



Article

Incremental Net Contribution from Targeted Coupons: A Budget-Constrained Policy Evaluation Framework

Yu Yan ^{*,†} , Shuang Qiu ^{*,†}, Sangbum Son, Xuyang Duan and Jian Li

Department of Global Business, Yeungnam University, Gyeongsan 38541, Republic of Korea

* Correspondence: 22550178@yu.ac.kr (Y.Y.); qiushuang@yu.ac.kr (S.Q.)

[†] These authors contributed equally to this work.

Abstract

Targeted coupons are widely used in e-commerce, yet campaign evaluation often emphasizes conversion or gross merchandise value without distinguishing realized coupon cost from the monetary exposure created by coupon allocation. This study presents an applied policy-evaluation framework that combines established propensity-adjustment, uplift-ranking, and offline evaluation tools to compare coupon-allocation policies under recipient-coverage and full-face-value budget constraints. The analysis uses 8414 anonymized assignment records representing 8402 users within a focal campaign window. Because the platform's accounting margin is unavailable, a 20% margin is used as the baseline scenario for an incremental net-contribution proxy. Statistical uncertainty is assessed through 2000 user-cluster bootstrap replications and five fixed user-level data partitions. At 20% coverage, PA-RF produced a less negative proxy value than response-oriented policies, although its absolute confidence interval included zero and its ordering relative to PO-RF varied across data partitions. At 30% coverage, PA-RF produced a smaller net-contribution loss than PO-RF in all five partitions, while its absolute value and ROI proxy remained uncertain. The complete monetary-budget grid further shows that value-density allocation can differ materially from score-based allocation, that economically unattractive recipients may leave part of the available budget unused, and that expanding expenditure can reduce marginal subsidy efficiency. The contribution lies in a transparent and decision-oriented evaluation structure for interpreting coupon-allocation performance under explicit margin, cost, overlap, and model assumptions.

Keywords: targeted coupons; incremental net-contribution proxy; budget-constrained policy evaluation; subsidy efficiency; observational policy evaluation; user-cluster bootstrap; e-commerce promotion



Academic Editor: Ting Chi

Received: 13 June 2026

Revised: 23 July 2026

Accepted: 27 July 2026

Published: 29 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

1. Introduction

E-commerce promotion systems increasingly connect customer histories, offer attributes, assignment records, and transaction outcomes at the record level. This infrastructure enables more targeted allocation, but it also creates a measurement problem: a policy can increase conversion or gross merchandise value (GMV) while using subsidies inefficiently. Digital transformation, campaign prediction, personalized marketing, and dynamic resource-allocation research therefore motivate decision criteria that link customer response to resource use rather than treating response as the sole campaign objective [1–6].

Coupon allocation also operates in a setting shaped by consumer inference, privacy, perceived fairness, and platform design [7–9]. These considerations do not enter the present dataset as psychological mechanisms, but they reinforce the need for auditable policy rules. Managers need to know what a ranking score optimizes, what costs it recognizes, how it uses a finite budget, and where the evidence remains conditional.

The empirical challenge is especially pronounced when campaign assignment is observational. Propensity adjustment and doubly robust scores can address imbalance in observed pre-treatment covariates, but they cannot eliminate bias from unobserved assignment factors. Similarly, an assumed margin and observed coupon face value can support a transparent subsidy-efficiency proxy, but they do not recover platform accounting profit. Treating these distinctions explicitly is more informative than describing a complex model as an automatic profit optimizer.

This study addresses three research questions. RQ1 examines how response-oriented, customer-history, profit-proxy, propensity-adjusted, potential-outcome, and doubly robust ranking policies compare under fixed recipient coverage. RQ2 examines how policy performance changes when recipient-share coverage is replaced by a monetary cap based on full-face-value coupon exposure. RQ3 examines whether the principal policy comparisons remain stable across alternative user-level data partitions, statistical uncertainty estimates, proxy assumptions, and feature specifications.

The contribution of the study is an applied evaluation structure rather than a new causal estimator, uplift algorithm, or optimization method. It brings established estimation and policy-evaluation tools into a common decision setting in which the subsidy objective, realized coupon cost, full-face-value exposure, allocation rule, statistical uncertainty, and implementation requirements are explicitly distinguished. In particular, the framework separates recipient coverage from monetary expenditure, specifies a common value-density allocation rule, reports the complete monetary policy grid, and connects robustness evidence to practical coupon-budget decisions.

The empirical analysis identifies three managerially relevant patterns. First, conversion- and GMV-oriented rankings can remain poorly aligned with incremental net contribution after coupon cost is considered. Second, the relative performance of advanced policies depends on the allocation constraint and is not uniform across data partitions. Third, expanding monetary exposure can add recipients with low or negative estimated incremental value. These findings support a conditional and decision-specific interpretation of coupon allocation rather than a general claim of model dominance. Figure 1 summarizes the analytical framework.




Figure 1. Decision-oriented policy evaluation framework for budget-constrained targeted coupons.

2. Literature Review and Research Positioning

Coupon research shows that promotion effects extend beyond a binary redemption event. Exposure, redemption, post-promotion behavior, category targeting, and spillovers can all affect customer value and campaign cost [10,11]. Response-dependent-cost and profit-uplift studies consequently argue that a high predicted response is not necessarily an economically attractive treatment when the action itself is costly [12–15].

Uplift and heterogeneous-treatment-effect methods provide ranking tools for settings in which customers respond differently to an intervention [16–26]. These methods are useful here as established components, not as a claimed theoretical advance. Their role is to produce policy scores that can be compared under a common evaluation target and under explicit observational-identification assumptions.

Offline policy evaluation and empirical welfare research emphasize that a predictive model should be assessed through the decisions it induces [19,27,28]. Cost-sensitive and predict-then-optimize research similarly shows that downstream constraints can change the value of a prediction [27,29]. For coupons, this implies that selecting a fixed share of users and spending a fixed monetary budget are different operational problems whenever coupon face values vary.

The e-commerce environment adds implementation and governance considerations. Interface personalization, recommendations, customer journeys, digital advertising, mobile applications, and platform power can shape how algorithmic allocation is experienced and managed [30–46]. The present study does not model those mechanisms directly. It focuses on the narrower operational question of how anonymized coupon-assignment records can be used to compare allocation rules without overstating causal or accounting conclusions.

Taken together, these research streams position targeted couponing as an allocation problem in which heterogeneous customer response must be evaluated jointly with subsidy cost and operational constraints. The present study contributes an applied protocol that connects a clearly defined incremental net-contribution proxy, established policy scores, separate coverage and monetary constraints, user-level statistical inference, complete-grid reporting, and managerial implementation guidance. The objective is not to advance a new theory of e-commerce promotion, but to improve the transparency and decision relevance of coupon-policy evaluation.

3. Campaign Setting, Data Construction, and Identification

3.1. Campaign Background, Record Unit, and Observable Scope

The dataset contains 50,000 anonymized campaign-assignment records. The unit of observation is an individual coupon-assignment opportunity rather than a unique customer. Restricting the data to the focal campaign window yields 8414 records assigned between 12 June and 12 September 2025. These records correspond to 8402 hashed customer identifiers. Twelve customers appear twice, and five of them have different treatment statuses across their two records. All records belonging to the same customer are kept together during data partitioning and are treated as a single cluster in statistical resampling.

Treatment equals one when a record was assigned a coupon and zero for a control assignment. The analytical sample contains 4832 treated and 3582 control records. Observable pre-treatment and policy-time dimensions include new/existing-user status, recency, purchase frequency, 90-day monetary history, prior coupon use, prior coupon face value, average price band, product category, channel, device operating system, province label, city tier, coupon type, face value, threshold, and validity days.

Platform, merchant, and enterprise identities are anonymized. The dataset also does not contain customer age, gender, household characteristics, actual accounting margin, merchant reimbursement, acquisition expenditure, fulfillment cost, notification cost, customer-

fatigue measures, or long-term customer value. Geographic information is limited to the province and city-tier labels contained in the records. The empirical findings therefore apply most directly to campaign environments with comparable assignment processes, coupon structures, customer-history variables, and market composition. Table 1 summarizes the observable campaign, customer, coupon, and market characteristics of the analytical sample.

