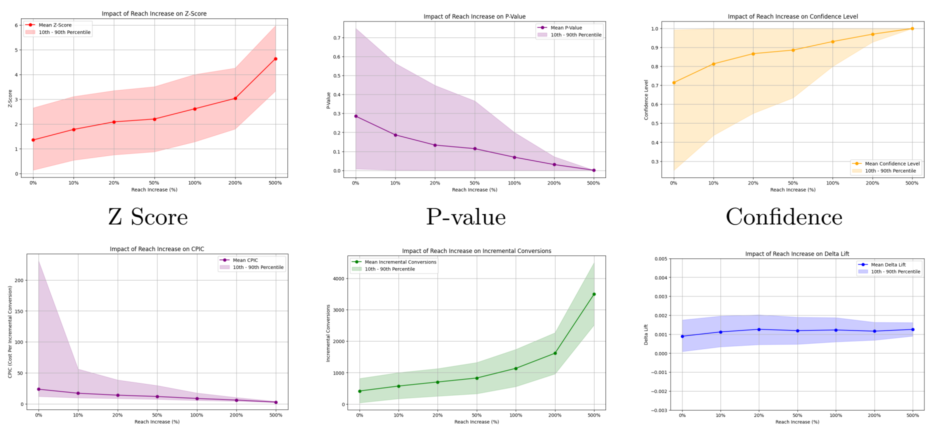

Z Score P-value Confidence

CPiC Incremental Conversions Delta Lift

Fig. 3: Impact of Reach Increase (mean and 10th–90th percentile bands).

6.2 Meta-Analysis: Reach and Lift

To ground the Reach surrogate empirically, we conducted a large-scale observational study across Conversion Lift experiments. Studies were classified as *low-reach* if their unique reach fell below the 25th percentile of the overall reach distribution, and as *high-reach* otherwise. The primary outcomes were (i) the continuous z-score summarizing lift magnitude and (ii) a binary indicator of a statistically significant positive (SSP) lift outcome ($p < 0.05$, $\hat{L} > 0$).

Because low reach is confounded with spend level, advertiser vertical, dominant bidding strategy, and study length, we applied propensity score matching (PSM) [12] using logistic regression with these covariates, followed by nearest-neighbor matching with replacement. Post-matching covariate balance was strong across all covariates.

The PSM results consistently show that low-reach studies underperform matched high-reach studies on both z-score and SSP rate, establishing a robust, confounder-adjusted association between insufficient reach and degraded lift performance. We note that PSM controls only for observed confounders and does not carry the identification guarantees of a randomized experiment; however, the magnitude and consistency of the effect across all adjustment schemes support treating Reach as a valid causal surrogate for Lift in the non-saturated regime.

6.3 Illustrative case study

An illustrative case study constructs a five-day planning task using historical campaign data from a representative market. The baseline ground truth utilizes a five-day lifetime budget of approximately $27K and yielded an overall reach of nearly 2 million users.

When tasked with optimizing the budget allocation, the agent recommends heavily front-loading the budget on the first day (allocating roughly $12K, more than double the baseline spend for that day) and distributing a lower, nearly flat budget of about $4K across each of the remaining four days. This optimized daily budget split yields a *predicted* overall reach of approximately 4 million users, representing roughly a 100% increase over the ground truth. This figure is a model estimate subject to the DRE's 18.8% MAPE; an online experiment is required to confirm realized gains. Under the reach-to-CPiC mapping produced by the simulation analysis (Fig. 3), this increase in reach corresponds to an estimated CPiC improvement of approximately 60%.

7 Discussion and limitations

While reach serves as a practical surrogate objective for campaign optimization, maximizing it does not guarantee maximized conversions or incremental value. However, broader campaign reach yields a crucial methodological advantage in experimentation: it expands the treatment groups in randomized controlled trials, providing the sample size and statistical power necessary to detect incremental effects.

Beyond objective formulation, our framework assumes the optimal execution of the underlying ranking, delivery, and bidding systems. In reality, extreme daily budget fluctuations can trigger unintended consequences, such as hitting minimum or maximum bid thresholds, or encountering inventory saturation under high-spend conditions. To mitigate these risks, our architecture relies on a strict separation of concerns: the multi-agent system determines daily allocations upstream, allowing the downstream pacing and ranking algorithms to independently manage intra-day auction execution.

