International Journal of Current Sciences and Engineering Intelligence
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: *https://ijcsi.com*
Volume 01, Issue 03, November 2026
OPEN ACCESS
E-ISSN :

Profit Sensitive Coupon Targeting Under Contact and Inventory Constraints

Shreya Iyer[1], Vaishnavi Patel[2], Divya Iyer[3]
[1,2] *Students, RSR Rungta College of Engineering and Technology, Bhilai, India*
[3] *Assistant Professor, RSR Rungta College of Engineering and Technology, Bhilai, India*

Article Info

Article History:
Published: 14 Nov 2026

Publication Issue:
Volume 01, Issue 03
November 2026

Page Number:
1-5

Corresponding Author:
Shreya Iyer

Abstract:
Coupon campaigns are often evaluated by response rate or additional purchases, although neither measure necessarily identifies the most profitable targeting policy. This paper develops an analytical framework that combines incremental response, discount expenditure, contact cost, and an expected inventory constraint. A constructed population of 10,000 customers is divided into three segments with different purchase probabilities under treatment and no treatment. Five campaign policies are compared using the same assumptions. A response ranked policy generates an expected incremental loss of ₹114,000, while an uplift ranked policy generates ₹22,000. A policy based on incremental contribution produces ₹74,000 using the same contact capacity. The difference arises because discounts are paid to all treated purchasers, including customers who would have purchased without an offer. Sensitivity analysis examines available inventory, discount size, and errors in assumed treatment effects. No campaign was conducted and no customer records were collected for this study. The results demonstrate a transparent decision framework whose practical value depends on credible causal response estimates and complete cost accounting.

Keywords: Coupon targeting; uplift modeling; incremental profit; causal marketing; customer segmentation; inventory constraints; campaign optimization.

Introduction

Digital marketing systems can assign a purchase probability to each customer and use that score to select recipients of a promotion. This approach appears reasonable because likely purchasers generate visible sales. However, a customer who would purchase without an offer may simply receive a discount on an existing intention. A campaign can therefore record high attributed revenue while creating little additional demand or reducing contribution.

The relevant comparison is between outcomes under an offer and outcomes under no offer. Even this incremental perspective is incomplete if the objective is profit. Two customers can have similar incremental purchase probabilities but different baseline purchase probabilities, causing different expected discount costs. Contact expense and operational capacity introduce further differences between an attractive statistical score and an economically suitable action.

The research question is how a retailer should allocate a limited number of coupon contacts when customers differ in both baseline response and treatment sensitivity. The study also asks how an inventory limit changes the preferred allocation. A finite segment model is used so that every policy and calculation can be examined without relying on an opaque prediction system.

The contribution is an explicit bridge between causal response assumptions and a constrained campaign decision. The study does not estimate causal effects from observed data. Instead, it holds hypothetical response probabilities visible throughout the analysis and shows how different ranking rules translate those assumptions into different financial outcomes. This distinction is necessary because optimization cannot correct response estimates that are biased or unsupported.

Copyright © 2026 The Author(s) : This is an open access article under the CC BY license (http://creativecommons.org/licenses/by-nc/4.0/)

Detailed Literature Review

Rosenbaum and Rubin [1] establish the role of the propensity score in observational causal inference under specified assumptions. Their propensity score concerns treatment assignment, which differs from a customer's predicted probability of purchase. Rubin and Waterman [2] apply causal reasoning to marketing interventions and explain why predicting outcomes alone does not establish the effect of a marketing action.

Gutierrez and Gérardy [3] review uplift modeling as estimation of incremental treatment impact. Devriendt, Moldovan, and Verbeke [4] compare uplift approaches and discuss their evaluation. These works motivate comparing treated and untreated outcomes, while also showing that model evaluation should reflect the intended application. A campaign ranking is useful only when its evaluation corresponds to the decision being made.

Athey and Imbens [5] develop partitioning methods for heterogeneous causal effects. Wager and Athey [6] develop causal forest methods, and Künzel and colleagues [7] organize approaches for estimating conditional treatment effects through metalearners. These methods could support a future empirical version of the present framework. They are not implemented here, and their statistical guarantees are not transferred to the constructed segment probabilities.

Ascarza [8] demonstrates that targeting customers on predicted risk can differ from targeting them on responsiveness to an intervention. Although the application concerns retention, the underlying distinction is relevant to coupon allocation: an outcome prediction and a treatment effect answer different questions. The current paper uses purchase response rather than churn as its modeled outcome.

Gubela, Lessmann, and Jaroszewicz [9] examine revenue uplift and campaign profit decomposition. Athey and Wager [10] connect causal estimation with policy learning. Together, these contributions support evaluating treatment assignment through a decision objective. The present paper provides a small analytical example that makes discount costs, contact capacity, and expected demand constraints explicit, allowing the source of a policy's advantage to be traced through arithmetic.