Table 1. Observable campaign, customer, coupon, and market characteristics.

Dimension	Observed Profile	Interpretation/Limitation
Record and user coverage	8414 assignment records; 8402 unique users	12 users have two records; 5 show treatment variation
Assignment window	12 June–12 September 2025	Observed assignment timestamps
Treatment composition	4832 treated (57.43%); 3582 control (42.57%)	Observational assignment
User tenure	4208 new; 4206 existing	No age or gender fields
Product categories	Electronics 1883; apparel 1490; home 1474; beauty 1281; grocery 1024; sports 843; baby 419	Observed category labels
Channels	App 5989; PC 1624; mini-app 801	Observed assignment channel
Device OS	Android 4643; iOS 2967; PC 804	Observed device label
Geography	Guangdong 1924; Jiangsu 1497; Zhejiang 1419; Shanghai 1212; Beijing 1086; Sichuan 764; Hubei 512	Province labels; country/enterprise identity unavailable
City tier	Tier 1: 2119; Tier 2: 2852; Tier 3: 2140; Tier 4: 1303	Observed market-tier label
Coupon type	Platform threshold 2946; store threshold 2083; platform direct 1294; store direct 1268; live-stream direct 823	Five observed types
Coupon face value	Mean 16.32; median 15; IQR 10–20; range 5–50	Observed nominal value
Coupon threshold	Observed for 5029 records; median 199; range 59–399	Structural missingness for direct coupons
Validity	Median 7 days; IQR 6–8; range 5–9	Observed coupon_valid_days
Customer history	Recency median 60 days; frequency median 0; 90-day monetary median 356.49; prior-use median 0.442	Pre-treatment fields
Average price band	Low 3791; medium 3303; high 1320	Categorical price-band field

3.2. Sample Construction and Campaign-Specific Analytical Sample

The analytical sample is defined by the focal campaign window. Of the 50,000 original records, 41,586 fall outside this window and are excluded, leaving 8414 campaign-period observations. Required-field, duplicate, and internal-consistency checks produce no material additional reduction. The restriction therefore establishes a common campaign environment rather than a population-representative sample of all records in the source dataset.

Figure 2 presents the sample-construction sequence and the treatment of repeated customers. The resulting observations form the campaign-specific sample used in all subsequent policy evaluations.




Figure 2. Sample construction and treatment of repeated users.

3.3. Outcomes, Proxy Construction, and Cost/Value Components

The observed behavioral outcomes are purchase conversion and transaction amount. The baseline realized net-contribution proxy is calculated as 0.20 times the observed transaction amount minus the coupon face value when an assigned coupon is redeemed. The 20% margin is a scenario assumption rather than an observed platform margin. Under the baseline specification, control observations and treated non-converters incur no realized coupon-redemption cost. The 20% baseline is not intended to represent a sector-wide e-commerce margin; appropriate rates may differ across product categories, fulfillment arrangements, and merchant–platform contracts.

Full-face-value exposure is conceptually and operationally different from realized redemption cost. Exposure equals the nominal face value of every selected coupon assignment and is used to define monetary budget caps. It is not deducted as though every selected coupon were redeemed. This separation allows coverage evaluation, realized-cost accounting, and budget reservation to be discussed without treating them as interchangeable.

The proxy is intended for relative policy comparison under explicitly stated assumptions and should not be interpreted as accounting profit. Alternative margin, cost-sharing, and fixed contact-cost specifications are examined in the sensitivity analyses. Logistics, customer acquisition, customer fatigue, budget-reservation opportunity cost, and long-term customer value remain outside the empirical measure because the required information is unavailable. Table 2 defines the variables and constructs used in the subsidy-efficiency evaluation.

Table 2. Variable definitions and subsidy-efficiency proxy construction.

Variable/Construct	Definition	Use
treatment	Coupon assignment indicator	1 = assigned coupon; 0 = control
outcome_binary	Observed conversion	Post-assignment outcome
outcome_amount	Observed transaction amount	Revenue-side input
coupon_face_value	Nominal coupon value	Realized-cost and budget-exposure input
full-face exposure	Face value for each selected assignment	Monetary budget unit
realized redemption cost	Face value $\times$ treated conversion	Baseline observed-cost rule
net-contribution proxy	$m \times \text{outcome_amount} - \text{realized redemption cost}$	$m = 0.20$ baseline scenario
policy value	Sum of user-level doubly robust incremental proxy over selected records	Conditional observational comparison

Table 3 summarizes the observed, assumed, sensitivity-tested, and unavailable cost and value components used in the evaluation.

Table 3. Cost and value components in the subsidy-efficiency evaluation.

Component	Observed	Baseline Proxy	Sensitivity	Coverage	Budget	Interpretation and Limitation
Baseline margin rate	No	Assumed: 20%	10% and 30%	Yes	Yes	Scenario parameter; not an observed margin
Realized coupon redemption cost	Partly: face value and conversion observed	Full face value for treated converters	50% cost-sharing scenario	Yes	Used in ROI proxy denominator	Observed-rule proxy; merchant reimbursement unavailable
Full-face-value budget exposure	Yes: coupon face value	Not deducted as realized cost	Score-only vs. density allocation	Reported	Budget constraint	Reservation exposure, not redemption cost
Budget-reservation cost	No	Excluded	Not evaluated	No	No	Operational opportunity cost remains unobserved
Merchant–platform cost sharing	No	Platform share set to 100%	50% cost-sharing scenario	Yes	Yes	Ranking/value can change with contract terms
Notification/contact cost	No	Excluded	Fixed 0.75 per selected record	Sensitivity analysis	Sensitivity analysis	Fixed scenario equal to 5% of the median face value; not estimated from observed operating costs
Logistics and fulfillment cost	No	Excluded	Not modeled	No	No	Order-level fulfillment data unavailable
Customer acquisition cost	No	Excluded	Not modeled	No	No	Acquisition source and spend unavailable
Customer fatigue	No	Excluded	Not modeled	No	No	Repeated-contact burden not measured
Long-term customer value	No	Excluded	Not modeled	No	No	No post-campaign CLV horizon

3.4. Covariate Balance, Propensity Adjustment, and Overlap

The treatment-control comparison is observational. In the primary test partition, the propensity model AUC is 0.552, indicating modest assignment predictability from available pre-treatment covariates. The common-support interval is [0.335, 0.760], and 0.119% of test records fall outside that interval. The largest unweighted absolute standardized mean difference among the reported numerical covariates is 0.249; after IPW weighting and 1st/99th-percentile weight clipping, the largest reported value is 0.075.

Propensity adjustment and doubly robust scoring address imbalance in observed pre-treatment covariates, but they do not eliminate bias from unobserved assignment factors. The resulting comparisons therefore remain conditional on ignorability, overlap, and model adequacy. Table 4 reports the covariate-balance and propensity-score diagnostics for the primary test partition.

Table 4. Covariate-balance and propensity diagnostics.

Covariate/Diagnostic	Unweighted Abs. SMD	IPW-Weighted Abs. SMD
is_new_user	0.249	0.050
r_days	0.199	0.058
f_30d	0.151	0.050
m_90d	0.223	0.075
prior_coupon_use_rate_90d	0.207	0.016
prior_coupon_facevalue_avg_90d	0.040	0.041
coupon_face_value	0.165	0.027
coupon_valid_days	0.032	0.073
Propensity AUC	0.552	—
Common support	[0.335, 0.760]	0.119% outside

Figure 3 visualizes covariate balance before and after IPW weighting and the overlap of the estimated propensity-score distributions between the treated and control groups.




Figure 3. Covariate balance and propensity-score overlap diagnostics. Note. In panel (A), the vertical dashed line marks the conventional absolute standardized mean difference threshold of 0.10; values to the left of the line indicate covariate balance within this threshold.

