Ivitskiy, I.,
PhD, Associate Professor,
Founder, Doctor Ads LTD, London

УДК 659.113.71
DOI: https://doi.org/10.64076/ihrc251010.21

A PORTFOLIO OPTIMIZATION MODEL FOR MEDIA INVESTMENT UNDER INCREMENTALITY UNCERTAINTY

Advertising portfolios are increasingly selected under noisy estimates of causal lift. The practical objective is to allocate a fixed budget across channels to maximize expected incremental profit while controlling downside risk that arises from uncertain incrementality, cross-channel interference, and diminishing returns. This abstract proposes a two-layer model that merges posterior estimates of incremental return on ad spend with portfolio optimization and robust control ideas from operations research [1–3,7].

Suppose there are K channels, with spend vector x in R_+^K and total budget B. Let $S_i(x_i)$ denote the incremental revenue response for channel i as a function of spend x_i. For small perturbations around the operating point, define the local marginal incremental return r_i and quantify its uncertainty with a posterior mean $\mathrm{mu_i}$ and covariance matrix Sigma obtained from randomized geo experiments or matched-pair designs [2]. The portfolio objective is then written as:

$$\max_{x\geq 0} E\left[\sum_{i=1}^{K} S_i(x_i)\right] - \lambda \sqrt{Var\left[\sum_{i=1}^{K} S_i(x_i)\right]}\ s.t.\ \sum_{i=1}^{K} x_i = B, \tag{1}$$

where λ is a risk-aversion coefficient. The first term aggregates expected incremental revenue across channels; the second penalizes uncertainty via the portfolio variance that propagates from the uncertainty in r_i and cross-channel correlations encoded in Sigma. This structure generalizes mean–variance portfolio selection to advertising and inherits its efficient-frontier intuition [1].

To connect with real media behavior, $S_i(x_i)$ should be concave and saturating. In practice, S_i is estimated by a media mix model that jointly learns carryover and saturation effects, for example the open-source Robyn approach, which fits adstock and Hill-type curves and returns uncertainty for response parameters [5]. In the local region, a first-order approximation yields $S_i(x_i) \approx S_i(x_i^0) + r_i(x_i - x_i^0)$, which maps the objective above to a convex program in x whenever the variance term is quadratic in x. For larger reallocations, one solves the same program on a discretized grid of x_i

values, using the nonlinear S_i values with bootstrapped uncertainty bands from the mix model [3,5].

Incrementality estimates enter from experiments. Randomized paired geo designs and the Trimmed-Match estimator provide distribution-free iROAS posteriors that are robust when the number of geos is small and heterogeneous [2]. The optimizer uses posterior means to set mu and uses posterior covariance to populate Sigma. Where experiments are scarce, the model falls back to robust optimization by imposing an uncertainty set around mu and solving a worst-case problem:

$$\max_{x\geq 0}\min_{\mu\in\mathcal{U}} \mu^\top x - \lambda\sqrt{x^\top \Sigma x} \quad s.t. \sum_{i=1}^{K} x_i = B, \tag{2}$$

where $\mathcal{U}$ is a budgeted or fuzzy set calibrated from expert judgments and platform telemetry [7]. This protects against optimistic lift estimates without requiring full probability models.

The model also acknowledges platform mechanics. Portfolio bid strategies in auction platforms couple individual campaigns through shared constraints and pacing; treating them as one instrument inside the optimizer is realistic and avoids double counting of risk when budgets are linked [4]. The dynamic nature of campaigns leads to a learning loop: a small exploration fraction epsilon of budget is routed to policies that maximize information gain, for example contextual bandit allocation that balances exploration and exploitation and updates the posterior for mu, Sigma online [6]. This loop reduces uncertainty over time and moves the operating point along the efficient frontier as data accumulates.

Computation proceeds in four steps. First, estimate S_i and uncertainty with a modern MMM, calibrating to any experiment-based priors where available [5]. Second, run at least one geo-level incrementality test per major channel cluster to pin down posterior means and covariances for local marginal returns [2]. Third, solve the convex program for x with either the mean–variance objective or its robust counterpart when evidence is sparse [1,3,7]. Fourth, deploy a contextual bandit to manage the exploration budget and refresh mu, Sigma on a rolling window [6]. Quantities appear directly in the optimizer: the budget B is the equality constraint; the risk trade-off lambda is set by acceptable confidence shortfall in incremental profit; any channel-level minimums or platform pacing rules enter as linear constraints, for example $x_i \geq m_i$ for legal or brand safety reasons.

A useful diagnostic is the implied incremental return frontier. Let $G(x) = \sum_i S_i(x_i)$ and define a variance proxy $V(x) = x^\top \Sigma x$. Plotting $E[G(x)]$ against $\sqrt{V(x)}$

for the optimal x at different lambda values yields a curve of achievable trade-offs that decision makers can compare to internal cost of capital. In settings with severe left-tail risk, replace variance with conditional value at risk by solving

$$\max_{x \geq 0} E[G(x)] - \eta \cdot CVaR_{\alpha}\big(L(x)\big), \quad \sum_{i} x_i = B, \tag{3}$$

where $L(x)$ is the shortfall of incremental profit relative to a floor and eta is the penalty weight. This is compatible with robust sets that represent worst-case lifts within an uncertainty budget [7].

The approach remains general enough for the “Economics and Technologies” track. It ties causal evidence to optimization, accommodates platform coupling through shared budget instruments, and supplies a disciplined way to trade expected incremental profit against uncertainty rather than relying on platform-reported proxy metrics [2–6].

References

1. Markowitz H. Portfolio selection. *The Journal of Finance*, 1952, 7(1), 77–91. DOI: 10.1111/j.1540-6261.1952.tb01525.x.

2. Chen A., Au C. T. Robust causal inference for incremental return on ad spend with randomized paired geo experiments. *Annals of Applied Statistics*, 2022, 16(1), 1–20. DOI: 10.48550/arXiv.1908.02922.

3. Luzon Y., Pinchover R., Khmelnitsky E. Dynamic budget allocation for social media advertising campaigns: optimization and learning. *European Journal of Operational Research*, 2022, 299(1), 223–234. DOI: 10.1016/j.ejor.2021.08.019.

4. Ivitskiy, I. Mastering Google Ads Portfolio Bid Strategies. *Doctor Ads Blog*, n.d. Available at: https://blog.thedoctorads.com/mastering-google-ads-portfolio-bid-strategies (accessed: 17.09.2025).

5. Runge J., Skokan I., Zhou G., Pauwels K. Packaging Up Media Mix Modeling: An Introduction to Robyn’s Open-Source Approach. *arXiv*, 2025. DOI: 10.48550/arXiv.2403.14674.

6. Han B., Arndt C. Budget allocation as a multi-agent system of contextual and continuous bandits. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, Article 10.1145/3447548.3467124. Available at: https://dl.acm.org/doi/10.1145/3447548.3467124 (accessed: 17.09.2025).

7. Nehézová T. S., Kvasnička R., Brožová H., Hlavatý R., Stanislavská L. K. A robust optimization approach to budget optimization in online marketing campaigns. *Central European Journal of Operations Research*, 2025. DOI: 10.1007/s10100-025-00984-x.