Detailed Methodology & Framework

The constructed population contains 3,000 customers in segment L, 4,000 in segment P, and 3,000 in segment N. Their baseline purchase probabilities are 0.60, 0.10, and 0.02 respectively. Under a coupon offer, the probabilities become 0.64, 0.20, and 0.04. The corresponding uplifts are four, ten, and two percentage points. The descriptive segment labels indicate different response patterns, not individually observed customer types.

Each order contributes ₹400 before the promotion. Every treated customer who purchases is assumed to redeem a ₹100 coupon. Contacting a customer costs ₹2 regardless of purchase. Each customer can place at most one modeled order during the campaign period. Returns, repeat purchases, cross product effects, and longer term customer value are excluded from this short period objective.

Incremental contribution per contacted customer equals prediscount margin multiplied by uplift, minus coupon value multiplied by treated purchase probability, minus contact cost. The discount term includes customers whose purchases would have occurred without treatment. Applying this expression gives contributions of minus ₹50 for L, plus ₹18 for P, and plus ₹2 for N.

$$\tau s = p1s - p0s$$

$$gs = m\tau s - dp1s - c = (m - d)\tau s - dp0s - c$$

$$\max \Sigma s\, xsgs\,;\ \Sigma s\, xs \le 5000$$

$$2260 + \Sigma s\, xs\tau s \le 2700\,;\ 0 \le xs \le ns$$

The index s identifies a customer segment. p0s and p1s are purchase probabilities without and with treatment, τs is uplift, m is prediscount contribution, d is coupon value, c is contact cost, xs is the integer contact allocation, and ns is segment population.

The campaign can contact at most 5,000 customers. Baseline expected demand is 2,260 orders, and the main scenario limits expected total demand to 2,700 orders. Consequently, the allocation can generate at most 440 expected incremental orders. This is an expectation constraint, not a guarantee against realized stockouts. A service level inventory model would be needed to control a specified stockout probability.

The optimization chooses nonnegative integer contact counts within each segment's population and maximizes total incremental contribution. For this small scenario, the optimum can be established directly. Segment P provides the greatest positive contribution per contact and per incremental order. Contacting all of P leaves room for 1,000 contacts, which are assigned to N because its contribution is positive while L's contribution is negative.

Comparison policies include no offer, an offer to everyone, ranking by treated purchase probability, ranking by uplift, and ranking by incremental contribution subject to the constraints. The universal offer policy is retained as an illustrative unconstrained calculation and is identified as infeasible. It is not presented as a valid competitor under the same capacity limits.

Sensitivity analysis lowers the inventory allowance, increases coupon value while holding response probabilities fixed, and reduces the actual uplift of segment P. The fixed response assumption in the discount sensitivity isolates accounting effects; it does not claim that customer response remains unchanged when a real coupon becomes more valuable. The framework diagram places causal estimation before contribution calculation and constrained allocation.

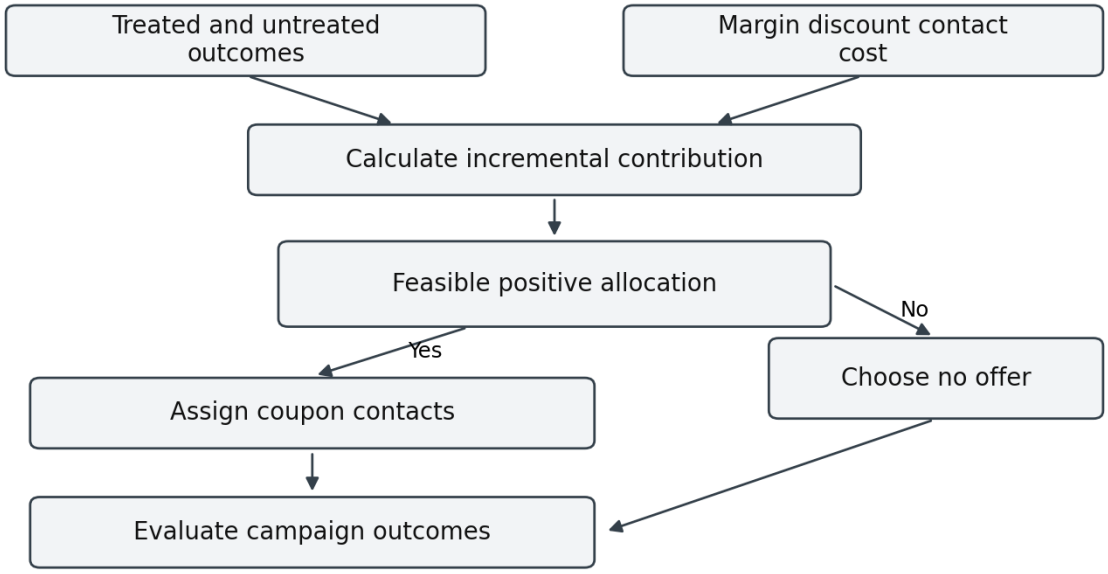


Figure 1. Connecting response estimates with campaign allocation

Detailed Results and Analysis

The response ranked policy contacts all 3,000 customers in L and 2,000 customers in P. It produces 320 expected incremental orders and 2,320 coupon redemptions. Incremental margin before promotion costs is ₹128,000, but discounts and contact costs total ₹242,000. The resulting incremental contribution is minus ₹114,000. A large number of coupon purchases therefore conceals an unfavorable comparison with making no offer.