3.5. User-Level Train-Validation-Test Separation

The data are partitioned at the customer level into 60% training, 20% validation, and 20% test samples. All observations belonging to a repeated customer remain in the same partition. Within each random-forest policy family, the validation sample is used to select

between two predefined depth and minimum-leaf specifications. The resulting specification is then fixed before test-sample evaluation. Five fixed user-level partitions are used to assess the stability of the principal policy comparisons.

4. Policy Evaluation Methodology

4.1. Potential-Outcomes and Policy-Value Target

Let $Y_i(t)$ denote the transaction amount and $C_i(t)$ the realized coupon cost for assignment record i under treatment state t . With margin scenario m , the net-contribution proxy is:

$$V_i(t) = mY_i(t) - C_i(t) \quad (1)$$

The conditional incremental proxy is:

$$\tau_i(X_i) = E[V_i(1) - V_i(0)|X_i] \quad (2)$$

The resulting quantity is an adjusted conditional policy-value estimand under the stated ignorability, overlap, and model-adequacy assumptions. It is not an observed accounting-profit effect.

4.2. Propensity Adjustment and Doubly Robust Evaluation

The propensity score $e(X_i) = P(T_i = 1|X_i)$ is estimated from pre-treatment features. Separate treatment and control outcome models estimate $\mu_1(X_i)$ and $\mu_0(X_i)$. The test-record doubly robust pseudo-outcome is:

$$\varphi_i = \mu_1(X_i) - \mu_0(X_i) + T_i[V_i - \mu_1(X_i)]/e(X_i) - (1 - T_i)[V_i - \mu_0(X_i)]/[1 - e(X_i)] \quad (3)$$

Propensity scores are clipped to [0.05, 0.95] in the baseline specification. Cross-fitted pseudo-outcomes are used to train the propensity-adjusted and doubly robust ranking models. These procedures reduce sensitivity to imbalance in observed pre-treatment covariates, but they cannot remove bias arising from unobserved assignment factors. The resulting policy comparisons remain conditional on ignorability, overlap, and adequate nuisance-model specification [16,19,26].

4.3. Coverage-Constrained Ranking

For a coverage level c , each policy ranks test records by its fixed score and selects the first $\text{round}(cN)$ records. Coverage is a recipient-share constraint. Policy value is the sum of the doubly robust incremental proxy among selected records; expected redemption cost is the sum of predicted treated conversion probability multiplied by coupon face value; ROI proxy is policy value divided by expected redemption cost.

4.4. Monetary Budget Allocation Rule

Monetary caps are defined as 5%, 10%, 20%, and 30% of total full-face-value exposure in the test population. The principal implementation is a deterministic value-density greedy rule. For net-contribution policies, density equals predicted incremental net contribution divided by full-face exposure. Response-oriented baselines use their policy-specific response score per unit of exposure so that the same allocation mechanics can be applied as a comparator.

Candidates are ordered by descending value density, followed by descending policy score, ascending coupon face value, and a stable record order for deterministic tie breaking. Policies whose scores can be negative exclude observations with nonpositive predicted value, whereas nonnegative response probabilities and response predictions remain eligible. A positive-score observation with zero face value would be placed before positive-cost

observations; however, no zero-face-value coupons occur in the analytical sample. Observations are added sequentially until the next eligible assignment would exceed the budget cap. Any unspent amount is retained and reported through the budget-utilization rate.

The same allocation mechanics are applied to all policies. The deterministic greedy rule provides a scalable approximation suitable for campaign-level batch allocation. A score-based greedy rule is examined separately to assess how policy conclusions change when ranking does not account for coupon exposure per unit of predicted value. For N eligible records, the greedy allocation requires $O(N \log N)$ time for sorting, followed by a linear scan for budget filling.

4.5. Policy Baselines and Representative Objectives

The comparison set represents distinct managerial objectives and levels of analytical complexity. Random allocation provides a no-targeting reference; face-value and historical RFM rankings represent simple business rules; conversion and GMV models represent response-oriented targeting; treated-profit and T-learner models incorporate economic outcomes; two-model, propensity-adjusted, and doubly robust models estimate different forms of incremental value. Evaluating these policies with the same outcome definition, coverage levels, budget caps, and allocation mechanics makes it possible to distinguish differences in policy objective from differences in evaluation design.

Together, these baselines span the principal managerial objectives considered in the study—no targeting, simple business rules, response maximization, treated-outcome value, two-model uplift, propensity-adjusted value, and doubly robust value—under a common evaluation protocol. The comparison is therefore designed to examine objective alignment and operational implementation rather than to benchmark the frontier of causal-machine-learning estimators. Table 5 summarizes the estimation objectives, training samples, assignment adjustments, validation procedures, allocation rules, and managerial interpretations of the evaluated policies.

Table 5. Policy baselines, estimation objectives, and managerial interpretation.

Policy	Score/Objective	Training Sample	Assignment Adjustment	Validation	Coverage	Budget	Interpretation
Random	Random permutation score	None	None	Not required	Descending policy score	Score/exposure density greedy	Reference allocation
Face-value	Coupon face value (descending)	None	None	Not required	Descending policy score	Score/exposure density greedy	Generous-coupon rule
Historical RFM	Standardized recency-frequency-monetary-history score	Training records	None	Fixed formula	Descending policy score	Score/exposure density greedy	Customer-history rule
Conversion RF	Treated-only conversion probability	Treated training records	No	Validation-based selection over the fixed coverage grid	Descending policy score	Response-score density greedy	Response-oriented baseline
GMV RF	Treated-only expected transaction amount	Treated training records	No	Validation-based selection over the fixed coverage grid	Descending policy score	GMV-score density greedy	Sales-value baseline
Profit RF	Treated-only expected net-contribution proxy	Treated training records	No	Validation-based selection over the fixed coverage grid	Descending policy score	Value-density greedy	Profit-proxy baseline

Table 5. *Cont.*

Policy	Score/Objective	Training Sample	Assignment Adjustment	Validation	Coverage	Budget	Interpretation
T-learner conversion	Difference in treated/control conversion predictions	Separate treatment/control training models	Outcome-model contrast	Fixed specification	Descending policy score	Score-density greedy	Conversion uplift baseline
T-learner profit	Difference in ridge net-contribution predictions	Separate treatment/control training models	Outcome-model contrast	Fixed specification	Descending policy score	Value-density greedy	Linear profit-uplift baseline
Two-model RF profit	Difference in RF net-contribution predictions	Separate treatment/control training models	Outcome-model contrast	Validation-based selection over the fixed coverage grid	Descending policy score	Value-density greedy	Nonlinear two-model profit baseline
PA-RF net CATE	RF prediction of propensity-adjusted transformed outcome	All training records	Propensity-adjusted pseudo-outcome	Validation-based selection over the fixed coverage grid	Descending policy score	Value-density greedy	Adjusted incremental-value ranking
PO-RF treated profit	Treated potential net-contribution prediction	Treated training records	No direct treatment-effect contrast	Validation-based selection over the fixed coverage grid	Descending policy score	Value-density greedy	Treated-outcome value ranking
DR-RF profit	RF prediction of doubly robust pseudo-outcome	All training records	Outcome + propensity components	Validation-based selection over the fixed coverage grid	Descending policy score	Value-density greedy	Doubly robust incremental-value ranking

4.6. Validation and Model Freezing

Within each random-forest policy family, two predefined combinations of maximum depth and minimum leaf size are compared using mean validation-sample policy value at 10%, 20%, and 30% coverage. The selected family specification is fixed before test-sample evaluation. In the repeated-partition analysis, the same model specifications and evaluation grids are retained while the models are re-estimated on each alternative training partition.

4.7. Statistical Inference, Partition Stability, and Computational Assessment

Statistical uncertainty is evaluated using 2000 user-cluster bootstrap replications. Customers are sampled with replacement, and all assignment records belonging to a sampled customer enter the bootstrap sample together. The fitted models, test-sample scores, and allocation rules remain fixed, while policy membership, utilization, policy value, ROI proxy, and paired differences are recalculated in each resampled test population. Models are not re-estimated within the bootstrap.