Finally, several boundary constraints remain. The model's ability to generalize across diverse advertisers and geographies relies heavily on comprehensive data coverage and historical delivery stability. Additionally, highly skewed budget allocations risk introducing delivery artifacts—a vulnerability our system actively constrains through the implementation of entropy-based guardrails.

8 Conclusion

We propose *Multi-Agent Causal Reasoning Framework*, a hierarchical system that optimizes daily budget allocation under a fixed lifetime budget to maximize predicted unique reach and improve incrementality outcomes. A prototype Daily Reach Estimator achieves strong predictive performance, and an illustrative end-to-end case suggests meaningful potential gains in reach and CPiC. Future work includes improved sequential modeling, broader advertiser generalization, and tighter coupling with bidding objectives.

A key remaining step is controlled online validation. A prospective A/B test would enroll advertisers running live incrementality experiments and randomly assign campaigns to two arms: a *control arm* following standard equal daily budget pacing (business-as-usual), and a *treatment arm* whose daily budgets are set each morning by the Multi-Agent System via the campaign management API. Both arms would target identical audiences, maintain the same holdout assignment logic, and operate under the same lifetime budget. The primary evaluation metric would be CPiC at campaign end, with secondary metrics including total unique reach, experiment z-score, and SSP rate. A successful validation would require a statistically significant reduction in CPiC in the treatment arm with no degradation in lift effect size, confirming that the reach gains driven by the agent translate to genuine incrementality improvements rather than delivery artifacts.

References

1. Johnson, G.A., Lewis, R.A., Nubbemeyer, E.I.: Ghost Ads: Improving the Economics of Measuring Online Ad Effectiveness. *Journal of Marketing Research* **54**(6), 867–884 (2017). https://journals.sagepub.com/doi/10.1509/jmr.15.0297.
2. Google Research: Challenges and opportunities in media mix modeling. https://research.google.com/pubs/archive/2d0395bc7d4d13ddedef54d744ba7748e8ba8dd1.pdf.
3. Google Research: Bayesian Methods for Media Mix Modeling with Carryover and Shape Effects. https://research.google.com/pubs/archive/46001.pdf.
4. Luzon, Y., Pinchover, R., Khmelnitsky, E.: Dynamic budget allocation for social media advertising campaigns: optimization and learning. *European Journal of Operational Research* (2022). https://ideas.repec.org/a/eee/ejores/v299y2022i1p223-234.html.
5. ACM Digital Library: An adaptable budget planner for enhancing budget allocation. https://dl.acm.org/doi/10.1145/3690624.3709414.
6. Ge, L., Xu, Y., Chu, J., et al.: Multi Task Combinatorial Bandits for Budget Allocation. *arXiv preprint* (2024). https://arxiv.org/abs/2409.00561.
7. Hojjat, A., Turner, J., Cetintas, S., Yang, J.: A Unified Framework for the Scheduling of Guaranteed Targeted Display Advertising Under Reach and Frequency Requirements. *Operations Research* (2017). https://doi.org/10.1287/opre.2016.1567.
8. Gao, Y., Qiao, M.: Reach Measurement, Optimization and Frequency Capping in Targeted Online Advertising Under k Anonymity. *arXiv preprint* (2025). https://arxiv.org/abs/2501.04882.
9. Zhao, Q., et al.: A Unified Framework for Marketing Budget Allocation. *arXiv preprint arXiv:1902.01128* (2019). https://arxiv.org/abs/1902.01128.
10. IDEAS RePEc: The business impact of campaign setup: Reducing media waste. https://ideas.repec.org/a/aza/ama000/y2025v11i2p136-142.html.

11. CP Algorithms: Stars and bars. *Algorithms for Competitive Programming*. https://cp-algorithms.com/combinatorics/stars_and_bars.html.
12. Stuart, E.A.: Matching methods for causal inference: A review and a look forward. *Statistical Science* **25**(1), 1–21 (2010). https://doi.org/10.1214/09-STS313.