The uplift ranked policy contacts all 4,000 customers in P and 1,000 in L. Its 440 incremental orders exactly use the expected inventory allowance. However, 1,440 expected redemptions reduce incremental contribution to ₹22,000. The contribution based allocation instead contacts P and 1,000 customers in N. It generates slightly fewer incremental orders, 420, but only 840 redemptions. Its expected incremental contribution is ₹74,000.

Table 1. Calculated campaign policy comparison

Policy	Contacts	Incremental orders	Total expected orders	Incremental contribution	Within constraints
No offer	0	0	2260	₹0	Yes
Offer everyone	10000	580	2840	−₹72000	No

Policy	Contacts	Incremental orders	Total expected orders	Incremental contribution	Within constraints
Response ranking	5000	320	2580	−₹114000	Yes
Uplift ranking	5000	440	2700	₹22000	Yes
Contribution based	5000	420	2680	₹74000	Yes

Note: Values are calculated from the stated analytical scenario.

The difference of ₹52,000 between these two feasible policies comes from the allocation of the final 1,000 contacts. Segment L provides greater uplift than N but loses ₹50 per contact, while N gains ₹2. This comparison explains why maximizing incremental orders can conflict with maximizing incremental contribution even when the same contact limit applies.

When expected inventory capacity falls to 2,500 orders, only 240 incremental orders can be accommodated. The preferred allocation contacts 2,400 customers in P and produces ₹43,200. Filling the entire contact allowance would be unnecessary. At a coupon value of ₹150 with response probabilities held fixed, contribution per P contact falls to ₹8, while N becomes economically neutral. Contacting P alone then yields ₹32,000.

Response estimation remains decisive. If P's actual uplift is three percentage points instead of ten, the original contribution based allocation produces a loss of ₹10,000. Recalculating the decision with that lower uplift would avoid P and favor the positive contribution available from N. Thus, the favorable main result demonstrates coherent optimization under assumptions, rather than evidence that an actual retailer will obtain the stated return.

Future Enhancement

An empirical extension should estimate response differences through a randomized coupon experiment with a clearly defined untreated comparison group. The outcome should include realized contribution after discounts, returns, and variable fulfillment costs. Evaluation on a later campaign or a reserved sample would help separate policy selection from policy assessment.

Further work can incorporate customer specific margins, several offer values, repeated contacts, and longer term effects such as delayed purchasing or dependence on discounts. Inventory can be represented through a probability of shortage rather than expected demand alone. Robust optimization could also restrict allocations when treatment effect estimates are uncertain. These extensions should retain the no offer option, since using all available campaign capacity is not necessarily economically justified.

Conclusion

The constructed comparison shows that high response, high uplift, and high incremental contribution can imply different coupon allocations. Discounting purchases that would occur anyway creates a cost that response based rankings can obscure. Under the stated assumptions, contribution based targeting generates ₹74,000, compared with ₹22,000 for uplift ranking and a loss for response ranking. The advantage is conditional on valid response probabilities and the specified accounting boundary. The framework provides a clear connection between causal marketing assumptions, operational constraints, and an economic decision, while identifying the estimation work required before practical use.

References

[1] Rosenbaum, P. R., and Rubin, D. B. (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika, 70(1), 41–55. Source

[2] Rubin, D. B., and Waterman, R. P. (2006). Estimating the Causal Effects of Marketing Interventions Using Propensity Score Methodology. Source

[3] Gutierrez, P., and Gérardy, J. Y. (2017). Causal Inference and Uplift Modelling: A Review of the Literature. Proceedings of Machine Learning Research, 67, 1–13. Source

[4] Devriendt, F., Moldovan, D., and Verbeke, W. (2018). A Literature Survey and Experimental Evaluation of the State-of-the-Art in Uplift Modeling: A Stepping Stone Toward the Development of Prescriptive Analytics. Big Data, 6(1), 13–41. Source

[5] Athey, S., and Imbens, G. (2016). Recursive Partitioning for Heterogeneous Causal Effects. PNAS, 113(27), 7353–7360. Source

[6] Wager, S., and Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests. Journal of the American Statistical Association, 113(523), 1228–1242. Source

[7] Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019). Metalearners for Estimating Heterogeneous Treatment Effects Using Machine Learning. PNAS, 116(10), 4156–4165. Source

[8] Ascarza, E. (2018). Retention Futility: Targeting High-Risk Customers Might Be Ineffective. Journal of Marketing Research, 55(1), 80–98. Source

[9] Gubela, R. M., Lessmann, S., and Jaroszewicz, S. (2020). Response Transformation and Profit Decomposition for Revenue Uplift Modeling. European Journal of Operational Research, 283(2), 647–661. Source

[10] Athey, S., and Wager, S. (2021). Policy Learning with Observational Data. Econometrica, 89(1), 133–161. Source