Training-sample sensitivity is assessed across five fixed user-level data partitions. For 20% and 30% coverage and for 10% and 20% monetary budgets, the analysis reports the mean, standard deviation, and directional consistency of the principal policy comparisons. Training time, batch-scoring time, serialized model size, and allocation time are measured in the same computing environment and are interpreted as within-study implementation diagnostics.

Two-sided bootstrap *p*-values are calculated from the sign distribution of the paired differences using a finite-sample correction. The repeated-partition analysis is used only to assess training-sample sensitivity and does not inform post hoc policy selection. Table 6 summarizes the statistical inference and reproducibility design, including data partitioning, bootstrap resampling, allocation rules, statistical reporting, and computational assessment.

Table 6. Statistical inference and reproducibility design.

Protocol Item	Implementation
Unit of analysis	The unit of analysis is the assignment record. All records belonging to the same customer are retained in the same data partition and are resampled together as a user cluster.
Data partition	A user-level 60%/20%/20% training–validation–test partition is used. The validation sample is used for model selection, and the test sample is reserved for final policy evaluation.
Model selection	Within each random-forest policy family, two predefined model specifications are compared using mean validation-sample policy value at 10%, 20%, and 30% coverage. The selected specification is fixed before test-sample evaluation.
Repeated-partition analysis	Five user-level data partitions are fixed before final evaluation. The models are re-estimated in each partition using the same feature definitions, model specifications, coverage levels, budget levels, and evaluation rules.
Bootstrap resampling	Statistical uncertainty is evaluated using 2000 user-cluster bootstrap replications. Each sampled customer contributes all associated assignment records to the bootstrap population.
Bootstrap model treatment	Fitted models and test-sample policy scores remain fixed during bootstrap resampling. Policy membership, budget utilization, policy value, ROI proxy, and paired differences are recalculated, but the models are not re-estimated.
Coverage levels	Coverage-constrained evaluation uses recipient-share caps of 10%, 20%, and 30%.
Monetary budget levels	Monetary caps are defined as 5%, 10%, 20%, and 30% of total full-face-value coupon exposure in the test population.
Monetary allocation rule	Candidates are allocated using a deterministic value-density greedy rule. Records are added in descending value-density order until the next eligible record would exceed the applicable budget cap. Unused budget is permitted.
Eligibility rule	Records with nonpositive predicted incremental value are excluded from policies whose scores can be negative. Response probabilities and other nonnegative response scores remain eligible by construction.
Tie-breaking rule	Ties are resolved by descending value density, descending policy score, ascending coupon face value, and ascending stable record order.
Statistical reporting	The analysis reports point estimates, paired policy differences, 95% user-cluster bootstrap confidence intervals, two-sided bootstrap <i>p</i> -values, repeated-partition means, repeated-partition standard deviations, and directional consistency.
Contact-cost sensitivity	A fixed cost of 0.75 per selected assignment, equal to 5% of the median coupon face value, is evaluated as a sensitivity scenario rather than as an observed operating cost.
Computational assessment	Training-sample size, encoded feature count, model size, fitting time, batch-scoring time, and budget-allocation time are evaluated in the same computing session. Reported times are within-study implementation diagnostics rather than hardware-independent benchmarks.

5. Coverage-Constrained Policy Evaluation

5.1. Main Policy Grid

At 20% coverage, random, conversion-RF, and GMV-RF allocation produce net-contribution proxy values of −2143.21, −3380.90, and −4008.64, respectively. PA-RF produces a less negative value of −242.29, although its 95% user-cluster bootstrap interval [−860.46, 257.37] includes zero. PO-RF and DR-RF produce −822.13 and −728.49. The large negative values of the response-oriented policies are consistent with misalignment

between response prediction and subsidy efficiency, while the advanced policy estimates remain uncertain in absolute terms.

In the primary partition, the paired PA-RF–PO-RF difference at 20% coverage is 579.84, with a bootstrap interval of [123.12, 989.91]. Across the five user-level partitions, however, the direction reverses once, and the magnitude varies substantially. The 20% comparison is therefore interpreted as favorable within the primary partition but sensitive to the composition of the training and test samples.

At 30% coverage, PA-RF produces -365.41 , compared with -1099.52 for PO-RF. The paired difference is 734.11, with a bootstrap interval of [143.83, 1245.13], and the difference is positive in all five user-level partitions. This evidence indicates a smaller conditional net-contribution loss for PA-RF within the reported 30% grid. It does not indicate positive accounting profit or stable ROI proxy superiority. Table 7 reports the coverage-constrained policy results, user-cluster confidence intervals, ROI proxies, and repeated-partition stability measures. Figure 4 plots the coverage-constrained incremental net-contribution proxy curves for the principal policies over the evaluated coverage range.




Figure 4. Coverage-constrained incremental net-contribution proxy curves.

Table 7. Coverage-constrained policy results, user-cluster inference, and partition stability.

Coverage	Policy	Proxy Value	95% CI	ROI Proxy	Partition Mean	Partition SD	Share of Partitions with a Higher Proxy Value than Random
20%	Random	−2143.21	[−3124.26, −1232.03]	−0.757	−1632.12	472.00	Reference
20%	Conversion RF	−3380.90	[−4270.58, −2583.87]	−0.752	−2864.57	362.47	0%
20%	GMV RF	−4008.64	[−5090.28, −2987.79]	−0.684	−3352.07	539.96	0%
20%	PA-RF net CATE	−242.29	[−860.46, 257.37]	−0.280	85.72	201.01	100%
20%	PO-RF treated profit	−822.13	[−1468.79, −199.34]	−0.531	−293.22	345.07	100%
20%	DR-RF profit	−728.49	[−1300.60, −250.48]	−0.731	−77.66	400.72	100%

Table 7. Cont.

Coverage	Policy	Proxy Value	95% CI	ROI Proxy	Partition Mean	Partition SD	Share of Partitions with a Higher Proxy Value than Random
30%	Random	−2870.00	[−3919.73, −1877.70]	−0.680	−2447.70	398.75	Reference
30%	Conversion RF	−4733.71	[−5774.90, −3809.35]	−0.728	−4072.38	410.80	0%
30%	GMV RF	−5276.85	[−6401.08, −4158.32]	−0.679	−4359.50	555.80	0%
30%	PA-RF net CATE	−365.41	[−1031.78, 197.66]	−0.261	−70.31	177.22	100%
30%	PO-RF treated profit	−1099.52	[−1812.39, −378.87]	−0.474	−481.40	372.83	100%
30%	DR-RF profit	−617.74	[−1312.73, −48.94]	−0.399	−114.98	316.14	100%

Note. The percentage indicates the share of the five user-level partitions in which the policy produced a higher net-contribution-proxy value than random allocation; it does not represent statistical significance.

5.2. Paired Policy Comparisons

Figure 5 reports paired PA-RF–PO-RF differences under the user-cluster bootstrap. At 20% coverage, the primary-partition paired difference is 579.84 [123.12, 989.91] (two-sided bootstrap $p = 0.011$), but the ordering reverses in one of the five user-level partitions and the magnitude varies substantially. The 20% result is therefore treated as favorable primary-partition evidence rather than as a stable policy advantage. At 30% coverage, the paired difference is 734.11 [143.83, 1245.13] ($p = 0.013$) and remains positive across all five partitions; this finding is interpreted as a grid-specific reduction in net-contribution loss, not as positive accounting profit or stable ROI-proxy superiority. The differences under the 10% and 20% monetary caps remain statistically inconclusive. Taken together, the results indicate that the relative ordering is more stable under the reported coverage setting than under monetary allocation, while remaining conditional on the data partition and evaluation grid.




Figure 5. Paired user-cluster bootstrap differences for PA-RF versus PO-RF. Note. Points indicate paired policy-value differences; horizontal error bars indicate 95% user-cluster bootstrap confidence intervals; the vertical line at zero indicates no difference between the two policies. Positive values favor PA-RF, whereas negative values favor PO-RF.

6. Monetary Budget-Constrained Policy Simulation

6.1. Full Value-Density Policy Grid

Table 8 reports the complete set of evaluated policies at each monetary budget level. At the 5% cap, T-learner profit produces the highest net-contribution proxy (248.79), followed by random allocation (216.10) and two-model RF profit (203.24). At the 10%, 20%, and 30% caps, the T-learner profit value remains 282.29 because the policy selects only 324 positive-score observations and uses 82.7%, 41.4%, and 27.6% of the available budgets. Two-model RF profit similarly stops after 146 observations because no additional candidates have positive predicted incremental value.

Table 8. Full monetary budget-constrained policy evaluation grid.

Budget	Policy	Selected	Exposure	Utilization	Proxy Value	ROI Proxy
5%	Random	245	1410	99.5%	216.10	0.379
5%	Face-value	28	1400	98.8%	−348.00	−0.440
5%	Historical RFM	181	1415	99.8%	−764.68	−0.839
5%	Conversion RF	253	1410	99.5%	−115.17	−0.149
5%	GMV RF	206	1410	99.5%	−572.20	−0.645
5%	Profit RF	242	1415	99.8%	−534.47	−0.687
5%	T-learner conversion	243	1415	99.8%	−337.14	−0.473
5%	T-learner profit	220	1415	99.8%	248.79	0.679
5%	Two-model RF profit	146	870	61.4%	203.24	0.843
5%	PA-RF net CATE	200	1170	82.5%	−373.06	−0.723
5%	PO-RF treated profit	242	1415	99.8%	−534.47	−0.687
5%	DR-RF profit	191	1410	99.5%	−544.82	−0.989
10%	Random	413	2830	99.8%	−66.06	−0.054
10%	Face-value	56	2800	98.8%	−633.15	−0.395
10%	Historical RFM	285	2830	99.8%	−1033.25	−0.554
10%	Conversion RF	434	2830	99.8%	−551.57	−0.360
10%	GMV RF	348	2825	99.6%	−966.09	−0.555
10%	Profit RF	455	2830	99.8%	−318.50	−0.230
10%	T-learner conversion	411	2830	99.8%	−579.98	−0.406
10%	T-learner profit	324	2345	82.7%	282.29	0.422
10%	Two-model RF profit	146	870	30.7%	203.24	0.843
10%	PA-RF net CATE	200	1170	41.3%	−373.06	−0.723
10%	PO-RF treated profit	455	2830	99.8%	−318.50	−0.230
10%	DR-RF profit	223	1685	59.4%	−555.55	−0.858
20%	Random	640	5670	100.0%	−782.35	−0.313
20%	Face-value	138	5660	99.8%	−1640.03	−0.540
20%	Historical RFM	446	5660	99.8%	−2093.89	−0.572
20%	Conversion RF	708	5665	99.9%	−1091.79	−0.374
20%	GMV RF	566	5665	99.9%	−1591.63	−0.470
20%	Profit RF	699	5660	99.8%	−674.11	−0.245
20%	T-learner conversion	667	5670	100.0%	−1246.80	−0.435
20%	T-learner profit	324	2345	41.4%	282.29	0.422
20%	Two-model RF profit	146	870	15.3%	203.24	0.843
20%	PA-RF net CATE	200	1170	20.6%	−373.06	−0.723
20%	PO-RF treated profit	699	5660	99.8%	−674.11	−0.245
20%	DR-RF profit	223	1685	29.7%	−555.55	−0.858
30%	Random	820	8500	99.9%	−1379.00	−0.363
30%	Face-value	232	8480	99.7%	−3158.04	−0.685
30%	Historical RFM	576	8490	99.8%	−3261.03	−0.604
30%	Conversion RF	891	8505	100.0%	−2206.71	−0.492
30%	GMV RF	755	8500	99.9%	−2606.17	−0.523
30%	Profit RF	895	8505	100.0%	−1432.58	−0.357
30%	T-learner conversion	854	8490	99.8%	−2186.48	−0.510
30%	T-learner profit	324	2345	27.6%	282.29	0.422
30%	Two-model RF profit	146	870	10.2%	203.24	0.843
30%	PA-RF net CATE	200	1170	13.8%	−373.06	−0.723
30%	PO-RF treated profit	895	8505	100.0%	−1432.58	−0.357
30%	DR-RF profit	223	1685	19.8%	−555.55	−0.858

For random and response-oriented policies, higher caps add lower-ranked recipients and make the proxy increasingly negative. At the 20% cap, random, conversion RF, and GMV RF produce −782.35, −1091.79, and −1591.63. Value-density ranking improves these policies relative to score-only greedy allocation, particularly where high response scores are

attached to expensive coupons. For policies whose positive-score set is exhausted before the cap, the result remains unchanged while utilization falls.

No policy dominates simultaneously in net-contribution value, ROI proxy, selected-customer coverage, and budget utilization. T-learner profit produces the highest conditional value in the monetary grid, PA-RF performs most favorably in the coverage grid, and two-model RF profit reports a higher ROI proxy while using only a small share of the available budget. The managerial choice therefore depends on whether the objective emphasizes aggregate conditional value, return per expected redemption cost, recipient coverage, or the willingness to leave unattractive budget unspent. Figure 6 visualizes the monetary-budget-constrained net-contribution and ROI frontiers of the principal policies across the evaluated budget caps.




Figure 6. Monetary budget-constrained value-density policy frontier.

6.2. Score-Only Versus Value-Density Allocation

The comparison between score-based and value-density allocation demonstrates that the monetary implementation rule materially affects policy value. At the 20% cap, value-density allocation improves the proxy relative to score-based allocation by 1327.85 for random allocation, 2034.00 for conversion RF, 692.94 for GMV RF, and 170.55 for PO-RF. The two rules produce the same result when the set of eligible positive-score observations is exhausted before the budget cap or when the score and density rankings coincide. These findings support value density as the principal rule when coupon face value represents the binding monetary resource.

7. Stability, Computational Profile, and Robustness

7.1. Repeated-Partition Stability

Figure 7 reports the PA-RF–PO-RF net-contribution-proxy difference across five user-level data partitions under two coverage settings and two monetary-budget settings. The difference is positive in four of the five partitions at 20% coverage and in all five partitions at 30% coverage. Under the 10% monetary cap, the direction changes across partitions. Under the 20% monetary cap, the difference remains positive but varies substantially in magnitude. These results indicate that the relative ordering is more stable under the reported coverage comparisons than under the lower monetary-budget setting and that its magnitude remains sensitive to the composition of the training and test samples.




Figure 7. Stability of PA-RF versus PO-RF differences across five user-level data partitions.

7.2. Computational Profile

The primary training partition contains 5046 records and 52 encoded features. Fitting and validation of the nine learned policy families require approximately 8.17 s in aggregate, while batch scoring of all 12 policy rules on the test sample requires approximately 0.12 s. The 24 allocations formed by six principal policies and four monetary caps require approximately 0.012 s in total. Serialized learned-model sizes range from approximately 0.004 to 0.388 MB. These values are aggregate implementation diagnostics obtained in a common computing session rather than model-specific or hardware-independent benchmarks. The results indicate that the evaluation can be implemented through offline campaign-level batch processing [47]. Table 9 summarizes the aggregate fitting and validation time, batch-scoring time, budget-allocation time, and serialized-model size of the analytical pipeline.

Table 9. Aggregate computational profile of the analytical pipeline.

Computational Operation	Analytical Scale	Elapsed Time (s)	Interpretation
Fitting and validation of all learned policy families	5046 training records; 52 encoded features; 9 learned policy families	8.165	Aggregate elapsed time for model fitting and validation; not a model-specific benchmark
Batch scoring of all policy rules	1685 test records; 12 policy rules	0.117	Aggregate test-sample scoring time
Four budget allocations for six principal policies	24 policy-budget allocations	0.0118	Aggregate deterministic greedy-allocation time
Serialized learned-model size	Principal learned models	0.004–0.388 MB	Approximate serialized-size range

7.3. Profit-Proxy, Cost, Overlap, and Ablation Sensitivity

At 20% coverage, the PA-RF proxy changes from −242.29 under the 20% baseline margin to −570.78 under a 10% margin and 86.20 under a 30% margin. A 50% coupon-cost-sharing specification produces 249.68, while a fixed contact cost of 0.75 per selected observation reduces the baseline value to −495.04. The sign and magnitude of the absolute policy value are therefore sensitive to the assumed margin and allocation of coupon-related cost.

Removing the combined price and coupon-feature block reduces the PA-RF value by 368.23 at 20% coverage and 430.69 at 30% coverage. The effect differs across policies and monetary grids, indicating that coupon structure and customer price characteristics are economically important inputs to the ranking models. The ablation result should be interpreted as a dependence diagnostic rather than as evidence that a single feature family determines policy superiority [48,49].

Figure 8 summarizes policy differences relative to random allocation under alternative proxy and adjustment assumptions. Figure 9 reports aggregated feature importance together with the effects of removing the price and coupon-feature block. The alternative propensity-clipping specification produces the same reported values because the estimated propensity scores remain within both clipping ranges in the evaluated sample. The fixed contact-cost specification changes the absolute value of each coverage-20 policy by the same amount because all policies select the same number of records; consequently, their differences relative to random allocation remain unchanged. Table 10 summarizes which findings are robust across the evaluated settings, which are assumption-dependent, and which are not established by the available data.

Table 10. Summary of robustness findings and substantive interpretation.

Evidence Pattern	Finding	Interpretation
Robust across evaluated settings	Coverage and monetary budgets are not interchangeable	Recipient share and full-face exposure yield different selections and rankings
Robust across evaluated settings	Response-oriented policies can be economically misaligned	Conversion and GMV rankings are strongly negative in the main coverage and expanded-budget grids
Robust across evaluated settings	Budget expansion can reduce marginal subsidy efficiency	Random and response policies deteriorate as low-ranked recipients are added; positive-score policies often stop early
Assumption-dependent	Absolute net-contribution proxy	Changes with margin, cost-sharing, and fixed contact cost
Assumption-dependent	PA-RF versus PO-RF ordering	Favorable in primary coverage partition; 20% has one reversal across the five partitions; budget differences include zero
Assumption-dependent	Budget-grid ranking and utilization	T-learner profit ranks first by value, while other policies differ in ROI proxy and budget use
Assumption-dependent	Overlap and model adequacy	Observed adjustment cannot remove unobserved assignment factors
Not established by the data	True accounting profit gains	Margin, reimbursement, logistics, acquisition, fatigue, and CLV are not fully observed
Not established by the data	Universal model dominance	Rankings vary by coverage, budget rule, utilization criterion, partition, and proxy assumptions
Not established by the data	Unconstrained causal superiority or stable ROI proxy superiority	Evidence is observational, and ROI proxy remains unstable across key grids




Figure 8. Sensitivity of coverage-20 policy differences to alternative proxy and adjustment assumptions.




Figure 9. Aggregated feature importance and price/coupon-bundle ablation.

8. Discussion

8.1. Research Contribution and Positioning

The study contributes an applied policy-evaluation architecture for targeted coupon allocation. Its value lies in connecting established estimation tools to an explicit subsidy objective and to the operational distinction between recipient coverage and monetary exposure. The framework is therefore positioned as a contribution to methodological transparency and decision design rather than as a new causal estimator, uplift algorithm, or optimization method.

Relative to existing profit-uplift and offline policy-evaluation research [28], the framework extends applied practice in four respects. It distinguishes realized redemption cost

from full-face-value exposure and other unobserved operating costs; treats coverage and monetary constraints as separate empirical settings; integrates user-cluster inference, repeated user-level partitions, complete policy-grid reporting, and computational assessment; and recognizes unused budget as an informative outcome when additional candidates have nonpositive predicted incremental value.

The findings also refine the interpretation of response-oriented personalization. Customers with high purchase probability or expected GMV are not necessarily the most subsidy-efficient recipients when coupon values differ across assignments [50–52]. At the same time, the sensitivity of the profit-aware rankings to margin, sample partition, and coupon features shows why estimated incremental value must remain tied to the assumptions and operational setting under which it is calculated.

These findings extend evidence on coupon redemption, spillovers, and category targeting by showing that the operational conclusion changes when heterogeneous coupon exposure is treated as a constrained resource [10,11]. They are also consistent with response-dependent-cost and value-driven uplift research [12–15,50], which cautions against equating customer response with economic value. The separate coverage and monetary results further complement predict-then-optimize and marketing resource-allocation studies by demonstrating that a ranking score must be evaluated through the budget-feasible decisions it induces. In platform settings, this decision layer should remain auditable because personalization also raises privacy, fairness, and customer-experience considerations [7–9,45,46].

8.2. Managerial Implementation

Operational implementation requires a stable customer identifier, treatment assignment, policy-time customer history, coupon face value and structure, conversion, transaction amount, and a clearly defined campaign window. Accounting margin, merchant reimbursement, notification expenditure, fulfillment cost, and other operating-cost fields should be incorporated whenever they are available. When they are unavailable, the associated assumptions should be reported explicitly and examined through scenario analysis.

Customer-history variables may be refreshed through a daily or weekly batch process, while models may be recalibrated for each campaign or whenever coupon design, customer composition, assignment practice, or outcome calibration changes materially. Before campaign launch, eligible customers are scored, coupon restrictions are applied, full-face-value exposure is calculated, and the value-density allocation rule is applied to the approved monetary cap.

Allocation should stop when the next eligible observation would exceed the budget or when no positive predicted-value candidates remain. A campaign dashboard should report selected customers, full-face exposure, budget utilization, expected redemption cost, conditional net-contribution value, statistical uncertainty, score distributions, and changes in customer and coupon composition. Post-campaign outcomes should be used to evaluate calibration and inform subsequent campaigns.

Simpler allocation rules remain preferable when sample size is limited, propensity overlap is weak, customer-history variables cannot be refreshed reliably, the value proxy is highly uncertain, or the advantage of a learned policy is unstable across data partitions. Governance requirements should include documented feature definitions, model versioning, allocation-rule approval, access controls for customer identifiers, monitoring of repeated-contact burden, and an auditable record linking campaign assignments to the applicable eligibility and budget rules.

9. Conclusions

This study develops a decision-oriented framework for evaluating targeted coupon policies under explicit recipient-coverage and monetary-budget constraints. Rather than treating customer response as the sole criterion for campaign success, the framework evaluates allocation decisions through an incremental net-contribution proxy that distinguishes realized coupon cost from full-face-value budget exposure. It combines observational assignment adjustment, user-cluster statistical inference, repeated user-level data partitions, transparent allocation rules, and complete policy-grid reporting within a common evaluation structure.

Three findings provide the principal empirical synthesis. First, conversion- and GMV-oriented rankings can remain poorly aligned with subsidy efficiency after coupon cost is incorporated. Second, recipient coverage and monetary expenditure represent different operational constraints and can produce materially different customer selections and policy rankings. Third, expanding monetary exposure can reduce marginal subsidy efficiency, while policies that exclude nonpositive predicted values may appropriately leave part of the available budget unused. The relative performance of PA-RF, PO-RF, T-learner, and other policies therefore depends on the evaluation grid and does not support universal model dominance.

The study's theoretical contribution is a decision-oriented reframing of targeted coupon allocation in which heterogeneous customer response, subsidy cost, and budget exposure must be evaluated jointly rather than as separate campaign outcomes. Its methodological contribution lies in integrating established policy-learning tools with explicit cost boundaries, separate coverage and monetary constraints, user-level uncertainty assessment, and full-grid evaluation. For practitioners, the framework provides an auditable workflow for defining the economic objective, scoring eligible customers, allocating finite coupon budgets, monitoring budget utilization, and stopping allocation when additional recipients no longer provide acceptable estimated value. These contributions make the framework relevant both to researchers examining promotion-policy evaluation and to managers designing accountable coupon-allocation systems.

10. Limitations and Future Research

This study is subject to several limitations. The analysis relies on observational assignment records from a single anonymized campaign window. Although propensity adjustment, common-support diagnostics, inverse-probability weighting, doubly robust scoring, and user-cluster inference help address observed imbalance and statistical uncertainty, they cannot eliminate bias arising from unobserved assignment factors. The findings should therefore be interpreted as conditional policy-evaluation evidence rather than as unconstrained causal effects.

The incremental net-contribution measure also depends on assumed margin scenarios because the dataset does not contain platform accounting margins, merchant reimbursement arrangements, fulfillment costs, acquisition expenditure, notification costs, or other complete operating-cost information. The sensitivity analyses clarify how alternative margins, coupon-cost sharing, and contact-cost assumptions affect policy values, but they cannot substitute for observed accounting data. Long-term customer value, repeated-contact fatigue, post-campaign behavior, and budget-reservation opportunity costs likewise remain outside the empirical outcome.

The anonymized platform and merchant identities, the single campaign period, the observed market composition, and the available coupon and customer-history variables further limit the generalizability of the results. Policy rankings also vary with data partitions, margin assumptions, the treatment of nonpositive predicted values, and the allocation rule

used under monetary constraints. Application in other settings therefore requires local recalibration, overlap assessment, completion of relevant cost components, and validation under the applicable campaign and budget rules.

Future research could validate the framework using multi-campaign and multi-platform datasets covering different product categories, market environments, coupon structures, and merchant–platform contracts. Incorporating observed accounting margins, contractual cost sharing, fulfillment expenditure, and customer lifetime value would permit a more complete evaluation of both short- and long-term economic outcomes. Dynamic or real-time extensions could examine sequential coupon allocation, changing customer states, budget pacing, and repeated-contact effects during active campaigns. Stronger causal designs, including randomized field experiments, natural experiments, instrumental-variable approaches, causal forests, and alternative doubly robust policy-learning procedures, could further assess treatment heterogeneity and policy value. The framework may also be extended to other digital marketing decisions, including personalized discounts, advertising exposure, recommendation incentives, retention offers, loyalty rewards, and cross-channel promotion allocation.

Author Contributions: Conceptualization, Y.Y., S.Q. and S.S.; methodology, Y.Y., S.Q. and S.S.; software, Y.Y.; validation, S.Q., S.S. and X.D.; formal analysis, Y.Y. and S.Q.; investigation, Y.Y. and X.D.; resources, S.S.; data curation, Y.Y. and X.D.; writing—original draft preparation, Y.Y. and S.Q.; writing—review and editing, Y.Y., S.Q., S.S., X.D. and J.L.; visualization, Y.Y., X.D. and J.L.; supervision, S.S.; project administration, S.Q. and J.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: The study was conducted in accordance with the Declaration of Helsinki, and approved by the Institutional Review Board of Yeungnam University (Project No. 2026-R-0512; IRB Exemption Confirmation No. 7002016-E-2026-106; 18 May 2026).

Informed Consent Statement: Informed consent was not required because the study involved secondary analysis of anonymized campaign records and was determined to be exempt from the Institutional Review Board of Yeungnam University.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to privacy, confidentiality, and data-use authorization restrictions. Requests for access to the datasets should be directed to the corresponding author and will be considered in accordance with the applicable ethical, legal, and institutional requirements.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (OpenAI, GPT-5 series) for language editing, structural refinement, and assistance with code and figure formatting. The authors reviewed and edited all outputs and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Verhoef, P.C.; Broekhuizen, T.; Bart, Y.; Bhattacharya, A.; Qi Dong, J.; Fabian, N.; Haenlein, M. Digital Transformation: A Multidisciplinary Reflection and Research Agenda. *J. Bus. Res.* **2021**, *122*, 889–901. [\[CrossRef\]](#)
2. Dwivedi, Y.K.; Ismagilova, E.; Hughes, D.L.; Carlson, J.; Filieri, R.; Jacobson, J.; Jain, V.; Karjaluoto, H.; Kefi, H.; Krishen, A.S.; et al. Setting the Future of Digital and Social Media Marketing Research: Perspectives and Research Propositions. *Int. J. Inf. Manag.* **2021**, *59*, 102168. [\[CrossRef\]](#)
3. Riegger, A.-S.; Klein, J.F.; Merfeld, K.; Henkel, S. Technology-Enabled Personalization in Retail Stores: Understanding Drivers and Barriers. *J. Bus. Res.* **2021**, *123*, 140–155. [\[CrossRef\]](#)
4. Basu, S. Personalized Product Recommendations and Firm Performance. *Electron. Commer. Res. Appl.* **2021**, *48*, 101074. [\[CrossRef\]](#)

5. Ayvaz, D.; Aydogan, R.; Akcura, M.T.; Sensoy, M. Campaign Participation Prediction with Deep Learning. *Electron. Commer. Res. Appl.* **2021**, *48*, 101058. [[CrossRef](#)]
6. Zhang, S.; Liao, P.; Ye, H.-Q.; Zhou, Z. Dynamic Marketing Resource Allocation with Two-Stage Decisions. *J. Theor. Appl. Electron. Commer. Res.* **2022**, *17*, 327–344. [[CrossRef](#)]
7. Zeng, F.; Ye, Q.; Li, J.; Yang, Z. Does Self-Disclosure Matter? A Dynamic Two-Stage Perspective for the Personalization-Privacy Paradox. *J. Bus. Res.* **2021**, *124*, 667–675. [[CrossRef](#)]
8. Hufnagel, G.; Schwaiger, M.; Weritz, L. Seeking the Perfect Price: Consumer Responses to Personalized Price Discrimination in e-Commerce. *J. Bus. Res.* **2022**, *143*, 346–365. [[CrossRef](#)]
9. Shin, J.; Yu, J. Targeted Advertising and Consumer Inference. *Mark. Sci.* **2021**, *40*, 900–922. [[CrossRef](#)]
10. Zhang, X.; Yao, Y.; Zhang, J.; Fang, C. Post-Promotion Redemption, Exposure, and Spillover Effects of Electronic Coupons: An Empirical Analysis. *Prod. Oper. Manag.* **2023**, *32*, 603–617. [[CrossRef](#)]
11. von Mutius, B.; Huchzermeier, A. Customized Targeting Strategies for Category Coupons to Maximize CLV and Minimize Cost. *J. Retail.* **2021**, *97*, 764–779. [[CrossRef](#)]
12. Haupt, J.; Lessmann, S. Targeting Customers under Response-Dependent Costs. *Eur. J. Oper. Res.* **2022**, *297*, 369–379. [[CrossRef](#)]
13. Baier, D.; Stöcker, B. Profit Uplift Modeling for Direct Marketing Campaigns: Approaches and Applications for Online Shops. *J. Bus. Econ.* **2022**, *92*, 645–673. [[CrossRef](#)]
14. Gubela, R.M.; Lessmann, S. Uplift Modeling with Value-Driven Evaluation Metrics. *Decis. Support Syst.* **2021**, *150*, 113648. [[CrossRef](#)]
15. Lessmann, S.; Haupt, J.; Coussement, K.; De Bock, K.W. Targeting Customers for Profit: An Ensemble Learning Framework to Support Marketing Decision-Making. *Inf. Sci.* **2021**, *557*, 286–301. [[CrossRef](#)]
16. Zhan, B.; Liu, C.; Li, Y.; Wu, C. Weighted Doubly Robust Learning: An Uplift Modeling Technique for Estimating Mixed Treatments' Effect. *Decis. Support Syst.* **2024**, *176*, 114060. [[CrossRef](#)]
17. Pinheiro, P.; Cavique, L. A Machine Learning Framework for Uplift Modeling through Customer Segmentation. *Decis. Anal. J.* **2025**, *17*, 100639. [[CrossRef](#)]
18. Wang, J.; Tan, Y.; Jiang, B.; Wu, B.; Liu, W. Dynamic Marketing Uplift Modeling: A Symmetry-Preserving Framework Integrating Causal Forests with Deep Reinforcement Learning for Personalized Intervention Strategies. *Symmetry* **2025**, *17*, 610. [[CrossRef](#)]
19. Athey, S.; Wager, S. Policy Learning With Observational Data. *Econometrica* **2021**, *89*, 133–161. [[CrossRef](#)]
20. Nie, X.; Wager, S. Quasi-Oracle Estimation of Heterogeneous Treatment Effects. *Biometrika* **2021**, *108*, 299–319. [[CrossRef](#)]
21. Knaus, M.C.; Lechner, M.; Strittmatter, A. Machine Learning Estimation of Heterogeneous Causal Effects: Empirical Monte Carlo Evidence. *Econom. J.* **2021**, *24*, 134–161. [[CrossRef](#)]
22. Athey, S.; Imbens, G. Recursive Partitioning for Heterogeneous Causal Effects. *Proc. Natl. Acad. Sci. USA* **2016**, *113*, 7353–7360. [[CrossRef](#)] [[PubMed](#)]
23. Wager, S.; Athey, S. Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests. *J. Am. Stat. Assoc.* **2018**, *113*, 1228–1242. [[CrossRef](#)]
24. Athey, S.; Tibshirani, J.; Wager, S. Generalized Random Forests. *Ann. Stat.* **2019**, *47*, 1148–1178. [[CrossRef](#)]
25. Künzel, S.R.; Sekhon, J.S.; Bickel, P.J.; Yu, B. Metalearners for Estimating Heterogeneous Treatment Effects Using Machine Learning. *Proc. Natl. Acad. Sci. USA* **2019**, *116*, 4156–4165. [[CrossRef](#)] [[PubMed](#)]
26. Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; Robins, J. Double/Debiased Machine Learning for Treatment and Structural Parameters. *Econom. J.* **2018**, *21*, C1–C68. [[CrossRef](#)]
27. Elmachtoub, A.N.; Grigas, P. Smart “Predict, Then Optimize”. *Manag. Sci.* **2022**, *68*, 9–26. [[CrossRef](#)]
28. Kitagawa, T.; Tetenov, A. Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice. *Econometrica* **2018**, *86*, 591–616. [[CrossRef](#)]
29. Petrides, G.; Moldovan, D.; Coenen, L.; Guns, T.; Verbeke, W. Cost-Sensitive Learning for Profit-Driven Credit Scoring. *J. Oper. Res. Soc.* **2022**, *73*, 338–350. [[CrossRef](#)]
30. Yang, X.; Lai, B.; Tang, C. Experiential Product Promotions on E-Commerce Platform: From the Perspective of Consumer Cognition and Emotion. *Sage Open* **2023**, *13*, 21582440231153857. [[CrossRef](#)]
31. Wasilewski, A.; Kolaczek, G. One Size Does Not Fit All: Multivariant User Interface Personalization in E-Commerce. *IEEE Access* **2024**, *12*, 65570–65582. [[CrossRef](#)]
32. Goel, P.; Parayitam, S.; Sharma, A.; Rana, N.P.; Dwivedi, Y.K. A Moderated Mediation Model for E-Impulse Buying Tendency, Customer Satisfaction and Intention to Continue e-Shopping. *J. Bus. Res.* **2022**, *142*, 1–16. [[CrossRef](#)]
33. Liu, P.; Li, M.; Dai, D.; Guo, L. The Effects of Social Commerce Environmental Characteristics on Customers' Purchase Intentions: The Chain Mediating Effect of Customer-to-Customer Interaction and Customer-Perceived Value. *Electron. Commer. Res. Appl.* **2021**, *48*, 101073. [[CrossRef](#)]
34. Kanani, R.; Glavee-Geo, R. Breaking the Uncertainty Barrier in Social Commerce: The Relevance of Seller and Customer-Based Signals. *Electron. Commer. Res. Appl.* **2021**, *48*, 101059. [[CrossRef](#)]

35. Lin, J.; Li, T.; Guo, J. Factors Influencing Consumers' Continuous Purchase Intention on Fresh Food e-Commerce Platforms: An Organic Foods-Centric Empirical Investigation. *Electron. Commer. Res. Appl.* **2021**, *50*, 101103. [[CrossRef](#)]
36. Cremer, S.; Loebbecke, C. Selling Goods on E-Commerce Platforms: The Impact of Scarcity Messages. *Electron. Commer. Res. Appl.* **2021**, *47*, 101039. [[CrossRef](#)]
37. Ruz-Mendoza, M.Á.; Trifu, A.; Cambra-Fierro, J.; Melero-Polo, I. Standardized vs. Customized Firm-Initiated Interactions: Their Effect on Customer Gratitude and Performance in a B2B Context. *J. Bus. Res.* **2021**, *133*, 341–353. [[CrossRef](#)]
38. Treiblmaier, H.; Sillaber, C. The Impact of Blockchain on E-Commerce: A Framework for Salient Research Topics. *Electron. Commer. Res. Appl.* **2021**, *48*, 101054. [[CrossRef](#)]
39. Li, Q.; Zhu, C.; Shi, T. Augmented Reality Advertising in an E-Commerce Model with Competition. *Electron. Commer. Res. Appl.* **2021**, *49*, 101092. [[CrossRef](#)]
40. Wang, L.; Wang, Z.; Wang, X.; Zhao, Y. Explaining Consumer Implementation Intentions in Mobile Shopping with SEM and fsQCA: Roles of Visual and Technical Perceptions. *Electron. Commer. Res. Appl.* **2021**, *49*, 101080. [[CrossRef](#)]
41. Li, C.-Y.; Fang, Y.-H. Toward Better Purchase Decision-Performance: Linking Person-Environment Fit to Explorative and Exploitative Use of Branded Applications. *Electron. Commer. Res. Appl.* **2021**, *48*, 101063. [[CrossRef](#)]
42. Guo, W.; Luo, Q. Investigating the Impact of Intelligent Personal Assistants on the Purchase Intentions of Generation Z Consumers: The Moderating Role of Brand Credibility. *J. Retail. Consum. Serv.* **2023**, *73*, 103353. [[CrossRef](#)]
43. De Biasio, A.; Navarin, N.; Jannach, D. Economic Recommender Systems—A Systematic Review. *Electron. Commer. Res. Appl.* **2024**, *63*, 101352. [[CrossRef](#)]
44. Jin, Z.; Ye, F.; Nedjah, N.; Zhang, X. A Comparative Study of Various Recommendation Algorithms Based on E-Commerce Big Data. *Electron. Commer. Res. Appl.* **2024**, *68*, 101461. [[CrossRef](#)]
45. Koronaki, E.; Vlachvei, A.; Panopoulos, A. Managing the Online Customer Experience and Subsequent Consumer Responses across the Customer Journey: A Review and Future Research Agenda. *Electron. Commer. Res. Appl.* **2023**, *58*, 101242. [[CrossRef](#)]
46. Chen, Q.; Wang, Y.; Gong, Y.; Liu, S. Ripping off Regular Consumers? The Antecedents and Consequences of Consumers' Perceptions of e-Commerce Platforms' Digital Power Abuse. *J. Bus. Res.* **2023**, *166*, 114123. [[CrossRef](#)]
47. Bertsimas, D.; Kallus, N. From Predictive to Prescriptive Analytics. *Manag. Sci.* **2020**, *66*, 1025–1044. [[CrossRef](#)]
48. Apley, D.W.; Zhu, J. Visualizing the Effects of Predictor Variables in Black Box Supervised Learning Models. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **2020**, *82*, 1059–1086. [[CrossRef](#)]
49. Lundberg, S.M.; Erion, G.; Chen, H.; DeGrave, A.; Prutkin, J.M.; Nair, B.; Katz, R.; Himmelfarb, J.; Bansal, N.; Lee, S.-I. From Local Explanations to Global Understanding with Explainable AI for Trees. *Nat. Mach. Intell.* **2020**, *2*, 56–67. [[CrossRef](#)] [[PubMed](#)]
50. Gubela, R.M.; Lessmann, S.; Jaroszewicz, S. Response Transformation and Profit Decomposition for Revenue Uplift Modeling. *Eur. J. Oper. Res.* **2020**, *283*, 647–661. [[CrossRef](#)]
51. Mulhern, F.J. Customer Profitability Analysis: Measurement, Concentration, and Research Directions. *J. Interact. Mark.* **1999**, *13*, 25–40. [[CrossRef](#)]
52. Venkatesan, R.; Kumar, V. A Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy. *J. Mark.* **2004**, *68*, 106–125. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.